Local Public Goods and Property Tax Compliance: Experimental Evidence from Street Pavement*

Manuel Fernández, Marco Gonzalez-Navarro, Climent Quintana-Domeque

August 15, 2025

Abstract

Developing countries may face a self-reinforcing cycle where weak tax compliance limits public goods provision, which in turn erodes citizens' incentives to pay taxes. This study tests whether improved local infrastructure can disrupt this cycle, using a randomized experiment involving first-time street paving in Acayucan, Mexico. Of 56 eligible street projects in poor neighborhoods, 28 were randomly selected for paving. A theoretical model informs our analysis, emphasizing two mechanisms: belief updating about government efficiency and reciprocity from direct benefits. We derive three testable implications: (1) belief updating should occur through exposure to street paving anywhere in the network; (2) compliance should increase among those exposed to this broader infrastructure; and (3) reciprocity should boost compliance among property owners directly benefiting from paving. Survey data support the belief-updating mechanism. Among residents initially dissatisfied with the local government, exposure to paving reduced dissatisfaction by 7.9 percentage points (ITT) and 8.8 points (LATE). No effect was observed among those initially satisfied. Administrative tax records further show that a one standard deviation increase in exposure to paving en route to downtown increased tax compliance by 1.5 percentage points (ITT) and 2.6 points (LATE)— the latter representing a 3% rise in baseline compliance. For the reciprocity mechanism, residents living directly on paved streets increased compliance by 3.2 percentage points (ITT) and 4.8 points (LATE), the latter representing a 5.5% rise from baseline. A back-of-the-envelope estimate suggests belief updating generated four times more tax revenue than direct reciprocity.

JEL Classification Codes: C93, H26, H41, H54, O12.
Keywords: taxpayer behavior, roads, infrastructure, belief updating, reciprocity, government efficiency, public investment, satisfaction with local government.

*Fernández: Universidad de los Andes & IZA, man-fern@uniandes.edu.co; Gonzalez-Navarro: University of California-Berkeley, BREAD, CEGA & J-PAL, marcog@berkeley.edu; Quintana-Domeque: University of Exeter & IZA, c.quintana-domeque@exeter.ac.uk. Daniel Lasso Jaramillo provided outstanding research assistance. We thank participants at Doshisha University, Hitotsubashi University, Institut d'Anàlisi Econòmica (IAE-CSIC), Stanford University, Universidad de Granada, Universidad del Rosario, Universidad de los Andes, Universitat Pompeu Fabra, University of Bristol, University of Exeter, University of Oxford, University of Tokyo, Catalan Economics Society Winter Workshop 2024, CEGA Research Retreat (R^2) at University of California Berkeley 2024, CEPR Development Economics Annual Symposium 2024, CSAE Annual Conference, Econometric Society World Congress 2023, Royal Economic Society Meetings 2025, and the Zurich Conference on Public Finance in Developing Countries 2024. The study protocol was approved by Princeton University Institutional Review Panel (Protocol 3104) and the University of Toronto Office of Research Ethics (Protocol 28782). Financial support from the Lincoln Institute of Land Policy is gratefully acknowledged. Any errors contained in the paper are our own.

1 Introduction

Governments depend on taxes to provide services and public goods. Yet, in many developing countries, there are concerns that low tax compliance and insufficient government funding create a negative feedback loop: citizens may resist paying taxes when public goods are absent or of low quality, and governments, in turn, struggle to provide these goods without adequate revenue. This cycle may be partially responsible for limited development outcomes worldwide.

In this paper, we examine whether a shock to public good provision enhances tax compliance. Anecdotal and correlational evidence suggests that inadequate public services can reduce citizens' incentives to pay taxes in both developed and developing countries. For instance, in 2010, a Detroit News investigation highlighted the frustration of property owners like Fred Philips, who asked, "Why pay taxes when they aren't supplying services?" Similarly, Afrobarometer data (2019–2021) shows that nearly half of surveyed Africans are willing to pay more taxes for better services. However, such evidence is merely suggestive, as governments' decisions to provide public goods often correlate with other factors influencing tax compliance (Kresch et al. 2023).

To address this causal identification challenge, we conducted a randomized experiment involving first-time road paving in underserved neighborhoods of a medium-sized Mexican city. Between 2007 and 2012, the city expanded its paved street network, randomly selecting half of 56 proposed asphalting projects for implementation (Gonzalez-Navarro and Quintana-Domeque 2016). Using administrative property tax records (2005–2012), we examine whether road paving improved compliance. Since the city's billing system prevents misreporting (Kleven et al. 2011), we can focus on the extensive margin of compliance with property tax payments as our main outcome variable.

We model the decision to pay taxes by building on prior research and the context of our study. First, we exclude penalty and punishment considerations (Becker 1968) because property tax enforcement in our setting is virtually absent.[1] Second, we incorporate non-pecuniary, tax morale, factors (Luttmer and Singhal 2014; Pomeranz and Vila-Belda 2019), emphasizing the role of reciprocity (Besley 2020) when the individual directly benefits from public good provision—by getting pave-

1. Tax compliance in economics is often analyzed within the framework of Becker's crime theory (Becker 1968), which predicts compliance based on a cost-benefit analysis of detection likelihood, penalties, and individual risk aversion (Allingham and Sandmo 1972). However, actual compliance frequently exceeds predictions based solely on these factors (Andreoni, Erard, and Feinstein 1998; Luttmer and Singhal 2014; Dwenger et al. 2016), a pattern observed in our study, where enforcement mechanisms are minimal.

ment in the street their property lines.[2] This acts as decreasing the effective tax rate when the government delivers public goods.[3] Third, we examine how public goods provision signals the local government's expenditure efficiency. This effect, which we believe is novel in the literature, should arise regardless of a citizen's proximity to newly paved streets, as it depends solely on observing government actions. Empirically, our model allows us to relax the Stable Unit Treatment Value Assumption by predicting two effects: (1) an exposure effect, where compliance increases from observing public goods delivery anywhere in the city, and (2) perceived direct benefit effects, which depend on proximity to the intervention. This framework captures nuanced behavioral responses to public goods provision.

We then proceed to the empirical analysis using both survey data and administrative property tax records. In our main specification, we define exposure as the fraction of nodes in the street network along the shortest path from a property to the city center that are located on, or intersect with, streets paved by the intervention. Three key findings emerge from our analysis. First, and consistent with the belief-updating mechanism, a one-standard deviation increase in exposure reduced dissatisfaction with the local government by 7.9 percentage points (ITT) and 8.8 percentage points (LATE), an effect which is driven by those initially dissatisfied. This finding aligns with our theoretical framework as long as satisfaction with the local government serves as a proxy for beliefs about local government efficiency in the spirit of Wang and Kim (2024). Moving to the findings on tax compliance, and consistent with the belief-updating mechanism, our second result is that a one-standard-deviation increase in exposure led to a 1.5 percentage point (ITT) and 2.6 percentage point (LATE) increase in compliance. Our third finding, supporting the reciprocity due to perceived direct benefits mechanism, is that property owners adjacent to paved streets experienced a 3.2 percentage points (ITT) and 4.8 percentage points (LATE) increase in property tax compliance.

We then perform a series of robustness checks. First, we show that the exposure mechanism does not affect satisfaction in domains beyond those related to government performance, nor does it appear to capture direct material benefits. While having a paved street in front of a property increases its value, exposure alone has no effect on property values or perceived travel time to the city center. These support our hypothesis that the exposure effect operates through the belief updating

2. Individuals reciprocate government actions, such as the provision of public goods, based on a sense of fairness or altruism rather than external incentives or strategic considerations (Sobel 2005; Besley 2020; Besley, Jensen, and Persson 2023). For example, individuals may view tax payment as fulfilling an implicit social contract with the state, contributing to the collective resources necessary for its effective functioning (Musgrave 1992; Besley 2020), even when evasion appears materially advantageous.

3. This sense of obligation can depend on the perceived fairness of government actions and the quality of services provided. In environments plagued by corruption, individuals may feel justified in withdrawing their cooperation and refusing to comply with tax obligations (Feld and Frey 2007; Besley and Persson 2014; Stantcheva 2021; Ajzenman et al. 2024).

channel. Second, to address the concern that exposure may reflect other direct benefits from traveling on paved streets—even when these streets are not adjacent to the property—we assess the robustness of our results using an alternative exposure measure. This measure excludes nodes along the shortest path and includes only those that intersect streets paved by the intervention, thereby capturing visual exposure independent of actual use. We find that results are statistically indistinguishable between the two metrics. Third, we demonstrate that our results remain robust when using alternative destinations for our measure of exposure, as well as a harmonized index that combines all of them. Fourth, accounting for non-random exposure by recentering exposure measures, following Borusyak and Hull (2023), yields very similar estimates. Finally, since our analysis on tax compliance excludes properties that change ownership during the study period—resulting in attrition—we address concerns about differential attrition by treatment status using the approach of Ghanem, Hirshleifer, and Ortiz-Beccera (2023), and find no evidence that selective attrition biases our results.

We conclude our empirical investigation with a back-of-the-envelope calculation of the fiscal consequences of the street paving intervention. This assessment considers not only the additional revenue from property owners whose streets were directly paved but also from those indirectly exposed to the paving of others' streets. While the absolute levels of tax revenue are understandably small due to the modest size of the property tax, the relative effects are significant and provide valuable insights. Specifically, we document substantial fiscal spillovers: the ratio of increases in property tax revenue obtained from those *not* residing adjacent to newly paved streets but who are exposed to the pavement program, is four times greater than the increase in tax revenue from citizens whose properties directly abut the newly paved streets. This finding underscores the broader fiscal benefits of infrastructure improvements beyond the directly treated areas.

Our paper contributes to three main strands of the literature: the literature on state capacity and taxation, which emphasizes the ability to effectively collect taxes as essential for the proper functioning of modern states and as a key marker of state capacity (Levi 1988; Besley 2020; Besley, Jensen, and Persson 2023); the literature on the role of public perception of government spending in incentivizing tax compliance, which highlights how government incompetence or corruption (Besley and Persson 2014; Timmons and Garfias 2015) can deter citizens from complying with taxes; and the broader literature on non-pecuniary factors influencing tax compliance, often grouped under the umbrella term of tax morale (Luttmer and Singhal 2014; Pomeranz and Vila-Belda 2019; Slemrod 2019).[4]

4. These factors include motivations such as honesty, fulfillment of civic duties, altruism, perceived fairness, valuation of government use of tax revenues, adherence to social norms, and peer effects. Despite their importance, these motivations are the hardest to measure and study empirically and, therefore, remain the least well-understood.

Extant studies on non-pecuniary motives for tax compliance rely on interventions in which taxpayers receive letters with randomly varying messages designed to prime specific dimensions of tax morale.[5] Results from these interventions tend to be mixed or null, with messages emphasizing traditional deterrence mechanisms proving more effective at inducing compliance than those appealing to intrinsic motivations (Kleven et al. 2011; Castro and Scartascini 2015; Dwenger et al. 2016; Bérgolo et al. 2017; Meiselman 2018; Antinyan and Asatryan 2024). In a recent review, Slemrod (2019) concludes: "In sum, a plethora of studies have failed to find evidence that appeals to tax morale, defined broadly, affect taxpayer behavior in the short run when delivered via a one-time mailing." As Luttmer and Singhal (2014) note, however, a null effect may result either from tax morale being irrelevant or from these messages failing to effectively alter it (Benabou and Tirole 2011).

Our study diverges from traditional letter-based interventions and informational approaches aimed at increasing the salience of public service provision (Montenbruck 2023) by focusing on the tangible impact of actual changes in public goods provision on tax compliance. Moreover, by analyzing a well-defined type of public good and leveraging granular, street-level project data, we overcome the limitations of prior research that examined larger areas and combined multiple interventions, yielding statistically insignificant results (Brockmeyer, Garfias, and Suárez Serrato 2024). In this way, our study offers new insights into how public goods delivery can break the cycle of low tax compliance and inadequate public services.

A closely related work to ours is Carrillo, Castro, and Scartascini (2021), who evaluate the effectiveness of announcing that citizens who are compliant with their property taxes will enter a lottery for new sidewalks. While they find that the announcement of the scheme induces very few additional taxpayers to comply in order to participate in the lottery, they do find that those who end up winning the sidewalks (as well as their neighbors) are less likely to be delinquent with their taxes in the future. We build on that paper in that we are able to better identify the reciprocity mechanism because the street pavement in this study was provided unconditional on tax compliance status, versus potentially confounding it with the strategic decision to comply in order to participate in future lotteries or due to individuals seeking public recognition as lottery winners.

The paper proceeds as follows. Section 2 presents a simple model to examine

5. For studies emphasizing the moral duty to pay taxes, see Hallsworth et al. (2017), Meiselman (2018), and Bott et al. (2020). For studies informing taxpayers about public goods recently financed through taxes, see Castro and Scartascini (2015), Chirico et al. (2016), Bérgolo et al. (2017), Hallsworth et al. (2017), Meiselman (2018), Bott et al. (2020), De Neve et al. (2021), Giaccobasso et al. (2022), Wilson and Rosenzweig (2023), and Sandholtz and Vicente (2024). For studies focusing on the compliance behavior of other citizens or fostering social recognition for compliance, see Castro and Scartascini (2015), Hallsworth et al. (2017), Perez-Truglia and Troiano (2018), and Drago, Mengel, and Traxler (2020). A meta-analysis of nudging interventions on tax compliance, based on studies that implement randomized control trials, is provided by Antinyan and Asatryan (2024).

how the provision of a local public good can influence property tax compliance, which delivers three testable predictions. Section 3 describes the institutional context and experimental design, defines the measure of exposure and the data sources used in our study. Section 4 discusses the econometric strategy to link the predictions of the model with the empirical analysis. Section 5 presents the main empirical findings of our paper. Section 6 assesses the validity of the exposure measure and contains additional robustness checks. Section 7 offers an analysis of the fiscal spillover effects of street paving. Section 8 concludes.

2 A Model of Property Tax Compliance

We propose a static tax compliance model that emphasizes the central role of perceived government expenditure inefficiency in individuals' compliance decisions. These perceptions are modeled as being directly influenced by observable government actions—specifically, expenditures on public goods in our context.

Setup. Consider a continuum of property owners of unit size, indexed by p. Each owner must decide whether to comply with property tax obligations ($C_p = 1$) or evade them ($C_p = 0$), where $\tau_p > 0$ denotes their individual tax liability. The revenue generated from property taxes is allocated towards the provision of local public goods. A central concept introduced in this framework is that taxpayers form an expectation regarding the fraction $\theta \in [0, 1]$ of total public expenditure G that is lost to inefficiencies, including corruption and mismanagement. This parameter, θ, directly affects the availability of funds for public goods and, as such, serves as a measure of government inefficiency. While total expenditure G is assumed to be common knowledge, the extent of inefficiency (θ) remains uncertain. This assumption is particularly relevant in institutional contexts where mayors are not allowed immediate reelection, thereby imposing persistent uncertainty regarding the inefficiency of government spending. To align closely with the empirical setting under consideration, we further assume that public good provision is limited to newly paved streets, with all non-wasted expenditure directed to this purpose.

Preferences. We assume citizen preferences can be represented by the following utility function u:

$$u(C_p) = y_p - \tau_p C_p + \lambda\Big(1 - \theta_p(E_p)\Big)GC_p + \alpha'_p B_p C_p, \tag{2.1}$$

where y_p denotes gross income, so that $y_p - \tau_p C_p$ reflects the fact that tax compliance reduces disposable income, and hence utility. Note that the tax bill τ_p is only paid when home owners choose to comply ($C_p = 1$).

We depart from the traditional Beckerian framework, which emphasizes fines as a primary deterrent against tax noncompliance, for several reasons. In most developing countries, including the context of our study, nonpayment of property taxes does not result in significant legal repercussions, such as home repossession or other forms of enforcement aimed at recovering delinquent taxes. Local governments maintain records of unpaid taxes but rarely, if ever, initiate legal proceedings against noncompliant property owners. The primary legal mechanism encouraging compliance operates at the time of property transactions: before a property can be registered under a new owner's name, the government requires that any outstanding tax arrears and penalties from the previous few years be settled. Given the absence of stringent legal enforcement, a key question emerges: why do we observe positive rates of compliance with property tax obligations?

To address this puzzle, we draw on the concept of tax morale, a notion widely explored in the contemporary literature on tax compliance (Luttmer and Singhal 2014). In contexts where enforcement is weak, tax morale can explain why individuals may still choose to pay property taxes due to intrinsic, non-pecuniary motivations. For example, taxpayers may perceive compliance as an obligation under an implicit social contract with the state, contributing to the collective financial resources necessary for effective governance (Musgrave 1992; Besley 2020). Additionally, some individuals may be willing to pay taxes in exchange for public goods and services, even when the pecuniary benefits of noncompliance exceed those of payment.

Building on Sobel (2005), Besley (2020), and Besley, Jensen, and Persson (2023), we formalize this idea by modeling taxpayer preferences as incorporating an intrinsic utility component, denoted by $\lambda > 0$, which reflects the extent to which individuals value expenditures on public goods. A higher λ implies that an individual derives greater utility from tax payments that contribute to public good provision, serving as a measure of "civic mindedness". However, this civic-mindedness only applies to those who choose to comply ($C_p = 1$).

The key innovation in our model lies in incorporating taxpayers' expectations about government inefficiency, represented by $\theta_p(E_p)$. Specifically, the term $(1 - \theta_p(E_p))$ captures the notion that individuals value only the portion of their tax contributions that they believe is effectively allocated to public goods, adjusting for perceived waste or mismanagement. Consequently, an individual's sense of obligation to comply is contingent upon their perception of governmental inefficiency. This framework aligns with empirical findings suggesting that in environments characterized by widespread corruption, individuals may feel justified in withdrawing tax payments as a rational response to perceived government failures (Feld and Frey 2007; Besley and Persson 2014; Stantcheva 2021; Ajzenman et al. 2024).

In summary, compliance with property tax obligations is more likely among individuals who exhibit stronger pro-public good preferences (higher λ), and hold more favorable perceptions of government expenditure inefficiency (higher $(1-\theta_p(E_p))$).

To anticipate the experimental variation we exploit in our empirical strategy, we posit that individuals' expectations regarding θ are influenced by their exposure to public goods, specifically the observation of newly paved streets ($E_p \geq 0$). In our framework, "exposure" refers to receiving a positive signal regarding the efficiency of government expenditures on public goods, which shapes individuals' beliefs. A more precise definition of our exposure measure is provided below.

The final component of equation (2.1), $\alpha_p' B_p C_p$, captures the perceived direct benefits of public good provision as a motivation for tax compliance. It is by now well established that the provision of local public goods increases property values (e.g. Gibbons and Machin (2005)). Accordingly, the direct benefits derived from public good provision, represented by B_p, serve as an additional motivation for taxpayers to comply. In the specific case of street pavement, Gonzalez-Navarro and Quintana-Domeque (2016) document property value increases of about 30% following infrastructure improvements, consistent with the direct benefits of paving streets. Comparable estimates are found in Sorin (2025) for Uganda. Our model incorporates this empirically significant relationship, allowing us to account for direct benefits, such as the capitalization of public good improvements into property values, as a potential determinant of compliance behavior. Once again, this reciprocity motive manifests through increased utility for individuals who choose to comply ($C_p = 1$).

In this expression, B_p can be a vector of indicator variables, where each element indicates whether a street within a specific distance from the property was paved, and α_p is the corresponding vector of utility weights. This spatial structure can capture the intuition that the benefits of paved streets decline with distance from the property. While receiving direct benefits from public goods is not contingent on complying with property tax payments, the term captures any motivation individuals experience that compels them to comply when such benefits are realized, so we multiply by the term C_p. Importantly, while receiving direct benefits necessarily implies exposure to paved streets, the converse is not true: property owners may observe paved streets without deriving direct benefits from them.

Individual Compliance. Given our setup, the compliance decision is based on a threshold condition. It is optimal to comply if:

$$\theta_p(E_p) \leq 1 - \frac{\tau_p}{\lambda G} + \frac{\alpha_p' B_p}{\lambda G}. \tag{2.2}$$

Rewritten as $\tau_p \leq \lambda(1-\theta_p(E_p))G + \alpha_p' B_p$, it is easy to see that the condition states that the likelihood of property tax compliance increases with individuals' civic-mindedness (λ), their perceptions of government efficiency $(1-\theta_p(E_p))$ and their perceived direct benefits of public good provision ($\alpha_p' B_p$). Conversely, compliance decreases as the tax liability (τ_p) rises. This expression offers an explanation for how compliance occurs despite the lack of enforcement mechanisms such as inspections or credible penalties (Andreoni, Erard, and Feinstein 1998; Luttmer and Singhal 2014; Dwenger et al. 2016). It underscores the role of non-pecuniary motivations in driving compliance while also highlighting how these intrinsic incentives are reinforced by the expectation that taxes are spent in a socially valuable manner.

Exposure. We now formalize our measure of exposure to newly paved streets. Suppose the city has a total of S unpaved streets. Normalizing the cost of paving a new street to one implies that the total number of streets paved is $(1-\theta)G \ll S$, where the inequality "much less than" reflects the fact that, in our context, the probability that any unpaved street is paved is low.[6]

Consistent with our experimental design, we assume that all unpaved streets have an equal and independent probability of being selected for paving. For any unpaved street s, we define an indicator variable d_s that takes a value of one if the street gets paved. The probability of this occurring is given by:

$$\pi \equiv Pr(d_s = 1) = \frac{(1-\theta)G}{S}. \tag{2.3}$$

The probability decreases with θ, indicating fewer paved streets in the presence of greater government inefficiency. Conversely, it increases with G as more resources are available for asphalting projects.

Next, consider a property owner who regularly travels along S_p unpaved streets. We focus on daily transit routes as these represent the subset of the city's street network that individuals consistently observe. Let X_p denote the total length of the owner's commute, defined as the sum of all streets (both paved and unpaved) along their route. For each street s, let the indicator variable $d_{p,s}$ equal one if the street is both newly paved ($d_s = 1$) and part of property owner p's route. We then define the exposure of property owner p to newly paved streets (E_p) as the fraction of their total trajectory that obtains pavement:

$$E_p = \frac{1}{X_p} \sum_s d_{p,s}. \tag{2.4}$$

Note that the number of newly paved streets observed by property owner p, given by the product $X_p E_p$, follows a Binomial distribution $X_p E_p \mid \theta \sim \text{Binomial}(S_p, \pi)$. This formalization captures how exposure to newly paved streets varies across property owners based on their location in the city and the stochastic nature of the paving process.

6. In Table A.1, we document that the percentage of unpaved streets in Acayucan that obtain pavement during the period of analysis averages 1.3% per year.

Beliefs. Property owners have prior beliefs about the unknown value of θ, which follow a Beta distribution $\theta_p \sim \text{Beta}(\kappa_{p,1}, \kappa_{p,2})$. The shape parameters $\kappa_{p,1}, \kappa_{p,2} > 0$ determine both the mean of the prior, denoted by θ_p^*, and its variance, which decreases with the concentration parameter $\kappa_p \equiv \kappa_{p,1} + \kappa_{p,2}$. This specification is flexible, as the shape parameters are individual-specific and can generate a wide range of prior belief distributions.[7]

Once public expenditure is realized, individuals gain access to new information that enables them to update their beliefs regarding government inefficiency. Specifically, the observation of newly paved streets serves as an indicator of a less inefficient government. We assume property owners update their beliefs using Bayes' rule and that the ratio G/S is small (i.e., $\pi \ll 1$). Exploiting the properties of conjugate distributions, we show in Appendix A.1 that the posterior expectation of θ, conditional on exposure and X_p, can be approximated by:

$$\theta_p(E_p) \approx \theta_p^* - \frac{\theta_p^* X_p}{\kappa_p} E_p \equiv \theta_p^* - \delta_p E_p, \tag{2.5}$$

where $\delta_p \equiv \frac{\theta_p^* X_p}{\kappa_p} \geq 0$ and the linear approximation error is of order $O((X_p E_p/\kappa_p)^2)$.

Equation (2.5) fully describes how citizens update their beliefs about government inefficiency. First, an increased exposure to newly paved streets lowers beliefs about government inefficiency (θ_p):

$$\frac{\partial \theta_p(E_p)}{\partial E_p} = -\delta_p \leq 0. \tag{2.6}$$

Since lower perceived inefficiency increases the likelihood of tax compliance, this belief-updating mechanism constitutes one channel through which public good provision can enhance tax compliance.[8]

Second, when individuals are not exposed to newly paved streets ($E_p = 0$), their posterior beliefs remain unchanged ($\theta_p(0) = \theta_p^*$). This persistence of prior beliefs reflects a low baseline expectation of public good provision—the absence of new infrastructure is viewed as the *status quo* rather than as information about government inefficiency.

7. For example, they can generate distributions that are uniform, bell-shaped, skewed, or monotonically increasing or decreasing.

8. Note that the magnitude of belief revision (δ_p) increases with more pessimistic prior beliefs (θ_p^*), as individuals with larger θ_p^* assign a lower ex-ante probability of being exposed to a newly paved street. In contrast, δ_p decreases with the concentration parameter (κ_p), which is inversely related to the prior variance. This implies that individuals with more uncertain initial beliefs are more responsive to new information, while those with highly concentrated priors require stronger evidence to update their views. This prediction helps explain the mixed results from tax compliance interventions that use informational letters, discussed in Section 1. As Luttmer and Singhal (2014) argue, messages in a letter may provide insufficient evidence to shift entrenched beliefs.

Direct Benefits. We now specify how direct benefits from public good provision enter the utility function. To capture this spatial dimension, we divide each property owner's commuting trajectory into distinct distance rings. For any street s in property owner p's regular commute, let $L(p, s)$ denote the distance between the property and the street, where $L(p, s) = 0$ for streets directly adjacent to the property. We define an indicator variable for each distance ring j:

$$b_{p,j} = \mathbf{1}\Big[\exists s : d_{p,s} = 1 \ \text{ and } \ L(p, s) \in [a_{j-1}, a_j)\Big], \tag{2.7}$$

where $\mathbf{1}[\cdot]$ is an indicator function that takes the value one if the condition inside holds, $[a_{j-1}, a_j)$ defines the boundaries of ring $j \in \{0, 1, \ldots, J\}$ and a_{-1} is set to zero. This indicator equals one if at least one newly paved street exists within a corresponding distance range.

The total direct benefits perceived by property owner p are then given by:

$$\alpha_p' B_p = \sum_j \alpha_{p,j} b_{p,j}, \tag{2.8}$$

where $\alpha_{p,j}$ captures the marginal benefit of having a paved street in ring j. These coefficients are assumed to be a function of distance. In particular, we expect them to decrease with distance, reflecting the diminishing impact of street improvements farther from the property. We will present empirical evidence not only of this negative relationship, but also that only streets paved directly in front of a property ($j = 0$) generate direct benefits large enough to influence compliance. This pattern is unsurprising in our setting, which consists of residential streets in peripheral neighborhoods, as opposed to trunk roads whose main beneficiaries are passers-by.

Updated Compliance Threshold. Substituting equations (2.5) and (2.8) into (2.2), the threshold condition for compliance becomes:

$$\theta_p^* \leq \eta_p + \delta_p E_p + \sum_j \gamma_{p,j} b_{p,j}. \tag{2.9}$$

Here, $\eta_p = 1 - \frac{\tau_p}{\lambda G}$ represents the baseline compliance threshold in the absence of exposure and direct benefits, and $\gamma_{p,j} = \frac{\alpha_{p,j}}{\lambda G}$. Compared to equation (2.2), this formulation explicitly incorporates the prior expectation of government inefficiency (θ_p^*), the realized level of exposure (E_p) and the perceived direct benefits from public good provision, captured through $\sum_j \gamma_{p,j} b_{p,j} = \sum_j \frac{\alpha_{p,j}}{\lambda G} b_{p,j}$.[9]

Binary Response Framework. The decision to comply with property tax obligations can be represented as a binary response model:

$$C_p = \mathbf{1}\left[\theta_p^* \leq \eta_p + \delta_p E_p + \sum_j \gamma_{p,j} b_{p,j}\right]. \tag{2.11}$$

This formulation emphasizes that compliance occurs, $C_p = 1$, if and only if the perceived government inefficiency θ_p^* is sufficiently low, given the individual's exposure to newly paved streets, E_p, and the direct benefits received, $b_{p,j}$.

Probability of Compliance. If we define the latent index:

$$y_p^* = \eta_p + \delta_p E_p + \sum_j \gamma_{p,j} b_{p,j},$$

and given that the perceived inefficiency θ_p^* is a continuous random variable with cumulative distribution function F,[10] the probability of compliance is:

$$\mathbb{P}(C_p = 1 \mid \eta_p, E_p, \mathbf{b}_p) = \mathbb{P}(\theta_p^* \leq y_p^*) = F(y_p^*).$$

The marginal effect of exposure to paved streets E_p on the probability of compliance is:

$$\frac{\partial \mathbb{P}(C_p = 1 \mid \eta_p, E_p, \mathbf{b}_p)}{\partial E_p} = f(y_p^*) \cdot \delta_p,$$

where $\mathbf{b}_p$ is the vector of direct benefits, and the effect of receiving benefit $b_{p,k}$ is:

$$\begin{aligned}&\mathbb{P}(C_p = 1 \mid \eta_p, E_p, b_{p,k} = 1) - \mathbb{P}(C_p = 1 \mid \eta_p, E_p, b_{p,k} = 0)\\ &\quad = F\big(y_p^* + \gamma_{p,k}(1 - b_{p,k})\big) - F\big(y_p^* - \gamma_{p,k} b_{p,k}\big),\end{aligned}$$

where f is the probability density function associated with the distribution F.

9. Note that the assumption of exposure effects being independent of distance to the plot and perceived direct benefits varying with distance is crucial for distinguishing the effects of exposure from those of direct benefits. Suppose that both effects were distance-independent, so that the benefit parameters $\alpha_{p,j}$ are constant and equal to α_p across all distance rings and each newly paved street belongs to a different ring from property p's perspective. Then $\sum_j \gamma_{p,j} b_{p,j} = \gamma_p \sum_s d_{p,s}$. Using the definition of exposure from equation (2.4), this simplifies to $\gamma_p X_p E_p$. The compliance condition then becomes:

$$\theta_p^* \leq \eta_p + (\delta_p + \gamma_p X_p) E_p = \eta_p + (\delta_p + \tilde{\gamma}_p) E_p = \eta_p + \tilde{\delta}_p E_p, \tag{2.10}$$

where $\delta_p \equiv \frac{\theta_p^* X_p}{\kappa_p}$ and $\tilde{\gamma}_p \equiv \gamma_p X_p$, so that $\tilde{\delta}_p$ would reflect a combination of effects of exposure and perceived direct benefits.

10. We have assumed that F is the Beta distribution: $\theta_p \sim \text{Beta}(\kappa_{p,1}, \kappa_{p,2})$.

Testable Implications. Our framework yields three testable predictions, which we evaluate empirically using our experimental setup described in the next section. The predictions are:

1. **Belief Updating:** An increased exposure to newly paved streets lowers beliefs about government inefficiency (equation (2.6)):

$$\frac{\partial \theta_p(E_p)}{\partial E_p} = -\delta_p \leq 0.$$

2. **Compliance via Updated Beliefs:** Property tax compliance increases with greater exposure to street paving anywhere on the network, as the observation of public goods provision reduces individuals' beliefs about government inefficiency (equations (2.5) and (2.9)):

$$\frac{\partial \mathbb{P}(C_p = 1 \mid \eta_p, E_p, \mathbf{b}_p)}{\partial E_p} = f(y_p^*) \cdot \delta_p \geq 0.$$

3. **Compliance via Direct Benefits:** Direct benefits from paved streets increase compliance, with effects that depend on the distance between the property and the improved roads (equations (2.8) and (2.9)):

$$\begin{aligned} &\mathbb{P}(C_p = 1 \mid \eta_p, E_p, b_{p,k} = 1) - \mathbb{P}(C_p = 1 \mid \eta_p, E_p, b_{p,k} = 0) \\ &= F\big(y_p^* + \gamma_{p,k}(1 - b_{p,k})\big) - F\big(y_p^* - \gamma_{p,k} b_{p,k}\big) \geq 0. \end{aligned}$$

3 Context, Experiment, Data and Measurement

3.1 Context and Experiment

Acayucan is a mid-sized city in southeastern Veracruz, Mexico, and serves as the seat of its namesake municipality. As of 2010, its population was approximately 85,000. The city features a paved central core surrounded by peripheral neighborhoods where street paving is gradually introduced—often long after residential construction and occupation. This sequencing is common across urban areas in Mexico and Latin America.

Responsibility for urban infrastructure in Mexico lies primarily with municipal governments, each elected to a three-year term and granted wide discretion over budget allocation. As a result, citizens must frequently update their beliefs about the efficiency of local public spending. In our setting, however, the same political family governed across three consecutive terms from 2005-2013. We therefore treat this as a single, uninterrupted tenure, abstracting from potential political dynamics across administrations.

The policy we study involves the first-time paving of residential, non-arterial streets in peripheral neighborhoods of Acayucan.[11] These projects vary in width from 8 to 15 meters and in length from 210 to 770 meters, typically accommodating two traffic lanes and one or two parking lanes. A typical street block includes approximately 16 adjacent properties. Paving materials consist of either hot-mix asphalt or Portland cement reinforced concrete. Upon completion, street maintenance remains the responsibility of the municipal government and is funded through general revenues.

Figure 1 illustrates an example of a paving project before and after completion. Urban street paving yields several benefits: it improves access for vehicles, pedestrians, and cyclists; creates parking space; facilitates commercial deliveries; and enhances the street's visual appearance. Field observations confirm that traffic congestion is minimal, consistent with the residential character of these roads. Despite these benefits, street paving is not privately provided due to its nature as a public good, which gives rise to free-rider problems that limit private investment.

Figure 1: Example of an Asphalting Project Before and After Pavement

(a) Before Pavement (b) After Pavement

Notes: The figure shows a typical treated street, before and after receiving the pavement. Source: Gonzalez-Navarro and Quintana-Domeque (2016).

Facing budget and time constraints, the Acayucan municipal government could pave only a subset of the city's unpaved streets. Under a memorandum of understanding with Gonzalez-Navarro and Quintana-Domeque, the public works office identified 56 eligible paving projects. Each project comprises a contiguous set of street blocks that, once paved, would connect to the existing paved road network. All candidate projects were located in populated areas and were eligible only if the streets had never been paved before.

11. The description of the context, pavement intervention and survey is adapted from Gonzalez-Navarro and Quintana-Domeque (2016), with substantial portions reproduced verbatim from the original source.

With available funds sufficient to complete approximately 28 of the 56 projects—each considered equally viable—the mayor and city council agreed to allow the research team to randomly select the projects to be implemented. Figure 2 displays the treatment and control streets. The paving program was not bundled with any other government initiative.

Figure 2: Location of the Asphalting Projects in Acayucan by Assignment

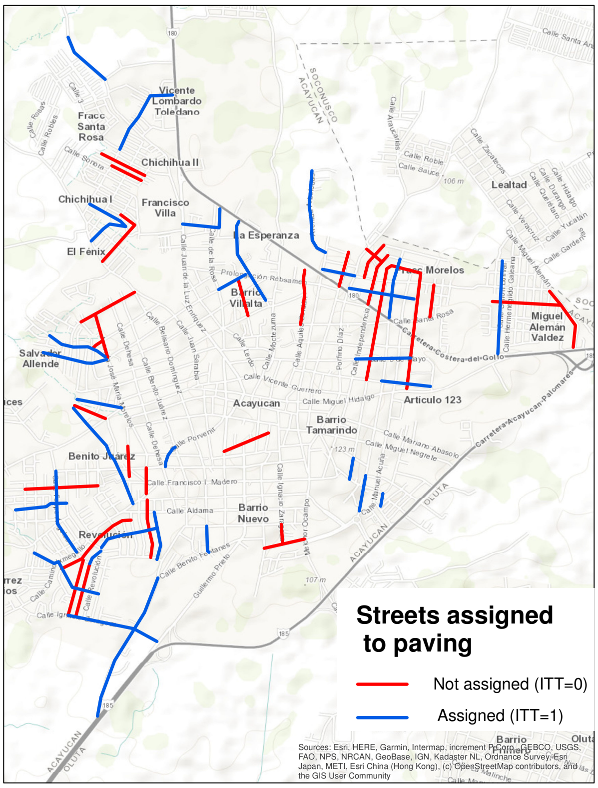

Notes:The figure illustrates the locations of streets designated for treatment (blue) and control (red) within Acayucan's street network.

Local governments in Mexico allocate a portion of their budgets to street paving, giving all unpaved streets some probability of eventual improvement. In practice, however, this probability is low: in Acayucan, an average of only 1.3% of streets were newly paved each year during the study period (see Table A.1). To preserve the integrity of the experimental design, the list of candidate and treated streets was not disclosed to the public, ensuring that control group expectations about future paving remained unaffected. Project identities were kept confidential because municipal budgeting occurs annually, whereas the study spanned multiple years. The selection thus served as an internal planning tool rather than a binding commitment. Residents were informed only upon the arrival of construction crews, minimizing the possibility of anticipation effects (Gonzalez-Navarro and Quintana-Domeque (2016) provides supporting evidence). Control streets did not receive any form of compensation.

Table A.1 and Figure A.1 present the implementation timeline for paving projects between 2007 and 2012. In the first three years, 19 of the 28 treatment streets (67%) were completed. However, weather and technical delays slowed progress: no projects were completed in 2010, only one in 2011, and six in the final year. These last seven had originally been assigned to the control group, resulting in two-sided non-compliance.

3.2 ASLS and Measuring Perceptions of Government Inefficiency

We use the Acayucan Standards of Living Survey (ASLS) to measure perceptions of local government inefficiency. The ASLS is a two-wave longitudinal household survey (2006, 2009) designed to evaluate the effects of street paving on low-income populations (Gonzalez-Navarro and Quintana-Domeque 2016). The target population included all occupied residential structures on streets selected for the experimental sample.[12]

To proxy beliefs about government efficiency, we use a self-reported measure of satisfaction with the local government, following Wang and Kim (2024). Respondents were asked to rate their satisfaction on a four-point Likert scale: (1) very unsatisfied, (2) unsatisfied, (3) satisfied, and (4) very satisfied. We construct a binary indicator that codes responses (1) and (2) as "dissatisfied" and (3) and (4) as "satisfied." We discuss the implications of using this proxy and a binary classification for our theoretical predictions in Appendix A.2.

The survey contains very detailed information. For instance, it allows us to quantify that only 5 percent of properties are occupied by renters, the rest being owner-occupied.

3.3 Property Tax Data and Measuring Compliance

We utilize administrative records of property tax payments for the municipality of Acayucan spanning the years 2005 to 2012. These data were merged with the 56 experimentally assigned street paving projects based on property addresses. Due to the peripheral location of some of these projects in recently urbanized areas, not all had been incorporated into the municipal cadastral system at the time of data collection. As a result, property tax information is available for 47 of the 56 projects.[13] Our primary tax compliance analysis focuses on a balanced panel of 752 plots located adjacent to either treatment or control streets, which together constitute the

12. The baseline survey included 1,231 households across 1,193 dwellings, with a 94% response rate. In 2009, 1,083 households were reinterviewed. In over 95% of cases, the respondent was either the household head or the spouse/partner of the head. Participants were not informed of the survey's ultimate purpose. Additional details are provided in Gonzalez-Navarro and Quintana-Domeque (2016).

13. Table A.2 shows that the 47 matched projects are very similar to the 9 unmatched projects.

experimental sample. This yields a total of 6,016 plot-year observations (752 × 8). Given our interest in the tax compliance response to public good provision, we restrict the analysis to plots that do not experience ownership changes over the study period. As detailed in subsection 6.5, we find no evidence of selective attrition by treatment status.

3.4 Measuring Exposure to Pavement

Main Exposure Measures. For each property p in our experimental sample (properties lining either treated or control streets), we measure exposure as follows: First, we select a destination node in the street network[14] and determine the shortest path between this destination and the property using a distance-weighted algorithm (hereafter, "shortest path to destination"; see Panel (a) of Figure 3 for an example). The number of nodes along this trajectory corresponds to X_p in our model, which remains constant over time.[15] We then compute two exposure measures:

1. *Intent-to-treat exposure* (E_p^Z): The share of nodes along the shortest path that lie either on streets assigned for paving or at intersections where such streets cross or end.

2. *Realized exposure* ($E_{p,t}$): Similar to the first measure, but based on streets that were actually paved rather than just assigned for paving.

Importantly, these measures are defined such that all properties in the experimental sample may have positive exposure, even if the plot itself is not directly adjacent to a treated street.

Panel (b) of Figure 3 illustrates this approach using a single property as an example and realized exposure as the metric. The shortest path, shown in violet, consists of 44 nodes (i.e., $X_p = 44$). The property is located in front of a paved street (depicted in yellow), which counts as one exposed node. The path includes three additional nodes on paved streets, further increasing exposure. Additionally, the path crosses one node where a paved street ends; while this paved street is not part of the shortest path, it is visible and therefore contributes to exposure. In total, 5 out of 44 nodes induce exposure, resulting in a value of $E_{p,t}$ equal to 11.36 percent.

Alternative Exposure Measures. To isolate visual exposure from direct use of paved streets, we define an alternative measure that excludes nodes in transited

14. Street network data are sourced from OpenStreetMap (Boeing 2025).

15. An underlying assumption is that the shortest path does not change due to the paving. This is likely the case since paved streets are located on the outskirts of the city and correspond to residential streets. Moreover, in Section 6.2, we present evidence that owners' perceptions of travel time to the city center remained unchanged, suggesting that their commuting routes did not change, that paving had minimal impact on travel times, or both.

streets along the shortest path—capturing only visual exposure to paving. This adjustment addresses concerns that our measure might be conflated with direct benefits from traveling on newly paved streets. A formal description of the alternative exposure measure is provided in Section 6.3.

Figure 3: Construction of the Exposure Measure with the City Center as the Destination Node

(a) Examples of Shortest Path to the City Center

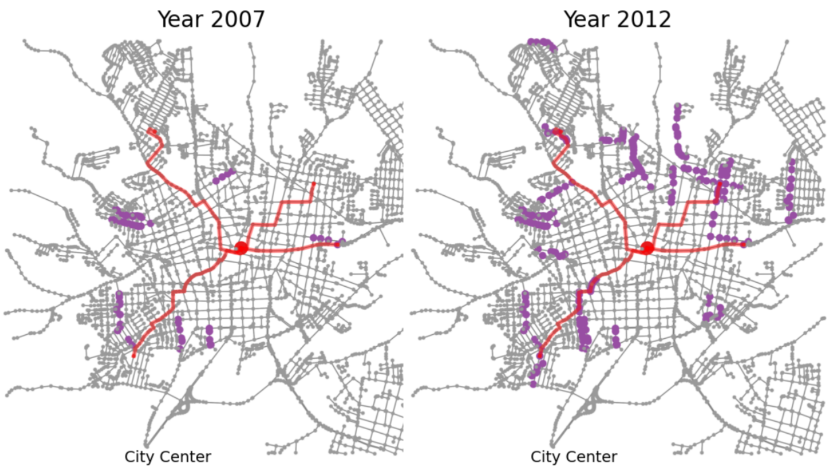

(b) Exposure Measure

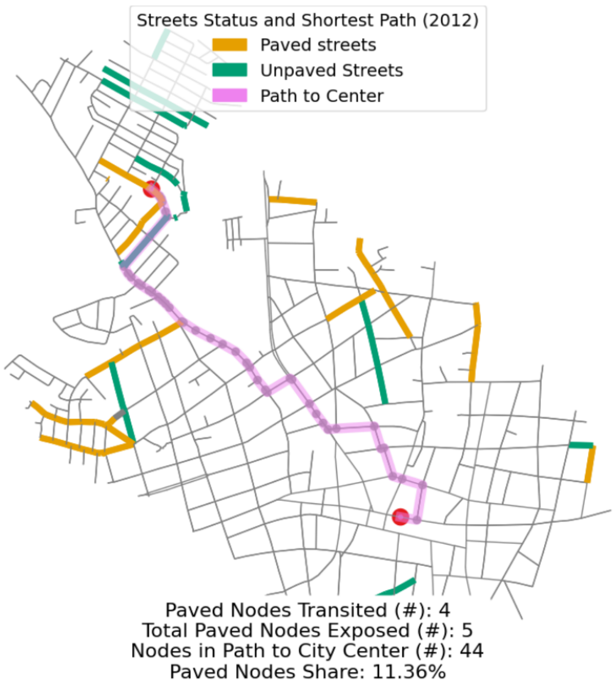

Notes: Panel (a) illustrates the construction of the exposure measure using Acayucan's city center as the destination node. The shortest paths from four experimental plots to the city center are computed using a distance-weighted algorithm based on OpenStreetMap data (Boeing 2025). Red lines represent the optimal routes, while purple circles denote nodes on streets paved through the intervention (on 2007 and 2012). Panel (b) visualizes the exposure measure for a single property. The shortest path from the property (northern red point) to the city center (southern red point) is shown in violet. In this example, the path passes through five (5) paved street nodes: one in front of the property, three along the shortest path, and one that belongs to paved streets that are not part of the shortest path but are visible. The realized exposure measure is 11.36 percent, calculated as the ratio of paved street nodes (5) to the total nodes in the path (44) multiplied by 100.

Our primary exposure measure uses Acayucan's downtown main square as the destination node. This central plaza, home to both the city hall and main catholic church, functions as the focal point for commercial and administrative activities, making it a natural reference point. The assumption underlying our measure is that individuals who regularly travel the shortest route between their residence and the city center will observe whether any streets along the route have been paved.

As a robustness check presented below, we also construct exposure measures using three additional key destinations: the municipal market, the main bank cluster (Banamex, Scotia and BBVA), and the principal health center. We also develop a composite exposure index that combines all four destination nodes.

Table A.3 and Figure A.2 present summary statistics for exposure measures to the city center. The intent-to-treat exposure averages 9.6%, while realized exposure increases from 2.57% to 12.4% as street paving progresses over time. As expected, properties adjacent to assigned-to-treatment streets (ITT=1) show higher average exposure, though properties on assigned-to-control streets (ITT=0) also exhibit positive exposure values. For ease of interpretation, we standardize all exposure measures by their standard deviation, so a one-unit increase corresponds to a one-standard-deviation change.

3.5 Measuring Proximity to Pavement

To measure proximity to pavement, we construct a series of indicator variables for each plot in our experimental sample. These indicators track the treatment status of streets along the ordered sequence from the plot to the destination node. For the street immediately adjacent to property p, we define two indicators: $b^Z_{p,0}$, which equals one if the street was assigned to treatment (zero otherwise), and $b_{p,0}$, which equals one if the street was actually paved (zero otherwise). Moving outward along the route to the destination node, we similarly define $b^Z_{p,1}$ and $b_{p,1}$ for the next street in sequence. This process continues sequentially until we reach the street containing the destination node.

These indicator variables serve as the empirical counterparts to equation (2.7), but are defined based on the sequence of streets in the shortest path rather than fixed distance rings. This approach ensures that no street is assigned to more than one distance ring and that, at most, only one street is paved or designated for paving within each ring. As a result, the average distance ring effectively corresponds to the average street length in the network.

3.6 Additional Variables

We use additional variables to further examine the potential benefits of street pavement, characterize the surveyed households, and extend our baseline balance analysis. The ASLS collects self-reported data on property and land values, as well as travel time to the city center (in minutes) using respondents' usual mode of transport. The survey also includes information on homeownership, civic knowledge (i.e., self-reported identification of the government level responsible for various public goods), and per-capita expenditure. In addition, we supplement self-reported property values with professional appraisals of land and property, conducted immediately after each survey round for approximately half the sample ($N = 578$).

4 Econometric Specifications

4.1 Econometric Specification to Test Prediction 1

Our model predicts that individuals exposed to public good provision update their beliefs about government inefficiency. Specifically, greater exposure to newly paved streets reduces perceived inefficiency (equation (2.6)). While we do not observe beliefs directly, we test this prediction using self-reported satisfaction with the local government from the ASLS as a proxy for beliefs about local government efficiency (Wang and Kim 2024).

Empirically, this requires two key departures from the theoretical framework: we replace a continuous measure of beliefs with a binary indicator, and we use reported dissatisfaction as a proxy for inefficiency perceptions. We formally discuss the implications of these changes in Appendix A.2. Intuitively, if latent dissatisfaction is positively correlated with perceived government inefficiency, then increased exposure to paved streets should cause some individuals to switch from reporting dissatisfaction to satisfaction. Importantly, since exposure reduces perceived inefficiency while no exposure leaves beliefs unchanged ($\theta_p(0) = \theta_p^*$), switching should occur only among those who were initially dissatisfied. Among those already satisfied, no change is expected. This leads to a modified empirical prediction: exposure should reduce dissatisfaction, with no corresponding change among the satisfied at baseline.

The ASLS is a longitudinal dataset with two waves, collected in 2006 and 2009. We test this prediction by estimating the following linear probability model of dissatisfaction with the local government for the property owner of plot p in the post-intervention year 2009, using OLS:

$$D_{p,09} = \rho_0^Z + \rho_1^Z D_{p,06} + \rho_2^Z E_{p,09}^Z + \rho_3^Z (D_{p,06} \times E_{p,09}^Z) + u_{p,09}, \tag{4.1}$$

where $D_{p,09} = 1$ if the property owner p reports dissatisfaction with the local government in 2009, $D_{p,06} = 1$ if the same owner reported dissatisfaction in 2006, and $E_{p,09}^Z$ is the *intent-to-treat* exposure measure (as defined in Section 3).

We then estimate the local average treatment effect (LATE) using 2SLS:

$$D_{p,09} = \rho'_0 + \rho'_1 D_{p,06} + \rho'_2 \widehat{E}_{p,09} + \rho'_3 (D_{p,06} \times \widehat{E}_{p,09}) + \upsilon_{p,09}, \tag{4.2}$$

where we instrument realized exposure $E_{p,09}$ with *intent-to-treat* exposure $E^Z_{p,09}$, so that $\widehat{E}_{p,09}$ captures variation in exposure induced by assignment to treatment.

The parameters of primary interest are ρ'_2 and ρ'_3 (and their reduced-form counterparts ρ^Z_2 and ρ^Z_3). The coefficient ρ'_2 captures the effect of exposure among those initially satisfied, which we expect to be zero:

$$\rho'_2 = \frac{\partial \mathbb{P}(D_{p,09} = 1 \mid D_{p,06} = 0, \widehat{E}_{p,09})}{\partial \widehat{E}_{p,09}} = 0.$$

The coefficient ρ'_3 captures the differential effect among those initially dissatisfied, which we expect to be negative:

$$\begin{aligned}
\operatorname{sign}(\rho'_3) &= \operatorname{sign}\left(\frac{\partial \mathbb{P}(D_{p,09} = 1 \mid D_{p,06} = 1, \widehat{E}_{p,09})}{\partial \widehat{E}_{p,09}} - \frac{\partial \mathbb{P}(D_{p,09} = 1 \mid D_{p,06} = 0, \widehat{E}_{p,09})}{\partial \widehat{E}_{p,09}} \right) \\
&= \operatorname{sign}\left(\frac{\partial \mathbb{P}(D_{p,09} = 1 \mid D_{p,06} = 1, \widehat{E}_{p,09})}{\partial \widehat{E}_{p,09}} \right) \leq 0.
\end{aligned}$$

4.2 Econometric Specifications to Test Predictions 2 and 3

Our property tax data is longitudinal. This allows us to control for plot/property fixed effects η_p:

$$C_{p,t} = \mathbf{1}\left[\theta^*_p \leq \eta_p + \delta_p E_{p,t} + \sum_j \gamma_{p,j} b_{p,t,j} \right],$$

and the latent index now can be defined as:

$$y^*_{p,t} = \eta_p + \delta_p E_{p,t} + \sum_j \gamma_{p,j} b_{p,t,j}.$$

Hence,

$$\mathbb{P}(C_{p,t} = 1 \mid \eta_p, E_p, \mathbf{b}_p) = \mathbb{P}(\theta^*_p \leq y^*_{p,t}) = F(y^*_{p,t}).$$

We estimate a linear probability model of tax compliance for property owner of plot p in year t using OLS. In its general form, we obtain the intention-to-treat (ITT) effects of pavement assignment with the following specification:

$$C_{p,t} = \eta_p + \phi_t + \delta^Z E^Z_{p,t} + \gamma^Z B^Z_{p,t} + \text{controls}_{p,t} + \epsilon_{p,t},$$

where $C_{p,t} = 1$ if plot p's property taxes are current in year t, η_p represents property fixed effects, and ϕ_t represents year fixed effects. The variable $E^Z_{p,t}$ is the *intent-to-*

treat exposure measure (as defined in Section 3), interacted with an indicator for $t \geq 2007$ to introduce time variation. $B^Z_{p,t} = 1$ if the street lining plot p is assigned to pavement, interacted with an indicator for $t \geq 2007$. The $\text{controls}_{p,t}$ include the cadastral property valuation and tax bill, both measured in 2005 and interacted with year dummies. These controls, combined with plot fixed effects, account for any baseline cross-sectional differences in tax liabilities (τ_p), perceptions of government inefficiency (θ^*_p), and the length of shortest-path trajectories (X_p) that may influence compliance according to our model.

When we take this model to the data, we begin with $\gamma^Z B^Z_{p,t} = \gamma^Z_0 b^Z_{0,p,t}$ and provide evidence that this is a good approximation, as additional benefits from distant assigned-to-be paved streets ($\gamma^Z B^Z_{p,t} = \gamma^Z_0 b^Z_{0,p,t} + \gamma^Z_1 b^Z_{1,p,t} + \gamma^Z_2 b^Z_{2,p,t}$) do not affect tax compliance. For this reason, our main specification is:

$$C_{p,t} = \eta_p + \phi_t + \delta^Z E^Z_{p,t} + \gamma^Z_0 b^Z_{0,p,t} + \text{controls}_{p,t} + \epsilon_{p,t}. \tag{4.3}$$

Next, we estimate an instrumented version of the model in which we utilize whether the street lining plot p has been assigned to pavement ($B^Z_{p,t}$) as an instrument for whether the street lining plot p is actually paved in year t ($B_{p,t}$). Similarly, we use *intent-to-treat* exposure ($E^Z_{p,t}$) as an instrument for *realized exposure* ($E_{p,t}$). This approach allows us to estimate local average treatment effects through the following second-stage equation using 2SLS:

$$C_{p,t} = \eta_p + \phi_t + \delta' \widehat{E}_{p,t} + \gamma' \widehat{B}_{p,t} + \text{controls}_{p,t} + \nu_{p,t}.$$

Once again, in the empirical analysis, we begin with $\gamma \widehat{B}_{p,t} = \gamma_0 \widehat{b}_{0,p,t}$ and provide evidence that additional benefits from distant paved streets do not appear to affect tax compliance. The main 2SLS specification becomes:

$$C_{p,t} = \eta_p + \phi_t + \delta' \widehat{E}_{p,t} + \gamma'_0 \widehat{b}_{0,p,t} + \text{controls}_{p,t} + \nu_{p,t}. \tag{4.4}$$

Under standard conditions — namely, that the link function is strictly increasing (as with a Beta *cdf*), the structural parameters δ_p and γ_p have consistent signs across units, and there is no saturation (i.e., outcomes are not clustered at 0 or 1) — the signs of the OLS and 2SLS coefficients are informative about the signs of the marginal effects. That is,

$$\text{sign}(\delta^Z) = \text{sign}(\delta') = \text{sign}\left(\frac{\partial \mathbb{P}(C_p = 1 \mid \eta_p, E_p, \mathbf{b}_p)}{\partial E_p}\right) \geq 0,$$

and

$$\text{sign}(\gamma^Z_0) = \text{sign}(\gamma'_0) = \text{sign}(\mathbb{P}(C_p = 1 \mid \eta_p, E_p, b_{p,j} = 1) - \mathbb{P}(C_p = 1 \mid \eta_p, E_p, b_{p,j} = 0)) \geq 0.$$

Remark. It is important to note that the instrument operates by increasing the probability that a street becomes part of the set of paved streets and, consequently, part of the streets defining an individual's exposure given their shortest-path trajectory. The monotonicity assumption is satisfied because assigning a street to treatment should not reduce the probability of it being paved, nor should it decrease exposure for individuals traveling along that trajectory.

Under this framework, the concept of a complier is tied to the street and, for exposure, to the set of individuals whose exposure depends on whether the street is part of a shortest path. Thus, the LATE interpretation applies to individuals whose exposure is defined by streets included in a shortest path. The plot/property fixed effects (η_p) are critical in this setting, as they condition the estimation on the shortest path relevant to each property owner p.

Table A.4 characterizes the complier population based on whether the street in front of the property was assigned to treatment. We focus on this variable because it is most directly tied to the randomization procedure and, being binary, allows us to implement the proportionality test proposed by Angrist and Pischke (2011). Specifically, we compute the sample analog of the ratio: $\frac{\mathbb{P}\left(b_{p,0}(1)>b_{p,0}(0)|Y_p=1\right)}{\mathbb{P}\left(b_{p,0}(1)>b_{p,0}(0)\right)}$, where $b_{p,0}(b^Z_{p,0})$ is an indicator for whether property p receives treatment under assignment $b^Z_{p,0} \in \{0,1\}$, and Y_p is a binary baseline covariate.[16]

The estimated size of the complier group, given by the sample analog of the denominator, is 55%. We find no evidence that compliers differ in the likelihood of baseline tax compliance or satisfaction with the local government. However, they are more likely to have below-median cadastral property values and shorter travel times to the city center. This suggests that the complier population—those who receive pavement under the experimental assignment but would not have received it otherwise—have slightly lower wealth but tend to live on streets that are more centrally located.

5 Results

5.1 Baseline Balance

Table 1 presents evidence of baseline balance across plot-level and project-level characteristics, from different data sources: survey, appraisals, property tax records, and OpenStreetMap data. The survey data come from the ASLS and include between 519 and 722 observations, depending on the variable. We examine four key pre-treatment (2006) household-level measures: dissatisfaction with local government, travel time to the city center (in minutes), self-reported property values, and per capita expenditure. We also incorporate professionally appraised property and land

16. For continuous variables, we define $Y_p = \mathbf{1}[Y_p > \text{median}(Y_p)]$, and label them as *High*.

values for a subsample of 334 plots. Across all survey-based and appraisal variables, we find no statistically significant differences between treatment and control groups at the 5% level. The magnitudes of the differences are small, and all balance individual tests suggest successful randomization. Since the ITT exposure measure applies to households adjacent to both treatment and control streets, we assess balance in ITT exposure by regressing standardized ITT exposure on each variable separately. Again, there is no evidence that exposure systematically predicts any of the household characteristics.

Next, we assess baseline balance using administrative plot-level data ($n = 752$), focusing on three pre-treatment (2006) indicators: assessed property values, annual property tax bills, and compliance with property tax obligations. For instance, the tax compliance rate was 85.5% for properties adjacent to streets assigned to treatment, compared to 84.6% for those near control streets (p-value = 0.830). The average government-appraised property value in the assigned to be treated group was 160,344.5 pesos (36,492 PPP USD in 2024), and not statistically different from that in the assigned to the control group.[17] This value is about 40 percent lower than the citywide average, reflecting the location of the paving projects in lower-income, peripheral neighborhoods. Cadastral values remained fixed over the period 2006–2012, including for newly paved properties, though tax rates were inflation-adjusted. As a result, any increase in market value from pavement represents a transfer to treated households. Likewise, changes in tax compliance cannot be attributed to increases in tax liability, as the tax bill remained unchanged. The annual property tax averaged was 159.44 pesos (35.68 PPP USD) in the assigned to be treated group, very similar to the one for the assigned to control group. Overall, we find evidence of balance when using data from the property tax records.

Finally, we assess balance in project-level characteristics ($n = 56$) derived from network analysis of OpenStreetMap data (Boeing 2025). These include street length, average Katz centrality, and shortest-path distance to the city center. In our context, Katz centrality captures how connected a plot is to the overall street network (Jackson 2008). These project characteristics are balanced across treatment and control streets. For example, the mean Katz centrality is statistically indistinguishable between groups (p-value = 0.58), and none of the project-level variables are correlated with ITT exposure.

Overall, the balance analysis supports the validity of the experimental design. The absence of statistically significant differences across a broad set of characteristics—spanning survey data, administrative records, and street network features—indicates that treatment and control groups are comparable at baseline.

17. PPP conversions use the World Bank's 2006 World Development Indicators and are adjusted for inflation using CPI data from the U.S. Bureau of Labor Statistics.

Table 1: Baseline Balance by Street ITT Status and ITT Exposure

	Street ITT Status ($b^Z_{0,p}$)				ITT Exposure (E^Z_p)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	ITT=1	ITT=0	Diff.	p-value	Coefficient	p-value	Obs.
Plot-Level ASLS Data and Property Appraisals							
Dissatisfied with the Local Government (0-1)	0.545	0.530	0.015	0.771	-0.014	0.542	722
	(0.4985)	(0.4998)	(0.0535)		(0.0243)		
Dissatisfied with Living in Acayucan (0-1)	0.145	0.139	0.006	0.847	-0.015	0.431	722
	(0.3530)	(0.3469)	(0.0310)		(0.0198)		
Log of Property Appraised Value (USD)	11.575	11.633	-0.058	0.557	0.076	0.246	334
	(0.650)	(0.796)	(0.099)		(0.065)		
Log of Property Owner Value (USD)	11.758	11.844	-0.087	0.535	0.076	0.332	519
	(0.978)	(0.955)	(0.138)		(0.078)		
Log of Land Appraised Value (USD)	10.272	10.139	0.133	0.199	0.067	0.291	334
	(0.608)	(0.499)	(0.103)		(0.063)		
Travel Time to the City Center (minutes)	19.699	20.442	-0.743	0.548	-0.098	0.904	722
	(10.563)	(12.047)	(1.228)		(0.808)		
Per Capita Expenditure (2006)	314.962	298.273	16.689	0.626	34.079	0.309	707
	(240.8181)	(248.8754)	(34.1060)		(33.2277)		
Plot-Level Administrative Data							
Tax Compliance (0-1)	0.855	0.846	0.008	0.830	0.010	0.582	752
	(0.353)	(0.361)	(0.038)		(0.018)		
Cadastral Property Value (USD)	36,492.395	42,158.541	-5,666.146	0.367	-3,521.507	0.253	752
	(39,254.489)	(45,788.269)	(6,231.454)		(3,047.475)		
Tax Owed (USD)	35.685	36.720	-1.035	0.499	-1.142	0.061	752
	(14.406)	(15.231)	(1.518)		(0.596)		
Project Level OpenStreetMap data							
Road Length (Meters)	414.661	545.826	-131.164	0.205	11.215	0.782	56
	(204.8076)	(501.841)	(102.433)		(40.413)		
Katz centrality	0.016	0.016	-0.0001	0.688	-0.0001	0.602	56
	(0.001)	(0.0008)	(0.0002)		(0.0002)		
Distance to City Center (Meters)	1,765.430	1,711.184	54.245	0.703	31.901	0.723	56
	(570.710)	(487.491)	(141.844)		(89.578)		

Notes: Authors' calculations are based on Acayucan's administrative data, the Acayucan Standards of Living Survey (ASLS), and appraisal valuations. Monetary values are reported in 2024 PPP USD. Robust errors, clustered at the project level, are reported in parentheses for the administrative and ASLS data. The ITT exposure variable (E^Z_p) is standardized by dividing by its standard deviation to ease interpretation of the coefficient (column 5). The Katz centrality is computed as the simple average of all the nodes within a project.

5.2 Testable Implication 1

Table 2 reports the ITT and LATE estimates based on OLS and 2SLS regressions of equations (4.1) and (4.2). Consistent with the theoretical predictions, we find that exposure to street paving significantly reduces dissatisfaction with the local government, but only among individuals who were dissatisfied prior to the intervention.

Columns 1 and 3 present specifications without interaction terms, estimating the average effect of a one–standard deviation increase in exposure across all respondents. These results indicate a reduction in dissatisfaction of 3.8 percentage points (ITT) and 4.3 percentage points (LATE). However, the effect is entirely driven by those who were dissatisfied at baseline, as shown in columns 2 and 4. Among this group, a one–standard deviation increase in exposure reduces dissatisfaction in 2009 by 7.9 percentage points (ITT) and 8.7 percentage points (LATE)—a 14% and 16%

decline, respectively, relative to the 2006 baseline dissatisfaction rate of 54%. We cannot reject a null effect of exposure among those who were satisfied at baseline.

Identification of the LATE relies on four assumptions: (a) random assignment of the instrument, (b) instrument relevance, (c) exclusion restriction, and (d) monotonicity. Assumption (a) is supported by the balance tests in Table 1. Assumption (c), the exclusion restriction, is not directly testable. However, in Section 6, we show that exposure does not significantly affect reported satisfaction with outcomes other than local government performance. Moreover, we find no evidence that exposure influences variables that could indicate direct private benefits—such as land or property valuations—or self-reported travel time to the city center. This supports the validity of our proposed mechanism. Assumption (d), monotonicity, is plausible in this setting: assignment to treatment should not decrease exposure along the travel path.

The Generalized First-Stage Statistic by Lewis and Mertens (2022), which nests the robust F-statistic or the effective F-statistic from Olea and Pflueger (2013) because there is only one single endogenous variable, is 218.01. This value exceeds the 39.25 threshold for 10% relative bias, so that we conclude that instrument relevance (assumption b) is satisfied.[18]

One concern is that the pattern observed in Table 2 could reflect mean reversion rather than belief updating. We provide two pieces of evidence against this interpretation. First, if mean reversion were at play, we would expect a negative coefficient on baseline dissatisfaction. Instead, the coefficients on baseline dissatisfaction are consistently large and positive across all specifications. In addition, the joint hypothesis that the sum of coefficients equals zero is strongly rejected in all cases (see table footnotes). Second, in Table A.5, we replicate the analysis from Table 2 using a restricted sample that excludes individuals who were either very satisfied or very dissatisfied at baseline—that is, the subsample least likely to exhibit mean reversion. The results remain virtually unchanged, suggesting that our estimated effects are not mechanically driven by reversion to the mean.

In sum, while we do not observe beliefs directly, the results align with the model's prediction that greater exposure enhances perceptions of government efficiency, thereby reducing dissatisfaction—particularly among those who were initially dissatisfied.

18. The effective F-statistic and its critical value are computed using the `weakivtest` routine in `Stata` (Pflueger and Wang 2015). The Generalized First-Stage Statistic and its critical value are from the `weakivtest2` routine (Zhou 2024).

Table 2: Effects of Street Paving Along the Shortest Path to the City Center on Satisfaction with the Local Government

	Dissatisfied with the Local Government = 1 (2009)			
	OLS		2SLS	
	(1)	(2)	(3)	(4)
Assigned to Pavement Exposure ($E^Z_{p,t}$)	-0.038*	0.012		
	(0.021)	(0.025)		
$E^Z_{p,t} \times D_{p,06}$		-0.091**		
		(0.031)		
Pavement Exposure ($\widehat{E}_{p,09}$)			-0.043*	0.013
			(0.022)	(0.027)
$\widehat{E}_{p,09} \times D_{p,06}$				-0.100**
				(0.032)
Dissatisfied in 2006 ($D_{p,06}$)	0.115**	0.221***	0.116***	0.218***
	(0.035)	(0.046)	(0.034)	(0.045)
Observations	722	722	722	722
Mean Dep. Var (2006)	.54	.54	.54	.54
Sum of Coefficients (Sum)		-.079		-.087
SE of Sum		.024		.024
P-value of Sum = 0		.002		0
Generalized First-Stage Statistic			690.9	218.01
Critical Value 10%			23.11	39.25

Notes: Authors' calculations use ASLS data from 722 georeferenced plots. Robust errors, clustered at the project level (53), are reported in parentheses. The *Dissatisfied* variable takes the value of 1 if the household head was very dissatisfied or dissatisfied, and 0 if they were satisfied or very satisfied. 2SLS uses the *Assigned to Paved* measure and its interaction with dissatisfaction at baseline as instruments for their corresponding *Paved* versions. The exposure measure is standardized by dividing by its standard deviation to ease interpretation. The p-value for testing H_0 : Sum of coefficients = 0 against H_1 : Sum of coefficients $\neq$ 0 measures the probability of observing a test statistic as extreme as (or more extreme than) the one computed from our sample, assuming the null hypothesis is true. The Generalized First-Stage Statistic is the one proposed by Lewis and Mertens (2022), which nests the robust F-statistic or the effective F-statistic from Olea and Pflueger (2013) since there is one endogenous variable. Estimation accounts for survey weights.

5.3 Testable Implications 2 and 3

In our model, tax compliance responds to public good provision through two primary channels. First, individuals are more likely to pay their property taxes when their exposure to street paving—anywhere in the network—is higher, as observed provision reduces beliefs about government inefficiency. Second, compliance increases when individuals derive direct benefits from paved streets. We test these predictions using administrative property tax records, OpenStreetMap data, and information on streets assigned for paving and ultimately paved.

Table 3 presents the results. Columns (1)–(3) report intent-to-treat (ITT) effects estimated via OLS. Column (1) shows that property owners on streets assigned to be paved are 4.8 percentage points (pp) more likely to pay their property taxes. Given a baseline compliance rate of 85%, this corresponds to a 5.6% increase. Column (2) estimates that a one-standard-deviation increase in the share of nodes along the shortest path to downtown that pass through streets assigned to be paved increases compliance by 2.3 pp. Column (3) reports joint ITT estimates from equation (4.3), which includes both channels and adjusts for potential bias in the separate estimates. The estimated effects are 3.2 pp for direct benefits and 1.5 pp for exposure.

Columns (4)–(6) present 2SLS estimates, using assignment to paving and assigned exposure as instruments for actual paving and exposure to newly paved routes. Columns (4) and (5) estimate each effect separately; Column (6) reports joint local average treatment effects (LATEs) from equation (4.4). As before, identification relies on four assumptions, but now we have two instruments and two treatments, so that the slighlty modified assumptions are: (a) random assignment of the two instruments, (b) relevance of the two instruments, (c) exclusion restrictions of the two instruments, and (d) monotonicity of the two instruments. Table 1 supports assumption (a). Exclusion restrictions (c) are not testable, in general. Monotonicity (d) is reasonable: assignment to treatment should not reduce the likelihood of paving or decrease exposure along the travel path.

The relevance of the two instruments (b) is assessed by means of the the Generalized First-Stage Statistic from Lewis and Mertens (2022), which accommodates multiple endogenous variables and heteroskedasticity. Instrument strength is robust throughout. The Generalized First-Stage Statistics are 50.18 and 37.05 in Columns (4) and (5), both exceeding the 23.1 threshold for 10% relative bias. In Column (6), the statistic is 29.44, above the critical value of 26.86.

The 2SLS estimates show that property owners whose streets were paved as a result of assignment are 4.8 pp more likely to comply. A one-standard-deviation increase in exposure to paved intersections along the shortest route to downtown—caused by the intervention—increases compliance by 2.6 pp. These estimates separately identify the direct effect of street paving and the exposure effect resulting from streets anywhere in the network being paved.

Taken together, our ITT and 2SLS results support testable implications 2 and 3 and provide new evidence that improvements in local public goods can generate economically meaningful increases in tax compliance. While direct effects are larger, exposure to public works along commonly traveled routes also raises compliance. This exposure effect points to a fiscal externality operating through the signaling value of government projects—a mechanism, to our knowledge, not previously documented in the context of local tax enforcement. We quantify the magnitude of this fiscal spillover in Section 7.

Table 3: Direct and Exposure Effects of Street Paving on Tax Compliance

	Tax compliance = 1					
	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
γ_0^Z: Plot p Assigned to Pavement $\times I_{\geq 2007}$ ($b^Z_{0,p,t}$)	0.048**		0.032*			
	(0.016)		(0.019)			
δ^Z: Assigned to Pavement Exposure $\times I_{\geq 2007}$ ($E^Z_{p,t}$)		0.023***	0.015**			
		(0.005)	(0.007)			
γ_0: Plot p Paved ($\widehat{b}_{0,p,t}$)				0.077**		0.048*
				(0.024)		(0.027)
δ: Pavement Exposure ($\widehat{E}_{p,t}$)					0.038***	0.026**
					(0.010)	(0.013)
Observations	6,016	6,016	6,016	6,016	6,016	6,016
Mean Dep. Var (2006)	0.85	0.85	0.85	0.85	0.85	0.85
Generalized First-Stage Statistic				50.18	37.05	29.44
Critical Value 10%				23.1	23.1	26.86

Notes: The table reports estimates of the direct benefit and exposure effects of street paving on tax compliance for the sample. Robust errors, clustered at the project level (47), are reported in parentheses. All regressions include plot fixed effects, year fixed effects and baseline controls (cadastral property value and the tax bill interacted with year dummies). The Pavement Exposure measure for each plot is defined as the fraction of intersections -observed or transited- along the shortest path trajectory between the closest node to the plot and the city center that passes through streets paved by the intervention. The exposure measure is standardized by dividing each value by the standard deviation computed within each year (2007–2012), to ease interpretation. The Assigned to Pavement Exposure measure is defined similarly, but uses the fraction of the shortest path trajectory that passes through streets assigned to be paved. The Generalized First-Stage Statistic is the one proposed by Lewis and Mertens (2022), which allows for the assessment of instrument strength when there are two instruments and two endogenous variables under non-homoskedasticity. In columns (4) and (5), one endogenous variable and one instrument, this statistic corresponds to the robust F-statistic or the effective F-statistic by Olea and Pflueger (2013).

While Table 3 uses only an indicator of whether the street in front of a household was paved (or assigned for pavement) as a proximity-based benefit indicator, in Table A.6, we estimate more flexible versions: we include $b^Z_{1\,p,t}$ and $b^Z_{2\,p,t}$ for OLS regressions, as well as $\widehat{b}_{1,p,t}$ and $\widehat{b}_{2,p,t}$ for 2SLS regressions—compared to our parsimonious specifications (4.3) and (4.4). These estimates do not alter our findings from Table 3. This is not only consistent with our assumption that direct benefits are a function of distance, but also suggests that they fade away when pavement occurs beyond the adjacent street. Thus, our parsimonious specifications (4.3) and (4.4) appear reliable, and we adopt them as our main ones.

6 Validity Checks

In this section, we perform several validity checks. First, we examine whether exposure affects satisfaction in domains other than the local government. Second, we examine whether exposure captures direct material benefits in addition to the hypothesized belief-updating mechanism. Third, we test the robustness of the exposure effect to alternative measurement strategies. Fourth, we redefine the exposure measure using alternative destinations and account for potential non-random exposure. Finally, we conduct additional analyses, including checks for selective attrition.

6.1 Does the Exposure Affect Satisfaction in Other Domains?

Our interpretation that exposure influences satisfaction with the local government through a belief-updating mechanism relies on the absence of alternative indirect channels. One potential concern is that exposure to paving may increase satisfaction in other domains, such as satisfaction with living in the city (Castelli et al. 2023), which could in turn be positively correlated with satisfaction with the local government. If this were the case, it would challenge the validity of our interpretation centered on belief updating.

To assess this possibility, Table 4 reports estimates of equations (4.1) and (4.2), using self-reported satisfaction with living in Acayucan—rather than satisfaction with the local government—as the outcome variable. Satisfaction with living in Acayucan is measured on a four-point Likert scale in the ASLS, which we dichotomize into a binary indicator equal to one if the respondent reports being either dissatisfied or very dissatisfied with living in Acayucan.

We find no evidence that exposure affects satisfaction with living in Acayucan, either on average or among respondents who reported being dissatisfied with living in Acayucan at baseline. This suggests that the effect of exposure is specific to perceptions of government performance, rather than reflecting broader changes in how individuals evaluate their overall satisfaction with life in Acayucan.

Table 4: Effects of Street Paving Along the Shortest Path to the City Center on Satisfaction with living with Acayucan

	Dissatisfied with Living in Acayucan in 2009 = $1(D^{Living}_{p,09})$			
	(1)	(2)	(3)	(4)
	OLS		2SLS	
Assigned to Pavement Exposure ($E^Z_{p,t}$)	0.009	0.007		
	(0.008)	(0.008)		
$E^Z_{p,t} \times D^{Living}_{p,06}$		0.012		
		(0.046)		
Pavement Exposure ($\widehat{E}_{p,09}$)			0.010	0.008
			(0.009)	(0.009)
$\widehat{E}_{p,09} \times D^{Living}_{p,06}$				0.011
				(0.047)
Dissatisfied with life in 2006 ($D^{Living}_{p,06}$)	0.094**	0.082	0.094**	0.083
	(0.041)	(0.055)	(0.040)	(0.053)
Observations	722	722	722	722
Mean Dep. Var (2006)	.14	.14	.14	.14
Sum of Coefficients		.019		.019
S.E. (Sum = 0)		.043		.044
P-value (Sum = 0)		.662		.66
F-generalized			697.7	236.74
Critical Value (10)			23.11	35.97

Notes: Authors' calculations use ASLS data from 722 georeferenced plots. Robust errors, clustered at the project level (53), are reported in parentheses. The *Dissatisfied* variable takes the value of 1 if the household head was very dissatisfied or dissatisfied, and 0 if they were satisfied or very satisfied with living in Acayucan. 2SLS uses the *Assigned to Paved* measure and its interaction with dissatisfaction at baseline as instruments for their corresponding *Paved* versions. The exposure measure is standardized by dividing by its standard deviation to ease interpretation. The p-value for testing H_0 : Sum of coefficients = 0 against H_1 : Sum of coefficients $\neq$ 0 measures the probability of observing a test statistic as extreme as (or more extreme than) the one computed from our sample, assuming the null hypothesis is true. The Generalized First-Stage Statistic is the one proposed by Lewis and Mertens (2022), which nests the robust F-statistic or the effective F-statistic from Olea and Pflueger (2013) since there is one endogenous variable. Estimation accounts for survey weights.

6.2 Does the Exposure Measure capture Direct Benefits?

A key assumption of the model is that, conditional on the treatment status of nearby streets, exposure should not generate direct material benefits. Empirically, we define exposure as the fraction of paved streets along the shortest path to a key destination (e.g., the city center). In practice, however, this measure could also capture direct benefits—such as improved access or perceived service provision—which may increase compliance through reciprocity rather than belief updating.

To test this possibility, we examine whether our exposure measure affects two outcomes plausibly influenced by paving: property values and commuting times. We first estimate the effect of residing on a street that was paved—or assigned to be paved—on property values and travel time to the city center. We then assess whether greater exposure to paving projects along the route to the city center affects either of these outcomes.

Table 5 presents the results. Consistent with Gonzalez-Navarro and Quintana-Domeque (2016), we find that paving the street in front of a property increases land and property values—by 55 and 18 log points, respectively, based on professional appraisal data. These magnitudes are comparable to those estimated by Sorin (2025) in Uganda. Since government-assessed property values and tax rates remained fixed during the study period, these effects reflect capital gains attributable to the intervention, suggesting a wealth transfer to property owners directly benefiting from paving. By contrast, and in line with model predictions, exposure to paving en route to the city center has no effect on land or property values. While the 2SLS (LATE) estimates using appraisal data may suffer from weak instruments, the OLS (ITT) results are qualitatively similar.

Turning to commuting outcomes, we find no evidence that either treatment affects travel time to the city center, as reported by respondents in the ASLS survey based on their usual mode of transportation. The estimated effects are both statistically insignificant and quantitatively small: approximately a half-minute reduction in average travel time for those living on a paved street, consistent with Gonzalez-Navarro and Quintana-Domeque (2016), and no measurable change associated with increased exposure. These null effects align with the residential character and peripheral location of the streets included in the paving program.

These findings suggest that our exposure measure (conditional on plot/property fixed effects, the direct benefit measure and the control variables) does not capture direct benefits—neither through higher property values nor through shorter commuting times. This supports our interpretation that the estimated effect of exposure operates primarily through a belief-updating mechanism, rather than through reciprocity motivated by personal material gains.

Table 5: Direct and Exposure Effects of Street Paving in Measures of Direct Benefits

	OLS				2SLS			
	Appraisal Valuation 2009		Owner Perception 2009		Appraisal Valuation 2009		Owner Perception 2009	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log Land Value	Log Property Value	Log Property Value	Travel Time to City Center (min)	Log Land Value	Log Property Value	Log Property Value	Travel Time to City Center (min)
Plot p Assigned to Pavement $\times$ $I_{\geq 2007}$	0.328***	0.107**	0.143	-0.405				
	(0.075)	(0.043)	(0.110)	(0.806)				
Assigned to Pavement Exposure $\times$ $I_{\geq 2007}$	-0.038	-0.021	0.048	-0.092				
	(0.031)	(0.015)	(0.058)	(0.735)				
Plot p Paved					0.545***	0.177***	0.225	-0.646
					(0.118)	(0.061)	(0.165)	(1.310)
Pavement Exposure					-0.057*	-0.029*	0.047	-0.083
					(0.034)	(0.017)	(0.062)	(0.813)
Observations	318	318	468	721	318	318	468	721
Mean Dep. Var (2006)	10.22	11.62	11.75	20.37	10.22	11.62	11.75	20.37
Generalized First-Stage Statistic					23.72	23.35	31.46	30.19
Critical Value 10%					24.33	24.2	23.18	23.03

Notes: The table reports estimates of the direct and exposure effects of street paving on four measures of direct benefits. Clustered errors at the project level (53), are reported in parentheses. The logarithm transformation in home and land values in 2009 for a sample of 318 appraised houses in Acayucan, Mexico. Also, owner's perception of the house value (for 468 plots) and travel time to the city center in minutes (for 721 plots). For further details on the appraisal procedure and the variables compilation see Gonzalez-Navarro and Quintana-Domeque (2016). The regression includes the value of the dependent variable at baseline (2006) as a control. For details regarding the exposure measures definitions, refer to the main text. The exposure measure is standardized by dividing by its standard deviation to ease interpretation. Estimation accounts for survey weights. The Generalized First-Stage Statistic is the one proposed by Lewis and Mertens (2022), which allows for the assessment of instrument strength when there are two instruments and two endogenous variables under non-homoskedasticity.

6.3 What If We Use an Alternative Measure of Exposure?

While greater exposure to pavement projects does not lead to increases in property values or reductions in commuting times, one could argue that our original measure of exposure might still capture other relevant margins of direct benefits. We address this potential concern by constructing an alternative exposure measure that excludes nodes located on streets transited along the shortest path, thereby capturing only visual exposure to paving. This measure is designed to exclude any benefits derived from directly using the paved streets.

More formally, let the shortest path from property p to the destination node be represented by the ordered sequence of nodes:

$$N_p = \{n_{p,j}\}_{j=0}^{J},$$

where $j = 0$ denotes the node adjacent to the property and $j = J$ the destination node. The path length is $X_p = |N_p|$, where $|\cdot|$ indicates the cardinality of the set.

Define $N_{p,t}^{pav} \subseteq N_p$ as the subset of nodes on this path that are located on streets paved by the experiment by time t. We classify these paved nodes into two categories:

- A node $n_{p,j} \in N_{p,t}^{pav}$ is classified as **transited** if it lies within a contiguous paved segment, that is, if either $n_{p,j-1}$ or $n_{p,j+1} \in N_{p,t}^{pav}$.
- A node $n_{p,j} \in N_{p,t}^{pav}$ is classified as **observed** if both neighboring nodes satisfy $n_{p,j-1}, n_{p,j+1} \in N_p \setminus N_{p,t}^{pav}$.

Let $N_{p,t}^{\text{tran}} \subseteq N_{p,t}^{pav}$ be the set of transited paved nodes at time t, and $N_{p,t}^{\text{obs}} \subseteq N_{p,t}^{pav}$ the set of observed paved nodes, where $|N_{p,t}^{pav}| = |N_{p,t}^{\text{tran}}| + |N_{p,t}^{\text{obs}}|$. The primary exposure measure is then given by:

$$E_{p,t} = \frac{|N_{p,t}^{\text{pav}}|}{X_p}.$$

In contrast, the alternative exposure measure is:

$$E_{p,t}^{\text{obs}} = \frac{|N_{p,t}^{\text{obs}}|}{X_p}.$$

Analogous measures based on *assignment to pavement*—rather than actual paving—are defined in the same way.

While this may stretch the limits of our data,[19] the results suggest that the belief updating mechanism is not only plausible in our context but also empirically relevant. Indeed, the analysis in Table 6 shows that the ITT and LATE estimates based on the alternative exposure measure are 0.012 (SE = 0.007) and 0.022 (SE = 0.013), respectively—very close to the corresponding estimates of 0.015 (SE = 0.007) and 0.027 (SE = 0.012) in Table 3. Thus, it appears that both measures of exposure—whether *used and observed* or *only observed*—capture the belief updating mechanism postulated in the model.[20]

19. The estimates in the first column for the 2SLS specification appear to show a relevant instrument. However, the estimates in the last two columns may suffer from weak instruments. Nevertheless, the OLS (ITT) estimates yield qualitatively similar findings.

20. Appendix Table A.7 decomposes the exposure measure into exposure to transited nodes and exposure to observed-only nodes. The results once again indicate that the belief updating mechanism is empirically plausible.

Table 6: Effects of Street Paving in Front of Properties and Along the Shortest Path to the City Center on Tax Compliance: Observed (not used) Intersections

	Tax compliance = 1					
	OLS			**2SLS**		
	(1)	(2)	(3)	(4)	(5)	(6)
γ_0^Z: Plot p Assigned to Pavement $\times I_{\geq 2007}$ ($b^Z_{0,p,t}$)	0.048***		0.047***			
	(0.016)		(0.017)			
δ^Z: Exposure to ITT Observed Intersections $\times$ $I_{\geq 2007}$		0.013	0.012			
		(0.008)	(0.007)			
γ_0: Plot p Paved ($\widehat{b}_{0,p,t}$)				0.077***		0.071***
				(0.024)		(0.025)
δ: Observed Paved Intersections Exposure					0.021*	0.022*
					(0.012)	(0.012)
Observations	6,016	6,016	6,016	6,016	6,016	6,016
Mean Dep. Var (2006)	0.85	0.85	0.85	0.85	0.85	0.85
Generalized First-Stage Statistic				50.19	19.38	16.49
Critical Values 10%				23.11	23.11	26.86

Notes: The table reports estimates of the direct benefit and exposure effects of street paving on tax compliance for the pooled sample. Robust errors, clustered at the project level (47), are reported in parentheses. All regressions include plot fixed effects, year fixed effects, and baseline controls (cadastral property value and the tax bill interacted with year dummies). The Pavement exposure measure excludes nodes on streets along the shortest path trajectory, ensuring that *only nodes on paved streets that are visible but not directly traversed contribute to exposure*. The Pavement exposure measure for each plot is defined as the fraction of nodes -observed- along the shortest path trajectory between the closest node to the plot and the city center that passes through streets paved by the intervention. The exposure measure is standardized by dividing each value by the standard deviation computed within each year (2007–2012), to ease interpretation. The Assigned to pavement exposure measure is defined similarly, but uses the share of nodes in the shortest path trajectory that passes through streets assigned to be paved. The Generalized First-Stage Statistic is the one proposed by Lewis and Mertens (2022), which allows for the assessment of instrument strength when there are two instruments and two endogenous variables under non-homoskedasticity. In columns (4) and (5), one endogenous variable and one instrument, this statistic corresponds to the robust F-statistic or the effective F-statistic by Olea and Pflueger (2013).

6.4 Alternative Destinations and Non-Random Exposure

Measuring exposure using the city center as the destination node is intuitive, but alternative landmark destinations—such as the municipal market, the main bank cluster (Banamex, Scotia, and BBVA), and the principal health center—could also be relevant. Table A.8, Panel A, in the appendix shows that the results are robust to using these alternative destinations, as well as to a harmonized index combining all of them.

Another potential concern with the exposure measure is the issue of non-random exposure (Borusyak and Hull 2023). However, in our context, this is unlikely to be a major concern, as the probability of any given street being paved—as previously mentioned—is very small (1.3%). Indeed, Table A.8, Panel B, in the appendix, which employs recentered exposure measures following Borusyak and Hull (2023), produces estimates that are very similar to those in Table 3.

6.5 Assessing Selective Attrition

Finally, we note that our primary sample for the analysis of tax compliance is a balanced panel of properties with consistent ownership from 2005 to 2012. To avoid confounding treatment assignment, we exclude properties with ownership changes during the experiment. However, this may introduce bias if excluded properties differ systematically or if attrition varies by treatment status. In Table A.9, columns 1–3, we find no evidence of differential attrition rates by treatment status over time, suggesting that attrition is not systematically related to treatment assignment. Nevertheless, even with identical attrition rates, the characteristics of properties leaving the sample may systematically differ between treatment and control groups. To address this, we follow the approach in Ghanem, Hirshleifer, and Ortiz-Beccera (2023) and report baseline compliance rates in columns 4–7 across four groups: assigned treatment stayers, control stayers, assigned treatment leavers, and control leavers. We cannot reject internal validity for stayers (i.e., equality of baseline outcome distributions between assigned treatment and control stayers, or between assigned treatment and control leavers) or for the study population as a whole (i.e., equality across all four subgroups). These exercises confirm that attrition does not compromise the validity of our findings. Full details are provided in Appendix Section A.3.

7 Fiscal Implications and Spillovers of Street Paving

We conclude our analysis with a back-of-the-envelope calculation to evaluate the fiscal impact of the intervention. This analysis includes both experimental streets—those that either received pavement or remained unpaved as part of the experiment—and non-experimental unpaved streets that were not directly included in the intervention but were indirectly exposed to the paving projects. Importantly, individuals living on both unpaved streets in the experimental group and non-experimental unpaved streets could experience indirect exposure.

Our calculation considers tax revenue from three groups of properties: (i) directly treated properties on paved experimental streets (P); (ii) indirectly affected properties on unpaved experimental streets (U); and (iii) indirectly affected properties on unpaved non-experimental streets (NE). Additional revenue is calculated as:

$$\text{Revenue}_{p,t} = \begin{cases} \text{Tax\_bill}_{p,t} \times \left(\hat{\delta}' \, E_{p,t} + \hat{\gamma'}_0 \, b_{p,0,t} \right), & p \in P \\ \text{Tax\_bill}_{p,t} \times \left(\hat{\delta}' \, E_{p,t} \right), & p \in U \cup NE \end{cases}$$

where $\hat{\delta}'$ and $\hat{\gamma'}_0$ are the 2SLS coefficients from equation (4.4), $E_{p,t}$ denotes realized exposure to paving, and $b_{p,0,t}$ is an indicator for whether the street directly in front of property p was paved.

While the absolute revenue gains are modest due to the limited size of the property tax base, the relative effects reveal substantial fiscal spillovers. As shown in Table 7, in 2012 the ratio of tax revenue generated by properties on unpaved experimental streets to those on paved streets is 0.40.

When we expand the comparison to include unpaved properties on non-experimental streets, the magnitude of the spillover becomes more apparent. In 2012, the combined revenue from unpaved experimental and non-experimental streets is 4.48 times greater than the revenue from directly treated streets.

These findings highlight the broader fiscal impact of the paving intervention and underscore the role of belief updating. Specifically, observed improvements in public infrastructure lead to substantial increases in compliance—even among those who do not receive direct material benefits. As a result, the revenue gains attributable to belief-driven spillovers exceed those from direct paving effects by a factor of more than four. These results suggest that local public investments can generate substantial fiscal returns through indirect behavioral channels.

Table 7: Increase in Revenue due to Paving, in 2024 PPP USD

Year	Revenue			Number of plots			Relative Revenue	
	Paved (P)	Unpaved (U)	Non-experimental Unpaved (NE)	P	U	NE	U/P	(U + NE)/P
2007	340.59	398.50	554.85	107	645	6037	1.17	2.80
2008	480.01	544.99	2419.35	140	612	6037	1.14	6.18
2009	870.64	451.38	3355.69	235	517	6037	0.52	4.37
2010	906.13	469.10	3491.57	235	517	6037	0.52	4.37
2011	967.76	484.86	3642.91	250	502	6037	0.50	4.27
2012	1372.25	543.26	5599.08	347	405	6037	0.40	4.48

Notes: The counterfactual individual's difference in revenue due to paving was obtained throughout the multiplication of the marginal increase in probability of tax compliance with each individual's tax burden. The LATE estimates from table 3 and the computed exposure to paving were used for the computation. Finally, the table presents the aggregation (sum) of these counterfactual differences.

While this is a back-of-the-envelope calculation, it illustrates the potentially large fiscal returns to public infrastructure investments once indirect behavioral responses are taken into account.

8 Conclusions

This study highlights the role of public goods provision in improving tax compliance. We develop a model in which compliance depends on both tax morale and perceptions of government efficiency in using tax revenues to deliver public goods. In our setting, public expenditure is represented by the paving of previously unpaved streets in the periphery of Acayucan, Mexico, where 28 out of 56 eligible street projects were randomly assigned to be paved.

We find that first-time paving of residential streets—a highly visible infrastructure improvement—leads to increased property tax compliance. Consistent with our model, which allows compliance to respond to both perceived direct benefits and the observability of government action, we document three key findings.

First, exposure to pavement reduces dissatisfaction with the local government, but only among individuals who were dissatisfied prior to the intervention.

Second, property owners whose streets were assigned to be paved—or were ultimately paved—are more likely to pay property taxes, consistent with compliance responding to perceived direct benefits.

Third, property owners are also more likely to comply when their frequent travel routes pass through streets that were assigned to be paved—or were ultimately paved—indicating that exposure alone, even absent direct benefits, strengthens beliefs about government quality and increases compliance.

A back-of-the-envelope calculation suggests large fiscal spillovers from the paving intervention, including on non-experimental streets indirectly exposed to paving. Our estimates imply that belief updating among indirectly exposed citizens generates four times as much additional revenue as reciprocity-driven compliance among direct beneficiaries. These findings underscore the importance of indirect effects in shaping fiscal outcomes.

From a policy perspective, the results highlight the fiscal and behavioral returns to investing in visible, high-quality public goods. Such investments not only demonstrate the effective use of taxpayer funds but also reinforce the social contract between citizens and the state, fostering a virtuous cycle of improved compliance and better public service delivery. This approach is especially relevant in developing countries, where breaking the cycle of low compliance and under-provision of public goods is critical for sustainable development. Whether similar effects extend to less visible public goods remains an open question for future research (Sandholtz and Vicente 2024).

References

Ajzenman, Nicolás, Martín Ardanaz, Guillermo Cruces, Germán Feierherd, and Ignacio Lunghi. 2024. *Unraveling the Paradox of Anticorruption Messaging: Experimental Evidence from a Tax Administration Reform.* Technical report. Inter American Development Bank. https://doi.org/10.18235/0012975.

Allingham, Michael G, and Agnar Sandmo. 1972. "Income tax evasion: a theoretical analysis." *Journal of Public Economics* 1, nos. 3-4 (November): 323–338. https://doi.org/10.1016/0047-2727(72)90010-2.

Andreoni, James, Brian Erard, and Jonathan Feinstein. 1998. "Tax Compliance." *Journal of Economic Literature* 36 (2): 818–860.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2011. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton, NJ: Princeton University Press.

Antinyan, Armenak, and Zareh Asatryan. 2024. "Nudging for Tax Compliance: A Meta-Analysis." *The Economic Journal* (October). https://doi.org/10.1093/ej/ueae088.

Becker, Gary S. 1968. "Crime and Punishment: An Economic Approach." *Journal of Political Economy* 76, no. 2 (March): 169–217. https://doi.org/10.1086/259394.

Benabou, Roland, and Jean Tirole. 2011. *Laws and Norms.* Working Paper, Working Paper Series 17579. Cambridge, MA: National Bureau of Economic Research, November. https://doi.org/10.3386/w17579.

Bérgolo, Marcelo, Rodrigo Ceni, Guillermo Cruces, Matias Giaccobasso, and Ricardo Perez-Truglia. 2017. *Tax Audits as Scarecrows: Evidence from a Large-Scale Field Experiment.* Technical report. Cambridge, MA: National Bureau of Economic Research, July. https://doi.org/10.3386/w23631.

Besley, Timothy. 2020. "State Capacity, Reciprocity, and the Social Contract." *Econometrica* 88 (4): 1307–1335. https://doi.org/10.3982/ECTA16863.

Besley, Timothy, Anders Jensen, and Torsten Persson. 2023. "Norms, Enforcement, and Tax Evasion." *Review of Economics and Statistics* 105, no. 4 (July): 998–1007. https://doi.org/10.1162/rest_a_01123.

Besley, Timothy, and Torsten Persson. 2014. "Why Do Developing Countries Tax So Little?" *Journal of Economic Perspectives* 28, no. 4 (November): 99–120. https://doi.org/10.1257/jep.28.4.99.

Boeing, Geoff. 2025. "Modeling and Analyzing Urban Networks and Amenities with OSMnx." Published online ahead of print on May 3, 2025, *Geographical Analysis,* https://doi.org/10.1111/gean.70009.

Borusyak, Kirill, and Peter Hull. 2023. "Nonrandom Exposure to Exogenous Shocks." *Econometrica* 91 (6): 2155–2185. https://doi.org/10.3982/ECTA19367.

Bott, Kristina M., Alexander W. Cappelen, Erik Ø. Sørensen, and Bertil Tungodden. 2020. "You've Got Mail: A Randomized Field Experiment on Tax Evasion." *Management Science* 66, no. 7 (July): 2801–2819. https://doi.org/10.1287/mnsc.2019.3390.

Brockmeyer, Anne, Francisco Garfias, and Juan Carlos Suárez Serrato. 2024. *The Fiscal Contract up Close: Experimental Evidence from Mexico City.* Working Paper, Working Paper Series 32776. National Bureau of Economic Research, August. https://doi.org/10.3386/w32776.

Carrillo, Paul E., Edgar Castro, and Carlos Scartascini. 2021. "Public good provision and property tax compliance: Evidence from a natural experiment." *Journal of Public Economics* 198 (June): 104422. https://doi.org/10.1016/j.jpubeco.2021.104422.

Castelli, Chiara, Beatrice d'Hombres, Laura de Dominicis, Lewis Dijkstra, Vito Montalto, and Nicola Pontarollo. 2023. "What makes cities happy? Factors contributing to life satisfaction in European cities." *European Urban and Regional Studies* 30 (4): 319–342. https://doi.org/10.1177/09697764231155335.

Castro, Lucio, and Carlos Scartascini. 2015. "Tax compliance and enforcement in the pampas evidence from a field experiment." *Journal of Economic Behavior & Organization* 116 (August): 65–82. https://doi.org/10.1016/j.jebo.2015.04.002.

Chirico, Michael, Robert P. Inman, Charles Loeffler, John MacDonald, and Holger Sieg. 2016. "An Experimental Evaluation of Notification Strategies to Increase Property Tax Compliance: Free-Riding in the City of Brotherly Love." *Tax Policy and the Economy* 30, no. 1 (January): 129–161. https://doi.org/10.1086/685595.

De Neve, Jan-Emmanuel, Clément Imbert, Johannes Spinnewijn, Teodora Tsankova, and Maarten Luts. 2021. "How to Improve Tax Compliance? Evidence from Population-Wide Experiments in Belgium." *Journal of Political Economy* 129, no. 5 (May): 1425–1463. https://doi.org/10.1086/713096.

Drago, Francesco, Friederike Mengel, and Christian Traxler. 2020. "Compliance Behavior in Networks: Evidence from a Field Experiment." *American Economic Journal: Applied Economics* 12, no. 2 (April): 96–133. https://doi.org/10.1257/app.20170690.

Dwenger, Nadja, Henrik Kleven, Imran Rasul, and Johannes Rincke. 2016. "Extrinsic and Intrinsic Motivations for Tax Compliance: Evidence from a Field Experiment in Germany." *American Economic Journal: Economic Policy* 8, no. 3 (August): 203–232. https://doi.org/10.1257/pol.20150083.

Feld, Lars P., and Bruno S. Frey. 2007. "Tax compliance as the result of a psychological tax contract: The role of incentives and responsive regulation." *Law and Policy* 29 (1). https://doi.org/10.1111/j.1467-9930.2007.00248.x.

Ghanem, Dalia, Sarojini Hirshleifer, and Karen Ortiz-Beccera. 2023. "Testing Attrition Bias in Field Experiments." *Journal of Human Resources* (August): 0920–11190R2. https://doi.org/10.3368/jhr.0920-11190R2.

Giaccobasso, Matias, Brad Nathan, Ricardo Perez-Truglia, and Alejandro Zentner. 2022. *Where Do My Tax Dollars Go? Tax Morale Effects of Perceived Government Spending.* Technical report. Cambridge, MA: National Bureau of Economic Research, February. https://doi.org/10.3386/w29789.

Gibbons, Stephen, and Stephen Machin. 2005. "Valuing rail access using transport innovations." *Journal of Urban Economics* 57, no. 1 (January): 148–169. https://doi.org/10.1016/j.jue.2004.10.002.

Gonzalez-Navarro, Marco, and Climent Quintana-Domeque. 2016. "Paving Streets for the Poor: Experimental Analysis of Infrastructure Effects." *Review of Economics and Statistics* 98, no. 2 (May): 254–267. https://doi.org/10.1162/REST_a_00553.

Hallsworth, Michael, John A. List, Robert D. Metcalfe, and Ivo Vlaev. 2017. "The behavioralist as tax collector: Using natural field experiments to enhance tax compliance." *Journal of Public Economics* 148 (April): 14–31. https://doi.org/10.1016/j.jpubeco.2017.02.003.

Jackson, Matthew O. 2008. *Social and Economic Networks.* Princeton University Press.

Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Soren Pedersen, and Emmanuel Saez. 2011. "Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark." *Econometrica* 79 (3): 651–692. https://doi.org/10.3982/ECTA9113.

Kresch, Evan Plous, Mark Walker, Michael Carlos Best, François Gerard, and Joana Naritomi. 2023. "Sanitation and property tax compliance: Analyzing the social contract in Brazil." *Journal of Development Economics* 160:102954. https://doi.org/https://doi.org/10.1016/j.jdeveco.2022.102954.

Levi, Margaret. 1988. *Of Rule and Revenue.* First Edit. University of California Press.

Lewis, Daniel J., and Karel Mertens. 2022. *A Robust Test for Weak Instruments for 2SLS with Multiple Endogenous Regressors.* Working Paper 2208. Revised 26 Sep 2024. Federal Reserve Bank of Dallas.

Luttmer, Erzo F. P., and Monica Singhal. 2014. "Tax Morale." *Journal of Economic Perspectives* 28, no. 4 (November): 149–168. https://doi.org/10.1257/jep.28.4.149.

Meiselman, Ben S. 2018. "Ghostbusting in Detroit: Evidence on nonfilers from a controlled field experiment." *Journal of Public Economics* 158 (February): 180–193. https://doi.org/10.1016/j.jpubeco.2018.01.005.

Montenbruck, Laura. 2023. *Fiscal Exchange and Tax Compliance: Strengthening the Social Contract Under Low State Capacity.* Working Paper.

Musgrave, Richard A. 1992. "Social contract, taxation and the standing of deadweight loss." *Journal of Public Economics* 49 (3). https://doi.org/10.1016/0047-2727(92)90074-P.

Olea, José Luis Montiel, and Carolin Pflueger. 2013. "A Robust Test for Weak Instruments." *Journal of Business & Economic Statistics* 31 (3): 358–369. https://doi.org/10.1080/00401706.2013.806694. eprint: https://doi.org/10.1080/00401706.2013.806694.

Perez-Truglia, Ricardo, and Ugo Troiano. 2018. "Shaming tax delinquents." *Journal of Public Economics* 167 (November): 120–137. https://doi.org/10.1016/j.jpubeco.2018.09.008.

Pflueger, Carolin E., and Su Wang. 2015. "A Robust Test for Weak Instruments in Stata." *The Stata Journal* (College Station, TX) 15 (1): 216–225.

Pomeranz, Dina, and José Vila-Belda. 2019. "Taking State-Capacity Research to the Field: Insights from Collaborations with Tax Authorities." *Annual Review of Economics* 11, no. 1 (August): 755–781. https://doi.org/10.1146/annurev-economics-080218-030312.

Sandholtz, Wayne Aaron, and Pedro C. Vicente. 2024. *Tax morale, public goods, and politics: Experimental evidence from Mozambique.* Nova SBE Working Paper Series wp671. Universidade Nova de Lisboa, Nova School of Business and Economics.

Slemrod, Joel. 2019. "Tax Compliance and Enforcement." *Journal of Economic Literature* 57, no. 4 (December): 904–954. https://doi.org/10.1257/jel.20181437.

Sobel, Joel. 2005. "Interdependent Preferences and Reciprocity." *Journal of Economic Literature* 43, no. 2 (May): 392–436. https://doi.org/10.1257/0022051054661530.

Sorin, Jeanne. 2025. *Public Roads on Private Lands: Land Costs and Optimal Road Improvements in Urban Uganda.* Working Paper.

Stantcheva, Stefanie. 2021. "Understanding Tax Policy: How do People Reason?" *The Quarterly Journal of Economics* 136, no. 4 (October): 2309–2369. https://doi.org/10.1093/qje/qjab033.

Timmons, Jeffrey F., and Francisco Garfias. 2015. "Revealed Corruption, Taxation, and Fiscal Accountability: Evidence from Brazil." *World Development* 70 (June): 13–27. https://doi.org/10.1016/j.worlddev.2014.12.011.

Wang, Weijie, and Taek Kyu Kim. 2024. "Do government performance signals affect citizen satisfaction?" *Journal of Policy Analysis and Management,* no. 43, 846–870.

Wilson, Nicole E., and Leah R. Rosenzweig. 2023. "Paying for growth or goods: Tax morale among property owners in Lagos." *Journal of Experimental Political Science* 105, no. 4 (July): 1–11. https://doi.org/10.1017/XPS.2023.27.

Zhou, Lingyun. 2024. *WEAKIVTEST2: Stata module to compute Robust Test for Weak Instruments with Multiple Endogenous Regressors.* Statistical Software Components S459310. Boston College Department of Economics.

A Online Appendix

A.1 Derivation of Posterior Beliefs

This appendix provides a formal derivation of the posterior expectations presented in equation (2.5). We proceed by applying Bayes' rule and exploiting the properties of conjugate distributions. The proof proceeds in several steps:

Setup. By Bayes' rule, the posterior distribution is proportional to the product of the likelihood function and the prior distribution:

$$\mathbb{P}(\theta|X_pE_p) \propto \mathbb{P}(X_pE_p|\theta) \times \mathbb{P}(\theta). \tag{A.1}$$

Under our assumptions that streets are paved with equal and independent probability (π), the number of paved streets observed by a property owner (X_pE_p) among the unpaved streets in their route (S_p) follows a Binomial distribution:

$$X_pE_p|\theta \sim \text{Binomial}(S_p, \pi), \tag{A.2}$$

where $\pi = \frac{(1-\theta)G}{S}$. The probability mass function is thus:

$$\mathbb{P}(X_pE_p = k|\theta) = \binom{S_p}{k}\pi^k(1-\pi)^{S_p-k} \tag{A.3}$$

On the other hand, the prior follows a Beta distribution:

$$\mathbb{P}(\theta) = \frac{\theta^{\kappa_{p,1}-1}(1-\theta)^{\kappa_{p,2}-1}}{B(\kappa_{p,1}, \kappa_{p,2})} \tag{A.4}$$

where $B(\cdot, \cdot)$ is the Beta function.

Poisson Approximation. Given G/S small, implying a low probability of a street being paved ($\pi \ll 1$), we can apply the Poisson approximation to the Binomial distribution:

$$\mathbb{P}(X_pE_p = k|\theta) \approx \frac{(S_p\pi)^k e^{-S_p\pi}}{k!} \tag{A.5}$$

Substituting $\pi = \frac{(1-\theta)G}{S}$:

$$\mathbb{P}(X_pE_p = k|\theta) \approx \frac{(S_p\frac{(1-\theta)G}{S})^k e^{-S_p\frac{(1-\theta)G}{S}}}{k!} \tag{A.6}$$

Posterior Distribution. The posterior distribution is proportional to:

$$\mathbb{P}(\theta|X_pE_p = k) \propto \frac{(S_p\frac{(1-\theta)G}{S})^k e^{-S_p\frac{(1-\theta)G}{S}}}{k!} \times \frac{\theta^{\kappa_{p,1}-1}(1-\theta)^{\kappa_{p,2}-1}}{B(\kappa_{p,1}, \kappa_{p,2})} \tag{A.7}$$

For small G/S, we approximate $e^{-S_p \frac{(1-\theta)G}{S}} \approx 1$. After this approximation and rearranging terms:

$$\mathbb{P}(\theta|X_pE_p = k) \propto \theta^{\kappa_{p,1}-1}(1-\theta)^{\kappa_{p,2}-1+k} \tag{A.8}$$

The form of this expression reveals that the posterior follows a Beta distribution with updated parameters:

$$\theta|X_pE_p \sim \text{Beta}(\kappa_{p,1}, \kappa_{p,2} + X_pE_p) \tag{A.9}$$

Posterior Expectation. Using the definition of the expected value of the Beta distribution, the posterior expectation is:

$$\mathbb{E}[\theta|X_pE_p] = \frac{\kappa_{p,1}}{\kappa_{p,1} + \kappa_{p,2} + X_pE_p}. \tag{A.10}$$

Given that $\theta_p^* = \frac{\kappa_{p,1}}{\kappa_{p,1}+\kappa_{p,2}}$ and $\kappa_p = \kappa_{p,1} + \kappa_{p,2}$, this expression can be written as:

$$\theta_p(E_p) = \frac{\theta_p^*}{1 + \frac{X_pE_p}{\kappa_p}}, \tag{A.11}$$

Finally, we apply a first-order approximation to the expectation. First, we rewrite the expectation in the form of a standard power series:

$$\theta_p(E_p) = \theta_p^* \left(1 + \frac{X_pE_p}{\kappa_p}\right)^{-1}. \tag{A.12}$$

Defining $Z_p \equiv \frac{X_pE_p}{\kappa_p}$, we have that when π is small,[21] $|\frac{X_pE_p}{\kappa_p}| < 1$, allowing us to use the expansion:

$$\theta_p(E_p) = \theta_p^*[1 - Z_p + Z_p^2 - Z_p^3 + ...] \tag{A.13}$$

For a first-order linear approximation, we keep only terms up to the first power and drop higher-order terms:

$$\theta_p(E_p) \approx \theta_p^*[1 - Z_p]. \tag{A.14}$$

Substituting and expanding the right-hand side:

$$\theta_p(E_p) \approx \theta_p^* - \frac{\theta_p^* X_p}{\kappa_p} E_p \tag{A.15}$$

The linear approximation error is of order $O((X_pE_p/\kappa_p)^2)$.

21. The equivalence between small E_p and small π follows from the fact that E_p is the observed fraction of paved streets, which has expected value π under our model assumptions.

A.2 (Binary) Dissatisfaction as a Proxy for Beliefs

This appendix shows how using a *binary* measure of dissatisfaction with the local government—as available in our ASLS survey—maps into the theoretical predictions in Section 2.

Latent Dissatisfaction. Let $D_p^*(E_p)$ denote property–owner p's continuous *latent* dissatisfaction, where larger values imply greater unhappiness with the local government. We assume

$$D_p^*(E_p) = \zeta_0 + \zeta_1 \theta_p(E_p) + \nu_p, \tag{A.16}$$

with $\zeta_1 > 0$, so dissatisfaction increases with the (latent) belief that government is inefficient, $\theta_p(E_p)$. For now, suppose $E_p \perp\!\!\!\perp \nu_p$ given the random assignment.

Under the posterior-belief approximation in equation (2.5),

$$\theta_p(E_p) = \theta_p^* - \delta_p E_p, \qquad \delta_p > 0.$$

Substituting gives

$$D_p^*(E_p) - D_p^*(0) = -\zeta_1 \delta_p E_p < 0, \tag{A.17}$$

that is, as long as $\zeta_1, \delta_p > 0$, exposure to more paved streets can only reduce latent dissatisfaction.

Binary Dissatisfaction. We observe neither $D_p^*(E_p)$ nor $\theta_p(E_p)$, but have a binary indicator $D_p = \mathbf{1}\{p \text{ is dissatisfied}\}$ based on self-reported answers to the ASLS survey. Consider a binary response model such that

$$\begin{aligned} D_p &= \mathbf{1}\{D_p^*(E_p) \geq T\} \\ &= \mathbf{1}\{D_p^*(0) - \zeta_1 \delta_p E_p \geq T\}. \end{aligned} \tag{A.18}$$

In words, people respond they are dissatisfied if their latent dissatisfaction is above a threshold T.

Testable Prediction. If $\zeta_1, \delta_p > 0$, a higher exposure E_p *weakly decreases* the probability that a respondent reports being dissatisfied:

$$\frac{\partial \mathbb{P}\big(D_p = 1\big)}{\partial E_p} \leq 0. \tag{A.19}$$

This is the alternative to the testable prediction of belief updating derived from equation (2.6) when using a binary measure of dissatisfaction. Finding a negative effect will be consistent with ζ_1, δ_p being *jointly* positive, simultaneously testing for the validity of the proxy and the sign of δ_p. In contrast, a null effect would be more consistent either with (i) latent dissatisfaction being uncorrelated with beliefs about inefficiency ($\zeta_1 = 0$) or (ii) beliefs being unresponsive to exposure ($\delta_p = 0$).

Who Can Switch from Dissatisfied to Satisfied? Because exposure only lowers D_p^*, individuals whose baseline latent dissatisfaction already lies *below* T will very likely remain satisfied after the intervention. Switching is therefore most likely among those initially above the threshold but not "too far" above it:

$$E_p \ \geq \ \frac{D_p^*(0) - T}{\zeta_1 \delta_p}.$$

Hence, we have a sharper prediction: we expect no change among baseline-satisfied respondents and a decrease in dissatisfaction among a subset of baseline-dissatisfied respondents.

Potential Alternative Channels. A potential confound arises if paving improves satisfaction in other domains (e.g., satisfaction with life in Acayucan), which may in turn be positively correlated with satisfaction with the local government. An alternative concern is that higher exposure generates private benefits—such as reduced travel times or increased property values—which could indirectly reduce dissatisfaction with the government. In either case, the assumption that $E_p \not\perp \nu_p$ would be violated, undermining our interpretation of the results as reflecting belief updating about government inefficiency. In Section 6, we show that exposure has no significant effect on overall satisfaction with life in Acayucan, property values, or self-reported travel times to the city center, helping to alleviate these concerns.

A.3 Assessing Selective Attrition

Our primary sample consists of a balanced panel of properties with consistent ownership from 2005 to 2012. To avoid confounding treatment assignments with ownership changes, we exclude properties from the estimating sample if ownership changed during the experiment. Although we do not directly observe the causes of ownership transfers, they could stem from sales, intra-family transfers, or administrative record discrepancies. This sample selection may introduce biases if the excluded properties systematically differ from those retained or by treatment assignment. For example, if properties with high compliance rates are more likely to change ownership, the remaining sample may be biased toward lower compliance rates. Since treatment affects compliance rates, this bias could vary by treatment assignment. Such selective attrition could potentially compromise both the internal and external validity of our results.

In this section, we conduct a series of robustness exercises to characterize attrition patterns and examine the sensitivity of the results to alternative samples.

Attrition Patterns and Selection Bias. Our initial dataset in 2005 comprised 1,063 plots, of which 752 retained consistent ownership throughout the study period. Notably, 14.4% of the original plots were dropped from the sample before the experiment began in 2007, meaning the actual attrition rate over the experimental period is closer to 17.5%, aligning with typical rates observed in field experiments (Ghanem, Hirshleifer, and Ortiz-Beccera 2023).

We first test for differential attrition rates by treatment assignment and year. The first three columns of Appendix Table A.9 report cumulative attrition rates relative to 2005 and the results of differential attrition tests by treatment assignment. Our findings show no evidence of differential attrition rates by treatment status over time, suggesting that attrition is not systematically related to treatment assignment.

However, equal attrition rates alone do not guarantee internal validity. Even with identical rates, characteristics of dropouts could systematically differ between treatment and control groups. If these differences correlate with the outcome of interest, attributing changes solely to the treatment becomes challenging. Furthermore, external validity may be compromised if the remaining sample systematically differs from the original sample or target population.

To address these concerns, we conduct the two joint tests proposed by Ghanem, Hirshleifer, and Ortiz-Beccera (2023):

1. **Internal Validity for Stayers (IVal-R) test:** This text examines the equality of baseline outcome distributions for assigned treatment and control stayers, as well as assigned treatment and control leavers.

2. **Internal Validity for the Study Population (IVal-P) test:** This test examines the equality of baseline outcome distributions across all four treatment/response subgroups.

Both tests leverage all available baseline data, encompassing both stayers and leavers. Furthermore, Ghanem, Hirshleifer, and Ortiz-Beccera (2023) demonstrate that these testable restrictions are sharp, indicating that they represent the strongest implications that can be tested given the available data.

Columns (4)-(7) of Appendix Table A.9 report baseline (2005) compliance rates across the four groups over time. We observe slightly higher compliance rates among leavers, irrespective of treatment assignment. However, these differences are small and statistically insignificant. Columns (8) and (9) report results for the IVal-P and IVal-R tests, respectively, and we cannot reject the null hypothesis for either test across all years.

The combination of no differential attrition rates and no significant differences in baseline compliance rates across subgroups suggests that our results are not attributable to non-random attrition bias. These findings provide reassurance that attrition does not compromise the internal or external validity of our estimates.

A.4 Tables and Figures

Figure A.1: Street Paving Rollout (2007–2012)

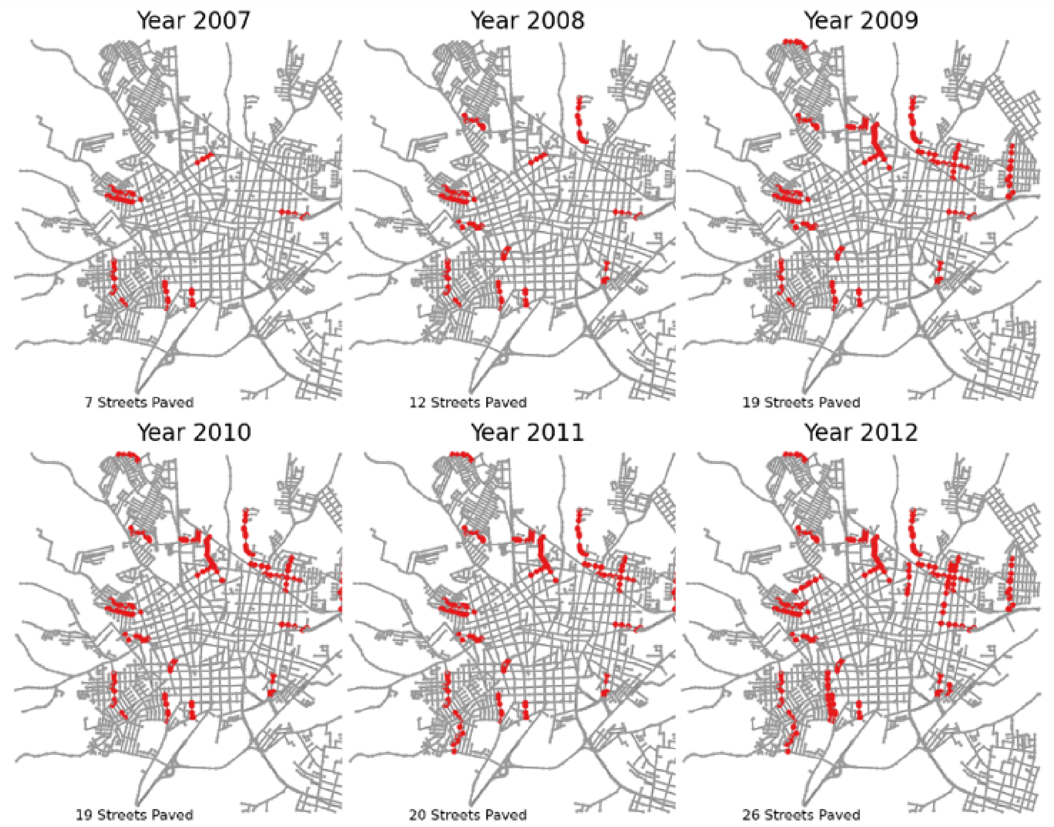

Notes: Authors' calculations based on Acayucan's administrative data, project rollout records, and Open Street Maps data, using the OSMnx library (Boeing 2025). Each red dot represents a node in the street network corresponding to a street paved under the experiment.

Figure A.2: Distribution of ITT and Realized Exposure in Route to City Center

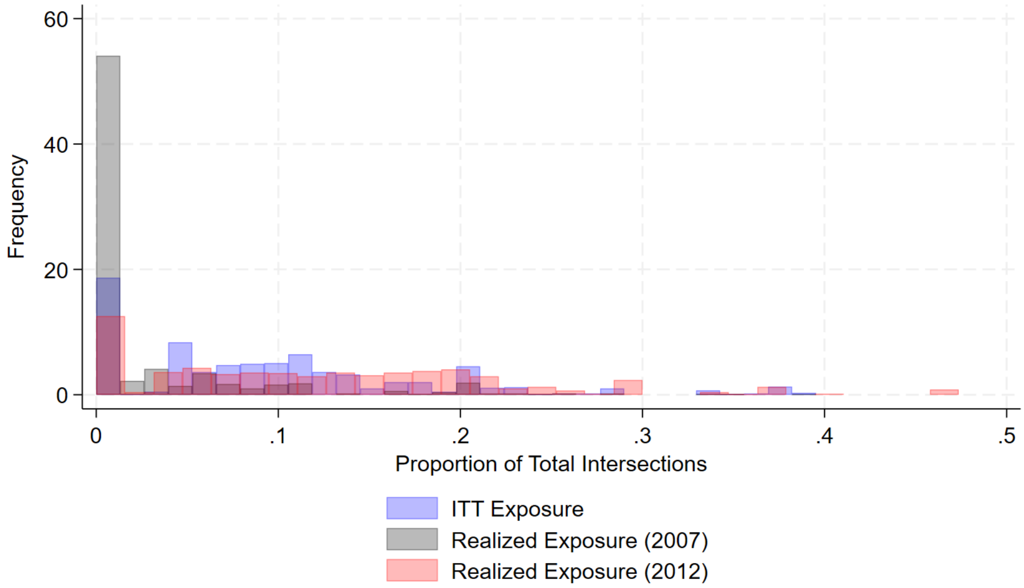

Notes: Author's calculations based on Acayucan's administrative dataa project rollout records. See Section 3 for a detailed description of the construction of the exposure measures.

Table A.1: Probability of Street Paving in Acayucan, Mexico (2007–2012)

	Number of Experimental & Non-Experimental streets Unpaved	Number of Streets Paved by Year	Fraction of Streets Paved
2007	345	7	0.021
2008	338	5	0.014
2009	333	7	0.021
2010	326	0	0.003
2011	326	1	0
2012	325	6	0.018

Notes: Authors' calculations based on Acayucan's administrative data and project rollout. The average probability of a street being paved throughout the 2007–2012 period is 1.3%.

Table A.2: Difference accross baseline variables between the 47 matched projects and the remaining 9 projects

	(1)	(2)	(3)	(4)	(5)
	Mean of the 47 matched projects	Mean of the 9 remaining projects	Difference in Means (Diff.)	Diff. p-value	Observations
Plot-Level ASLS Data and Property Appraisals					
Dissatisfied with the Local Government (0-1)	0.514	0.706	-0.192	0.001	722
	(0.500)	(0.457)	(0.047)		
Dissatisfied with living in Acayucan (0-1)	0.131	0.217	-0.085	0.173	722
	(0.338)	(0.414)	(0.062)		
Log of Property Appraised Value (USD)	11.568	11.836	-0.268	0.199	334
	(0.695)	(0.856)	(0.206)		
Log of Property Owner Value (USD)	11.790	11.862	-0.072	0.825	519
	(0.942)	(1.118)	(0.325)		
Log of Land Appraised Value (USD)	10.149	10.604	-0.455	0.015	334
	(0.477)	(0.862)	(0.182)		
Travel Time to the City Center (minutes)	19.944	20.684	-0.740	0.779	722
	(11.226)	(11.557)	(2.629)		
Per Capita Expenditure (USD)	297.111	378.701	-81.589	0.466	707
	(223.295)	(353.792)	(111.185)		
Plot-Level Administrative Data					
Tax Compliance (0-1)	0.868	0.878	-0.010	0.809	1,364
	(0.338)	(0.329)	(0.042)		
Cadastral Property Value (USD)	39,135.701	32,458.425	6,677.276	0.587	1,364
	(52,776.681)	(40,259.424)	(12,219.024)		
Tax Owed (USD)	36.797	35.037	1.759	0.372	1,364
	(18.797)	(9.411)	(1.956)		
Project Level OpenStreetMap data					
Road Length (Meters)	480.6492	478.1280	2.5212	0.9819	56
	(403.6312)	(291.3161)	(110.5007)		
Katz centrality	0.0168	0.0163	0.0006	0.1804	56
	(0.0008)	(0.0012)	(0.0004)		
Distance to City Center (Meters)	1,686.3083	2,009.8605	-323.5522	0.1530	56
	(488.3613)	(660.5158)	(223.2393)		

Notes: The table compares the means of 13 variables between the 47 matched projects used in the main tax compliance analysis and the remaining 9 projects. Column 1 presents the mean values of the variables for the observations that belong to the 47 projects. Column 2 presents the same mean but for the observations that belong to the 9 remaining projects. Column 3 presents the difference in means and Column 4 the p-value associated to the test of the difference in means. The results indicate no major differences between the two groups. Plots that belongs to the 9 projects excluded from the main analysis show higher dissatisfaction with the local government and higher land appraised values. However, for the other 11 variables, there is no statistical evidence of differences between the groups. Authors' calculations are based on Acayucan's administrative data, the Acayucan Standards of Living Survey (ASLS), and appraisal valuations. Monetary values are reported in 2024 PPP USD. Robust errors, clustered at the project level, are reported in parentheses for the administrative and ASLS data.

Table A.3: Exposure to Paved Streets Along the Route to the City Center

	by front street Assigned to Pav. status ($b^Z_{0\ p,t}$)				**Overall Average**	
	2007		**2012**		**2007**	**2012**
	$b^Z_{0\ p,2007}$=1	$b^Z_{0\ p,2007}$=0	$b^Z_{0\ p,2012}$=1	$b^Z_{0\ p,2012}$=0		
Realized Exposure ($E_{p,t}$)	5.1%	1.0%	14.4%	11.1%	2.57%	12.40%
ITT Exposure (E^Z_p)	15.0%	6.1%	15.0%	6.1%	9.6%	

Notes: The table presents exposure statistics by the front street ITT status for 2007 and 2012 ($b^Z_{0\ p,t} \in (1, 0)$ and $b^Z_{0\ p,2012} \in (1, 0)$), along with the overall average exposure across all experimental observations. Since ITT exposure is time-invariant, its 2007 value is repeated for 2012.

Table A.4: Complier Profiles Based on Assignment of Adjacent Street to Treatment

	Proportionality test	**Lower 95% confidence limit**	**Upper 95% confidence limit**
Plot-Level Administrative Variables (Y_p) in 2006			
Tax Compliance (0-1)	0.996	0.938	1.053
High Property Value (0-1)	0.790	0.664	0.916
High Tax Owed (0-1)	0.794	0.487	1.102
Plot-Level ASLS Variables (Y_p) in 2006			
Dissatisfied with the Local Government (0-1)	1.065	1.006	1.124
Dissatisfied with Living in Acayucan (0-1)	1.220	1.065	1.376
High Property Appraised Value (0-1)	1.064	1.024	1.104
High Property Owner Value (0-1)	1.048	0.992	1.104
High Land Appraised Value (0-1)	1.056	1.014	1.098
High Travel Time to City Center (0-1)	0.772	0.658	0.887
High Per Capita Expenditure (USD)	1.042	0.977	1.108

Notes: The table reports a proportionality test to characterize compliers based on binary baseline covariates Y_p, following the approach of Angrist and Pischke (2011). Specifically, it computes the sample analog of the ratio: $\frac{\mathbb{P}\left(b_{p,0}(1)>b_{p,0}(0)|Y_p=1\right)}{\mathbb{P}\left(b_{p,0}(1)>b_{p,0}(0)\right)}$, where $b_{p,0}(b^Z_{p,0})$ is an indicator for whether property p receives treatment under assignment $b^Z_{p,0} \in \{0,1\}$, and Y_P is a binary baseline covariate. Variables labeled as *High* are equal to 1 if the corresponding continuous variable is above its median, and 0 otherwise. If the values in the first column are above one, compliers are more likely to have the corresponding characteristic Y_p. Bootstrap simulations with 100 replications were used to derive the confidence limits.

Table A.5: Effects of Street Paving on Satisfaction with the Local Government, Excluding Respondents Who Were Either Very Satisfied or Very Dissatisfied at Baseline

	Dissatisfied with the Local Government = 1 (2009)			
	OLS		2SLS	
	(5)	(6)	(7)	(8)
Assigned to Pavement Exposure ($E^Z_{p,t}$)	-0.032	0.014		
	(0.023)	(0.025)		
$E^Z_{p,t} \times D_{p,06}$		-0.102**		
		(0.037)		
Pavement Exposure ($\widehat{E}_{p,09}$)			-0.035	0.016
			(0.025)	(0.028)
$\widehat{E}_{p,09} \times D_{p,06}$				-0.107**
				(0.038)
Dissatisfied in 2006 ($D_{p,06}$)	0.107**	0.225***	0.108**	0.218***
	(0.042)	(0.059)	(0.042)	(0.057)
Observations	585	585	585	585
Mean Dep. Var (2006)	.46	.46	.46	.46
Sum of Coefficients		-.088		-.091
S.E. (Sum = 0)		.033		.033
P-value (Sum = 0)		.01		.005
Generalized First-Stage Statistic			1164.47	218.25
Critical Values 10%			23.11	36.85

Notes: Authors' calculations use ASLS data. The restricted sample (585 out of 722 plots) excludes plots either very dissatisfied or very satisfied with the Local Government in 2006. Robust errors, clustered at the project level (52), are reported in parentheses. The *Dissatisfied* variable takes the value of 1 if the household head was very dissatisfied or dissatisfied, and 0 if they were satisfied or very satisfied. 2SLS uses the *Assigned to Paved* measure and its interaction with dissatisfaction at baseline as instruments for their corresponding *Paved* versions. The exposure measure is standardized by dividing by its standard deviation to ease interpretation. The p-value for testing H_0 : Sum of coefficients = 0 against H_1 : Sum of coefficients $\neq$ 0 measures the probability of observing a test statistic as extreme as (or more extreme than) the one computed from our sample, assuming the null hypothesis is true. The Generalized First-Stage Statistic is the one proposed by Lewis and Mertens (2022), which nests the robust F-statistic or the effective F-statistic from Olea and Pflueger (2013) since there is one endogenous variable. Estimation accounts for survey weights.

Table A.6: Accounting for direct benefits from distant paved streets

	Tax compliance = 1					
	(1)	(2)	(3)	(4)	(5)	(6)
		OLS			2SLS	
δ^Z: Assigned to Pavement						
Exposure $\times$ I_{2007} ($E^Z_{p,t}$)	0.015**	0.016**	0.016**			
	(0.007)	(0.007)	(0.007)			
γ^Z_0: Plot p Assigned						
to Pavement $\times$ I_{2007} ($b^Z_{0,p,t}$)	0.032*	0.031*	0.032			
	(0.019)	(0.019)	(0.019)			
γ^Z_1: Plot p Assigned						
to Pavement $\times$ I_{2007} ($b^Z_{1,p,t}$)		-0.006	-0.006			
		(0.018)	(0.018)			
γ^Z_2: Plot p Assigned						
to Pavement $\times$ I_{2007}($b^Z_{2,p,t}$)			0.003			
			(0.023)			
δ: Pavement Exposure ($\widehat{E}_{p,t}$)				0.026**	0.027*	0.026*
				(0.013)	(0.014)	(0.015)
γ_0: Plot p Paved ($\widehat{b}_{0,p,t}$)				0.048*	0.048*	0.049
				(0.027)	(0.027)	(0.031)
γ_1: Plot p Paved ($\widehat{b}_{1,p,t}$)					-0.003	-0.004
					(0.029)	(0.033)
γ_2: Plot p Paved ($\widehat{b}_{2,p,t}$)						0.004
						(0.048)
Observations	6,016	6,016	6,016	6,016	6,016	6,016
Mean Dep. Var (2006)	0.85	0.85	0.85	0.85	0.85	0.85
Generalized First-Stage Statistic				29.44	21.84	7.63
Critical Values 10%				26.87	28.24	32.68

Notes: The table reports estimates of the direct benefit and exposure effects of street paving on tax compliance for the pooled sample. Robust errors, clustered at the project level (47), are reported in parentheses. Columns 1 and 4 present the same specification and results as in table 3, while columns 2, 3, 5 and 6 also presents the series of indicator variables $b^Z_{n,p}$ and $b_{n,p}$ if the n-th street crossed in the shortest path to the city center is assigned to pavement or paved, respectively. All regressions include plot fixed effects, year fixed effects and baseline controls (property value and the tax bill interacted with year dummies). The exposure measure is standardized by dividing each value by the standard deviation computed within each year (2007–2012), to ease interpretation. The Generalized First-Stage Statistic is the one proposed by Lewis and Mertens (2022), which allows for the assessment of instrument strength when there are two instruments and two endogenous variables under non-homoskedasticity.

Table A.7: Decomposition into Transited and Observed Nodes

	Tax compliance = 1					
	(1)	(2)	(3)	(4)	(5)	(6)
	ITT			LATE		
Exposure to ITT Transited Intersections × I_{2007}	0.021***		0.020***			
	(0.005)		(0.005)			
Exposure to ITT Observed Intersections × I_{2007}		0.013	0.012			
		(0.008)	(0.008)			
Exposure to Paved Transited Intersections				0.031***		0.034***
				(0.009)		(0.010)
Exposure to Paved Observed Intersections					0.020	0.027**
					(0.012)	(0.013)
Observations	6,016	6,016	6,016	6,016	6,016	6,016
Mean Dep. Var (2006)	0.85	0.85	0.85	0.85	0.85	0.85
Generalized First-Stage Statistic				83.53	19.38	11.28
Critical Values 10%				23.11	23.11	26.39

Notes: The table reports estimates of the direct benefits and exposure effects of street paving on tax compliance for the pooled sample. Robust errors, clustered at the project level, are reported in parentheses. We additively separate the principal exposure measure described in section 3 into the exposure to: transited and observed intersections, as exposed in section 6.3. Each exposure measure is standardized by dividing each value by its standard deviation computed within each year (2007–2012), to ease interpretation. The Assigned to Pavement Exposure measure is defined similarly, but uses the streets assigned to be paved. The Generalized First-Stage Statistic is the one proposed by Lewis and Mertens (2022), which allows for the assessment of instrument strength when there are two instruments and two endogenous variables under non-homoskedasticity. In columns (4) and (5), one endogenous variable and one instrument, this statistic corresponds to the robust F-statistic or the effective F-statistic by Olea and Pflueger (2013). All regressions include plot fixed effects, year fixed effects and baseline controls (property value and the tax bill interacted with year dummies).

Table A.8: Effects using alternative destinations and accounting for non-random exposure

	Tax compliance = 1									
Panel A: Exposure	**OLS**					**2SLS**				
In Route to:	City Center	Main Bank	Main Hospital	Main Market	Index	City Center	Main Bank	Main Hospital	Main Market	Index
Plot p Assigned	0.032**	0.033**	0.043***	0.033**	0.031*					
to Pavement ($b_{p,0}$)	(0.019)	(0.020)	(0.017)	(0.019)	(0.019)					
Exposure to Assigned	0.015**	0.014**	0.006	0.014**	0.017**					
to Pavement in Path (E_p)	(0.007)	(0.008)	(0.012)	(0.007)	(0.009)					
Plot p Paved						0.048**	0.058**	0.069***	0.053**	0.052**
($b_{p,0}$)						(0.027)	(0.029)	(0.025)	(0.028)	(0.027)
Exposure to Paved						0.026**	0.017*	0.008	0.020**	0.022**
in Path (E_p)						(0.013)	(0.011)	(0.011)	(0.010)	(0.011)
Observations	6,016	6,016	6,016	6,016	6,016	6,016	6,016	6,016	6,016	6,016
Generalized First-Stage Statistic	.	.	.	.	.	29.442	36.984	15.378	38.235	33.774
Critical Values (10)	.	.	.	.	.	26.868	24.892	27.116	26.942	23.765

	Tax compliance = 1									
Panel B: Exposure re-Centered	**OLS**					**2SLS**				
In Route to:	City Center	Main Bank	Main Hospital	Main Market	Index	City Center	Main Bank	Main Hospital	Main Market	Index
Plot p Assigned	0.032**	0.033**	0.043***	0.033**	0.033**					
to Pavement ($b_{p,0}$)	(0.019)	(0.020)	(0.017)	(0.019)	(0.019)					
Exposure to Assigned	0.015**	0.014**	0.006	0.014**	0.014*					
to Pavement in Path (E_p)	(0.007)	(0.008)	(0.012)	(0.007)	(0.010)					
Plot p Paved						0.049**	0.058**	0.069***	0.054**	0.056**
($b_{p,0}$)						(0.027)	(0.029)	(0.026)	(0.028)	(0.027)
Exposure to Paved						0.025**	0.016*	0.007	0.019**	0.017**
in Path (E_p)						(0.012)	(0.011)	(0.011)	(0.010)	(0.010)
Observations	6,016	6,016	6,016	6,016	6,016	6,016	6,016	6,016	6,016	6,016
Generalized First-Stage Statistic	.	.	.	.	.	31.312	36.987	15.544	38.197	27.562
Critical Values 10%	.	.	.	.	.	26.939	25.011	27.141	27.021	24.731

Notes: The table presents multiple robustness analysis for the main results in Table 3. For all estimates, standard errors, clustered at the project level (47), are reported in parentheses. Panel A presents the main results after changing the destinations that plots takes to other relevant -most common- places in the City Network. For instance, the main Bank (Banamex), the main Hospital (Acayucan's Health Center II), Main Market and an index that corresponds to the geometric mean of the previous locations. Panel B provides the recentered by the expected exposure to paving measures following Borusyak and Hull (2023). The expected exposure is computed as the fraction of unpaved total intersections -in the route to each destination- multiplied by the probability of a street being paved (on average 1.2%). All regressions include baseline property value and the tax bill interacted with year dummies as controls. For details regarding the exposure measures definitions, refer to the main text.

Table A.9: Assessing Internal and External Validity in the Presence of Attrition

Years	Attrition Rates			Baseline (2005) Compliance Rates by ITT				Test of IVal-R	Test of IVal-P
	Share	Differential by ITT	P-value	ITT=1 Stayers	ITT=0 Stayers	ITT=1 Leavers	ITT=0 Leavers	P-value	P-value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
2006	0.105	0.006	0.770	0.875	0.892	0.935	0.939	0.721	0.327
2006-2007	0.144	0.008	0.706	0.872	0.896	0.937	0.900	0.405	0.411
2006-2008	0.182	0.024	0.314	0.871	0.892	0.928	0.918	0.588	0.327
2006-2009	0.211	0.027	0.292	0.878	0.896	0.896	0.898	0.695	0.838
2006-2010	0.252	0.025	0.360	0.871	0.893	0.912	0.910	0.629	0.506
2006-2011	0.268	0.022	0.430	0.868	0.890	0.916	0.916	0.634	0.335
2006-2012	0.293	0.013	0.655	0.868	0.893	0.913	0.908	0.573	0.430

Notes: Authors' calculation using Acayucan administrative data. Table presents the internal validity for the respondent (stayers) and for the population tests as described in Ghanem, Hirshleifer, and Ortiz-Beccera (2023). For a further explanation of the methodology, please refer to section A.3.