Taxing Property in Developing Countries: Theory and Evidence from Mexico*

Anne Brockmeyer
IFS, UCL, World Bank & CEPR

Alejandro Estefan
University of Notre Dame

Karina Ramírez Arras
Secretaría de Hacienda

Juan Carlos Suárez Serrato
Stanford GSB & NBER

September 3, 2023

Abstract

We model and estimate the welfare effects of the property tax—the most under-utilized tax in developing countries. Our model shows property tax hikes impact welfare by reducing compliance and exacerbating liquidity constraints. Enforcement impacts welfare by subjecting non-compliant taxpayers to threats of fines and property seizure. The model yields a sufficient-statistics condition indicating when tax hikes raise welfare more than enforcement. In our setting, administrative data, sharp tax increases, and an enforcement experiment show both policies increase revenue. Tax hikes also raise welfare because revenue gains surpass liquidity costs. However, enforcement reduces welfare as threat costs overshadow revenue increases.

Keywords: property taxation, tax compliance, administrative capacity, liquidity constraints.
JEL codes: H71 , H26 , H21 , O23.

*We are grateful to the Ministry of Finance of Mexico City for granting us access to property tax records, to Arturo Herrera for initiating innovative analytical work on tax compliance in Mexico, and to Roberto Carlos Fernández González, Gabriel Portilla Orozco, Lizbeth Rodríguez Luna, Alejandra Aguirre, Emilio Barriga, and Carlos Cardoso for their guidance. We appreciate helpful comments from Stuart Adam, Hunt Allcott, Peter Arcidiacono, David Agrawal, Michel Azulai, Pierre Bachas, Pat Bayer, David Berger, Michael Best, Augustin Bergeron, Pierre Boyer, Marius Brüllhart, Stefano Curto, Don Davis, Manon Delvaux, Aureo de Paula, Rebecca Diamond, Gilles Duranton, Fernando Ferreira, Lucie Gadenne, François Gerard, Guillermo Gruces, Irem Güçeri, Steve Hamilton, Jessie Handbury, Mariaflavia Harari, Andreas Haufler, Anders Jensen, Matt Johnson, Ben Lockwood, Juliana Londoño Vélez, Mohammed Mardan, Tim McQuade, Jakob Miethe, Enrico Moretti, Joana Naritomi, Oyebola Okunogbe, Ben Olken, Andreas Peichl, David Phillips, Dina Pomeranz, Giacomo Ponzetto, Diego Puga, Sander Ramboer, Pablo Saavedra, Dominik Sachs, Guttorm Schjelderup, Modibo Sidibe, Divya Singh, Joel Slemrod, Chris Timmins, Gabriel Tourek, Nick Tsivanidis, Nancy Wallace, Mazhar Waseem, Jonathan Weigel, Francis Wong, Maisy Wong, Owen Zidar, Eric Zwick, Floris Zoutman, and seminar/conference participants at APPAM, Barcelona Public Economics Meeting, Cities Workshop, CURE, EEA, Georgetown University, George Mason University, George Washington University, Hoover Institution, IAAE, IADB, IFS, IIPF, Imperial College, Johns Hopkins University, Queen Mary University of London, NBER Summer Institute, NBER Public Economics Fall Meeting, NEUDC, Northwestern University, Norwegian School of Economics, the Online Public Finance Seminar, Oxford Center for Business Taxation, Richmond Fed, RIDGE conference, Stanford University (SITE), Syracuse University, UC Berkeley (Real Estate), UCSD, University of Maryland, University of Michigan, University of Munich, University of Notre Dame, University of Warwick, Wharton School, World Bank, ZEW, and the Zurich Public Finance in Developing Countries conference. Magaly Sáenz Somarriba and Meghan Howat provided excellent research assistance. The work was funded by the World Bank, by UKAID from the UK government via the Centre for Tax Analysis in Developing Countries (TaxDev) and through Brockmeyer's UKRI Future Leaders Fellowship (grant reference MR/V025058/1). The findings, interpretations, and conclusions expressed in this work do not necessarily reflect the views of the World Bank, its Board of Executive Directors, or the governments that they represent, nor do they reflect the views of the Ministry of Finance of Mexico City or the Federal Ministry of Finance in Mexico. All errors are our own.

We study the design and effectiveness of the most under-utilized tax in developing countries: the property tax. While lower-income countries generally raise less revenue as a share of GDP than higher-income countries, Figure 1 shows that this pattern is much more pronounced for the property tax than for any other tax. This paper studies whether governments in the developing world can increase property tax revenue to provide local public goods and fund infrastructure investment. We focus on the two main instruments used to raise revenue: tax rate hikes and enforcement. Do these instruments lead to increases in tax revenue? Which instrument is more effective at raising welfare?

We answer these questions by combining administrative tax data, multiple sources of policy variation, and a model of optimal property taxation. Our analyses cover all residential properties in Mexico City—the second-largest city in the Western Hemisphere. In this setting, 40 percent of taxpayers are delinquent on their property taxes. We first evaluate whether the government has the administrative capacity to increase tax collections. Using multiple quasi-experimental tax increases and a field experiment varying enforcement messages delivered to delinquent taxpayers, we find that both tax rate increases and additional enforcement raise tax revenue. Our results show that even in settings with significant under-compliance, governments can raise property tax revenue.

To study the welfare effects of property taxation, we develop a model that formalizes the government's trade-off between raising revenue to provide public goods and the welfare cost of different policy instruments. Tax rate increases can generate additional welfare costs to taxpayers in the presence of liquidity constraints, which are prevalent in developing countries. Property taxes may exacerbate liquidity constraints because they are based on an illiquid stock rather than on a flow of income or consumption. We validate these concerns empirically by documenting imperfect insurance of household consumption to temporary income shocks and by showing that, following a tax increase, taxpayers are more likely to pay in installments. In our model, accounting for the presence of liquidity constraints moderately raises the welfare cost of property taxation. Nonetheless, our model combined with our empirical estimates indicates that tax rate increases are welfare-enhancing under most assumptions on the value of public goods.

An alternative approach to increasing revenue is to more strictly enforce existing taxes on delinquent taxpayers. Our model evaluates the welfare effects of enforcement by comparing the revenue gains from enforcement to the private costs that non-compliant taxpayers experience through threats of fines and property seizures. To quantify this private cost, we use tax rate and enforcement elasticities to compute the tax increase that would raise the same revenue as the enforcement action. While enforcement raises revenue, our estimates imply that the utility cost to non-compliant taxpayers exceeds the welfare gains from the increase in tax revenue. Our empirical estimates therefore imply that, on the margin, a welfare-maximizing government would prefer to raise property tax rates rather than rely on revenues from additional enforcement.

Our analysis proceeds in four steps. We first construct a simple model of optimal policy design, building on Keen and Slemrod (2017). The model is accessible, can be tied to our empirical estimates, and speaks to policymakers' desire to raise revenue while limiting taxpayer hardship.

The model studies optimal tax rate and enforcement policies in a context with liquidity constraints. When households are subject to liquidity constraints, their consumption may decline when they pay property taxes. While the government observes the tax liability, households may not comply with the tax due to liquidity constraints or low tax morale.

The model provides two guiding insights for policy. First, liquidity constraints increase the welfare cost of taxation and yield lower optimal tax rates. Second, while enhanced enforcement raises tax revenue, it also lowers welfare by increasing the private costs of tax delinquency through threats of fines and property seizure. These insights imply that a full welfare analysis of the property tax system requires empirical estimates that capture the costs and benefits of relying on different policy instruments, including tax and compliance elasticities, consumption changes, and private costs of enforcement.

In our second step, we estimate the sufficient statistics that allow us to implement our model and evaluate the welfare effects of tax rate changes and enhanced enforcement. Our analysis leverages large quasi-experimental tax hikes, which affect properties in specific cadastral value bands. Figure 2 shows that depending on the year, tax rates increased by between 18 to 47 percent. We use a regression discontinuity (RD) design to estimate short-term responses and a difference-in-differences (DiD) design to estimate medium-term responses. Both research designs yield similar estimates and are bolstered by a number of checks, including that cadastral values are not manipulated, that property characteristics do not change discontinuously around band thresholds, and that treated and never-treated properties have similar pretrends. We also obtain similar results when running panel regressions that exploit variation in tax rates across the full range of value bands. While our estimates show that governments can raise revenue, we also find that tax hikes significantly reduce compliance with the property tax.

We then study whether enforcement actions can succeed at raising revenue from delinquent taxpayers. To deal with the fact that 40 percent of taxpayers were delinquent on their property taxes, the government of Mexico City applies fines and interest to late payments, occasionally seizes the property of delinquent taxpayers, and conducts enforcement campaigns to encourage compliance. We evaluate the effects of one such campaign using a field experiment in which the tax authority sent enforcement letters emphasizing sanctions and fines to 80,000 delinquent taxpayers. We compare delinquent taxpayers who received an enforcement letter to a passive control group, since this is the policy-relevant comparison for evaluating enforcement. We find that the treatment tripled taxpayers' likelihood of making a payment. These empirical results demonstrate that the government is able to raise property tax revenue through either tax rate increases or enforcement. However, our model also considers that these policy tools can cause hardship, either by exerting private costs or by exacerbating household liquidity constraints.

The third step of our analysis provides two pieces of evidence that liquidity constraints affect the ability of households to comply with the property tax. First, we use standard methods from the consumption literature to show that, in Mexico City, household consumption is only imperfectly insured against temporary idiosyncratic income shocks. The same methods also show that property tax payments are sensitive to temporary income shocks. Second, we show that tax hikes affect taxpayers' choice of payment modality, conditional on payment. Since households that pay their annual tax liability in full and before a given date receive a discount on their liability, the choice to pay in installments is equivalent to taking out a loan from the government. We exploit an unusual inflection point in the tax rate schedule to show that higher tax rates are associated with a higher likelihood of taxpayers paying in installments instead of paying their annual liability all at once. Conditional on attempting to pay in installments, the likelihood that a taxpayer completes all installment payments is also inversely related to the tax rate. We complement this evidence using a DiD design that shows that tax hikes increased taxpayers' propensity to pay in installments.

The fourth and final step of our analysis implements our model of optimal property taxation and provides guidance to policymakers based on our empirical results. The model shows that even though compliance is low, raising revenue through stricter enforcement reduces welfare. This result follows from the fact that the disutility that delinquent taxpayers face from additional enforcement exceeds the utility gains from additional revenue.[1] By contrast, raising property tax rates to provide public goods can increase welfare. In our model, when we account for the presence of liquidity constraints, tax rate increases may lead to consumption declines for constrained households, which increases the welfare cost of taxation. Empirically, the effects of tax hikes on the welfare cost of taxation are small when considering plausible welfare costs of consumption drops and the fact that property taxes constitute a small share of overall household consumption. Our model and empirical estimates therefore show that existing tax rates are below the optimal tax rate.

Overall, this paper shows that developing countries can increase their reliance on property taxation to provide public goods and fund infrastructure investment. Although additional enforcement increases tax revenue, we find that marginal enforcement actions reduce welfare. On the contrary, because tax hikes increase both tax revenue and welfare, tax rate increases are more desirable in our setting. Our model provides guidance for governments by using sufficient statistics to determine whether tax hikes or enforcement are desirable in a given setting and highlights the importance of accounting for revenue effects and welfare costs of property tax policies.

Our results are internally valid for Mexico City and should be interpreted as the effects of large interventions. The tax hikes that we analyze were part of reforms that coincided with a 36 percent increase in property tax revenue (see Figure B.1, Panel A). Our rigorous evaluations of these reforms show that a large part of this increase can be attributed to the causal impacts of tax increases. Similarly, by contacting 80,000 taxpayers, our field experiment reached close to 14 percent of delinquent taxpayers. Finally, because our enforcement intervention was part of a regular enforcement campaign, our results can be interpreted as in-equilibrium effects of existing policies.

[1] This result is consistent with Allcott et al. (2022) and Allcott and Kessler (2019), who model interventions such as ours and show that they can lower welfare even if they succeed at changing behavior. Our result is robust to under-weighting the disutility that delinquent households experience through enforcement.

Mexico City is a very useful laboratory for studying property taxation in a developing country context. Tax administrations in lower-income countries display weaker capacity. For example, the number of tax audits per capita—a widely used proxy for tax administration capacity—increases with GDP per capita (see Panels A1 and A2 of Figure B.2). Mexico fits this pattern well, exhibiting a level of administrative capacity similar to that of other middle-income countries. Households in lower-income countries also suffer from liquidity constraints. The share of households that have a bank or credit card accounts or that can access loans to finance unexpected expenses also increases with GDP per capita.[2] As with tax capacity, measures of liquidity constraints in Mexico City are broadly representative of places with similar income levels, suggesting that our results are likely externally valid for other developing countries. For example, fewer than 20 percent of households in Mexico City have access to a credit card.

This paper contributes to the emerging literature on taxation in developing countries (Besley and Persson 2013, Pomeranz and Vila-Belda 2019). This literature has argued that the optimal mix of tax instruments may diverge from traditional public finance theory prescriptions in a context of limited enforcement capacity (Best et al., 2015). We build on this literature as well as on the tax systems approach of Slemrod and Gillitzer (2013) to think holistically about the revenue and welfare effects of different policy instruments.[3] Similar to Basri et al. (2021), we compare the revenue and welfare effects of tax rate and tax administration changes. While they show that tax administration investments can be superior to tax increases, our model shows that the welfare trade-offs in the case of income taxes differ from those of property taxes. In particular, our enforcement treatment generates a welfare cost by subjecting household to threats of fines or property seizure, while the administrative investment in Basri et al. (2021) gives firms access to a medium taxpayer office meant to reduce compliance costs. In our setting, these forces lead us to conclude that tax rate increases are more effective than enforcement actions at raising welfare.[4]

Our paper is also related to the tax compliance literature (see Slemrod 2019 for a recent survey). This literature has traditionally been concerned with the accurate reporting of liabilities and the impact of detection but has recently noted that payment enforcement is a separate and similarly important challenge, even in the US (Versprille, 2020). Recently, researchers have started evaluating novel tools for the enforcement of outstanding payments (Organ et al., 2022; Kessler, 2020; Perez-Truglia and Troiano, 2018; Dusek et al., 2022). Our optimal tax model considers the welfare gains from policies that target delinquency relative to the gains from other tax policies.

In addition, this study contributes to research that uses property taxes to study questions related

[2]Households in developing countries experience more income volatility (Morduch 1995), have less access to insurance (Jack and Suri 2014; Townsend 1995), and have less access to the formal financial system (Morduch and Karlan 2009; Demirguc-Kunt et al. 2017). Table B.1 describes consumer debt in Mexico City, and Panels B1-B4 in Figure B.2 compare measures of liquidity constraints in Mexico with countries around the world.

[3]Our approach is related to Meiselman (2018) and Brockmeyer et al. (2019), who build on Keen and Slemrod (2017) to examine the effectiveness of enforcement letters to taxpayers in Detroit and Costa Rica, respectively.

[4]Relative to the set of papers that focus on the role of information in enforcing taxes (Pomeranz, 2015; Naritomi, 2019), this paper studies compliance in a setting where the government has full information.

to taxation and development (e.g., Khan et al., 2016, 2019; Okunogbe, 2021; Weigel, 2020; Balan et al., 2022). Best et al. (2020) study the implications of horizontal inequities for tax morale and compliance. Bergeron et al. (2023) study the effects of tax rates and enforcement on compliance and how these instruments jointly determine the revenue-maximizing rate. Consistent with our results, they find that responses to tax rate changes are likely driven by liquidity constraints. Dzansi et al. (2022) also show that liquidity constraints influence households' propensity to pay the property tax in Ghana. Our paper shows that governments in developing countries can improve the design of the property tax by accounting for both revenue and welfare effects of tax and enforcement policies.[5]

Finally, there is a large body of work on property taxes in the US, reviewed in Agrawal et al. (2020).[6] While this paper emphasizes aspects of property taxation that are more salient in developing countries, our results may be applicable to developed country settings. Cabral and Hoxby (2023) show that property taxes are less popular when households lack escrow accounts to smooth tax payments. Wong (2020) also shows that small property tax increases in the US can lead to financial hardship, including mortgage delinquencies and declines in consumption. Similarly, several US cities have trouble collecting property taxes and feature noncompliance rates above 10 percent (Chirico et al., 2019).

The rest of the paper is structured as follows. Section 1 presents our model of optimal property taxation. Sections 2 and 3 describe property taxation in Mexico City and our administrative tax data. We study the effect of tax rate changes on tax revenue in Section 4, the effect of enforcement on revenue in Section 5, and the role of liquidity constraints in Section 6. In Section 7, we implement the model using our empirical estimates and discuss policy implications. Section 8 concludes.

1 Optimal Property Tax Administration with Liquidity Constraints

This section develops a model of optimal property tax administration that considers both tax rate and enforcement policies building on the work of Keen and Slemrod (2017). To match our empirical setting, the model features an observed and fixed tax liability and focuses on compliance along the extensive margin.[7] Since the property tax liability is not tied to a cash flow, the model allows for taxpayer liquidity constraints to impact tax compliance and the welfare cost of taxation. The model delivers expressions for the welfare effects of changes in tax rates and enforcement that are functions of estimable elasticities and yields a condition that determines when tax rate increases are preferable to enhanced enforcement.

[5] Appendix J estimates the revenue-maximizing tax rate, which is significantly greater than existing tax rates.

[6] A concern in this literature is the impact of property taxes on the real estate market. We document that new construction is not designed to target property tax thresholds and that tax increases are not likely to impact investment in existing housing units. Property taxes are unlikely to affect household location decisions, as school funding is not tied to neighborhood-level taxes and internal migration is lower in lower-income countries. This is especially true in our setting, as Mexico City offers unique amenities.

[7] In contrast, taxpayers in the model of Keen and Slemrod (2017) can take costly actions to "hide" taxable income. Table C.1, Panel B, shows that most taxpayers either pay their tax in full or do not pay; very few pay partially.

1.1 Model Setup

Households i live two periods. They consume a private good c and a public good g and have uncertain income y in the first period. Households start owning a property of value H, pay a tax t in the first period, and use the liquidated value of the asset for consumption in the second period. These assumptions represent an initial state where households have committed to a level of housing consumption and a second period where households re-optimize housing and consumption.[8]

While the government observes the tax liability, households may decide not to pay the property tax. $\mathbb{I}[\text{Delinquent}_i]$ denotes the event of household i being delinquent. Households face two types of costs when they are delinquent. First, they incur a private or psychological cost of non-payment $M_i(\alpha) = m(\alpha) + \varepsilon_i$. The idiosyncratic disutility from not paying taxes $\varepsilon_i \sim F(\cdot)$ captures variation in tax morale.[9] The common cost of non-payment, $m(\alpha)$, measures the utility cost taxpayers experience from threats of fines and property seizure and is an increasing function of the enforcement level α.[10] Second, when households are delinquent, the value of their property depreciates by a factor z. This factor corresponds to back taxes collected by the government. Households solve the problem:

$$\begin{aligned} \max_{s,\text{Delinquent}} \quad & u(c) + \beta u(c') + v(g) - M_i(\alpha) \times \mathbb{I}[\text{Delinquent}_i] \\ c &= y - tH(1 - \mathbb{I}[\text{Delinquent}_i]) - s \\ c' &= s(1+r) + H(1 - z \times \mathbb{I}[\text{Delinquent}_i]) \quad \text{and} \quad s \geq 0, \end{aligned}$$

where savings receive an interest rate r and the level of public goods g is set by the government.

Conditional on the decision $j \in \{\text{Pay}, \text{Delinquent}\}$, optimal consumption is determined by the first-order condition:

$$u'(c_s^j) \geq (1+r)\beta u'(c_s'^{,j}),$$

where we index consumption by savings, s. When $s > 0$, this equation holds with equality. When $s = 0$, the household is liquidity constrained, and this equation holds with a strict inequality.

Households decide whether to comply with the property tax by comparing indirect utilities from paying and being delinquent. Denoting the mean indirect utility of payment decision j by V^j, the overall utility for household i is then $V_i = \max\{V^{\text{Pay}}, V^{\text{Delinquent}} + \varepsilon_i\}$. Finally, let V denote the population expectation over V_i.

Let $N^{\text{Pay}} = \mathbb{P}r(V^{\text{Pay}} > V^{\text{Delinquent}} + \varepsilon_i)$ be the fraction of households that pay the tax. Following Busso et al. (2013), we note that $\frac{dV}{dV^j} = N^j$. This expression shows that because households have

[8] We present a simple model for clarity of exposition and discuss extensions in Section 1.5.

[9] Singhal and Luttmer (2014) discuss the importance and drivers of tax morale.

[10] The enforcement letters we study emphasize potential costs and legal actions to non-compliant households. Glaeser (2006) argues that these type of interventions act as emotional taxes—a cost that needs to be taken into account in welfare evaluations (see also Caplin, 2003; Loewenstein and O'Donoghue, 2006). In their taxonomy of the welfare effects of nudges, Allcott and Kessler (2019) also consider information, reminder, and warm glow effects. As we discuss in Appendix A.2, the intervention in our setting does not feature these effects. Meiselman (2018), Organ et al. (2022) and Butera et al. (2022) also consider the role of private costs and public recognition in welfare analyses.

already optimized over being delinquent or paying, the overall effect on welfare from marginal changes to mean utilities does not depend on changes in delinquency status (i.e, $\frac{dN^{\text{Pay}}}{dV^{\text{Pay}}}$).

Consider now the effects of taxes on mean utilities:

$$\frac{\partial V^{\text{Delinquent}}}{\partial t} = 0 \quad \text{and} \quad \frac{\partial V^{\text{Pay}}}{\partial t} = -u'(c_s^{\text{Pay}})H.$$

By the first expression, tax increases do not impact the mean utility of being delinquent. While the second equation applies both when $s = 0$ and when $s > 0$, the interpretation differs across these cases. When $s > 0$, the envelope theorem holds, and households readjust their savings, but the welfare effect of this readjustment cancels out. When $s = 0$, there is no such readjustment, as households are constrained from consuming from their future wealth. Because marginal utility is higher when $s = 0$, the welfare cost of raising taxes is also higher in this case.

Consider now the effects of increasing enforcement by raising α:

$$\frac{\partial V^{\text{Delinquent}}}{\partial \alpha} = -m'(\alpha) \quad \text{and} \quad \frac{\partial V^{\text{Pay}}}{\partial \alpha} = 0.$$

The last expression shows that tax-paying households are not directly affected by changes in enforcement. In contrast, delinquent households experience a drop in their mean utility as the private cost of not paying taxes increases with enforcement. Note, however, that the welfare costs of enforcement do not depend on consumption changes since households that pay taxes in response to enforcement are on the margin between compliance and delinquency. Hence, their change in status does not have a first-order effect on welfare.

A welfare-maximizing social planner would incorporate the private costs of non-compliant households when evaluating the welfare effects of different policies. However, as in Hebous and Keen (2021), it is possible to deviate from the principle of horizontal equity and allow for the possibility that the welfare costs of non-compliant taxpayers may be less valued by society than other costs of property taxes. To incorporate this "fairness perspective," we multiply the term $\frac{\partial V^{\text{Delinquent}}}{\partial \alpha}$ by $\omega \leq 1$, which potentially down-weights this welfare cost.

Finally, the government uses tax revenue to provide public goods g and enforcement α at a cost $a(\alpha)$. Its budget constraint is:

$$g + a(\alpha) = tHN^{\text{Pay}} + zH(1 - N^{\text{Pay}}),$$

where z represents the back-taxes that the government will eventually collect from delinquent taxpayers. While the effects of policy changes on welfare do not depend on changes in the decision to pay or be delinquent, the effects of t and α on N^{Pay} are crucial drivers of tax revenue.

1.2 Optimal Property Tax

The government maximizes V subject to its budget constraint. To simplify exposition, we consider the case where income can take two values, one where households are liquidity constrained, $s = 0$, and one with positive saving, $s > 0$. Substituting the government's constraint for g, a tax increase has the following impact on welfare:

$$-N^{\text{Pay}}H[\pi_s^{\text{Pay}}u'(c_s^{\text{Pay}}) + \pi_0^{\text{Pay}}u'(c_0^{\text{Pay}})] + v'(g) \times \left\{HN^{\text{Pay}} + (t-z)H\frac{\partial N^{\text{Pay}}}{\partial t}\right\},$$

where π_s^{Pay} denotes the share of households who pay taxes with $s = 0$ and $s > 0$.

We express the impact on revenue in the second term by defining the tax elasticity of compliance $\varepsilon_t^{\text{Pay}} = \frac{\partial N^{\text{Pay}}}{\partial t}\frac{t}{N^{\text{Pay}}} < 0$ and letting $\tilde{z} \leq 1$ be the fraction of back-taxes collected in the future. The effect on tax revenue is then $HN^{\text{Pay}}\{1 + (1-\tilde{z})\varepsilon_t^{\text{Pay}}\}$.[11] As in Chetty (2006), we approximate the marginal utility of consumption in the first term with $u'(c)(1 - \gamma\Delta c)$, where $\gamma = \frac{-u''(c)c}{u'(c)}$ is the coefficient of relative risk aversion and Δc measures the average decrease in consumption in response to the property tax increase.[12] We can then manipulate this expression to obtain:

$$MVPF_t = \underbrace{\frac{v'(g)}{u'(c)}}_{\text{Value of Public Goods}} - \underbrace{\frac{1-\gamma\Delta c}{1+(1-\tilde{z})\varepsilon_t^{\text{Pay}}}}_{\text{Welfare Cost Per Dollar of Revenue}}. \tag{1}$$

Equation 1 shows that the marginal value of public funds (MVPF) (e.g., Atkinson and Stern, 1974; Slemrod and Yitzhaki, 2001; Hendren, 2016) from raising taxes depends on the value of public goods and the distortions associated with raising revenue. The first term measures the welfare gain of using \$1 of consumption c to increase the provision of the public good g. The second term measures the welfare cost per dollar of revenue.[13] Absent liquidity constraints, this term takes the familiar inverse-elasticity form $\frac{1}{1+(1-\tilde{z})\varepsilon_t^{\text{Pay}}}$. The costs of raising tax revenue are larger when taxpayers are less likely to comply with a tax increase (large $\varepsilon_t^{\text{Pay}}$ in absolute value) and when the government is less able to collect back-taxes in future periods (low value of $\tilde{z}$). Since the property tax is a tax on a stock rather than on a flow of income or consumption, the welfare cost per dollar also includes the factor $(1 - \gamma\Delta c)$, which incorporates interactions with liquidity constraints. The coefficient $\gamma \geq 0$ captures how costly consumption changes are for taxpayers. The welfare cost of property

[11] Tax hikes increase revenue as long as $\varepsilon_t^{\text{Pay}} > -\frac{1}{1-\tilde{z}}$. As with other welfare models of tax administration (e.g., Keen and Slemrod, 2017; Basri et al., 2021), our model relies on first-order effects of tax changes and does not depend on cross-derivatives between tax rate changes and enforcement, which can be present in models of revenue maximization (e.g., Bergeron et al., 2023). In Appendix J, we consider the case of revenue-maximization and show that current tax rates are below the revenue-maximizing rate.

[12] Note that Δc depends on the share of consumption by households that pay taxes with different values of s. See Appendix A for details.

[13] Equation 1 captures the welfare effect of a budget-balanced policy that raises tax revenue to provide public goods. As Hendren and Sprung-Keyser (2020) show, it can also be valuable to consider the MVPF of policies without taking into account a balanced budget. The MVPF of raising tax revenue through tax hikes is captured in Equation 1 by the welfare cost per dollar of revenue.

taxation increases when tax rate increases lead to larger drops in consumption (more negative Δc). A positive $MVPF_t$ implies that the welfare gains from increasing the provision of public goods exceed the welfare cost of raising tax rates.

We obtain an expression for the optimal property tax by setting $MVPF_t = 0$. Writing $\Delta c = -\eta_t^c t$, where $\eta_t^c > 0$ is the average tax semi-elasticity of consumption, the optimal tax rate is:

$$t = \frac{\left[1 + (1 - \tilde{z})\varepsilon_t^{\text{Pay}}\right] \frac{v'(g)}{u'(c)} - 1}{\gamma \eta_t^c}. \tag{2}$$

The optimal tax is larger when public goods are more valuable $\left(\text{a larger value of } \frac{v'(g)}{u'(c)}\right)$ and when taxes have a smaller effect on delinquency ($\varepsilon_t^{\text{Pay}}$ is close to zero). Similarly, because the average consumption response is likely to be greater when there is a larger share of liquidity-constrained households, the government can set higher property taxes when the fraction of liquidity-constrained households is smaller.

1.3 Optimal Enforcement

Consider now the government's choice to expend resources on enforcement. Increasing α has the following effect on welfare:

$$-(1 - N^{\text{Pay}})\omega m'(\alpha) + v'(g) \times \left\{(t - z)H\frac{\partial N^{\text{Pay}}}{\partial \alpha} - a'(\alpha)\right\}.$$

In contrast to the effect of changing tax rates, the welfare effects of changes in α do not depend on consumption changes.[14] Defining the enforcement elasticity of compliance $\varepsilon_\alpha^{\text{Pay}} = \frac{\partial N^{\text{Pay}}}{\partial \alpha}\frac{\alpha}{N^{\text{Pay}}} > 0$, the MVPF from enforcement is then:

$$MVPF_\alpha = \frac{v'(g)}{\omega m'(\alpha)} - \frac{1 - N^{\text{Pay}}}{(1 - \tilde{z})N^{\text{Pay}}\frac{\varepsilon_\alpha^{\text{Pay}} Ht}{\alpha} - a'(\alpha)}. \tag{3}$$

The first term in this expression is the value of public goods relative to the welfare cost of enforcement. This value is higher when the welfare costs of enforcement $m'(\alpha)$ are low or when these costs are not valued by society (low ω). The second term is the welfare cost of raising revenue from a $1 - N^{\text{Pay}}$ fraction of households. The term $(1 - \tilde{z})\frac{\varepsilon_\alpha^{\text{Pay}} Ht}{\alpha}$ captures the revenue gains from enforcement net of the reduction in future back taxes, $\tilde{z}$.[15] We can use Equation 3 to evaluate whether an

[14] To see this, note that the welfare costs of enforcement are borne by non-compliant households, whose consumption is not affected by taxes. While some households become compliant in response to the enforcement, the welfare effect of this change is second order since these households were indifferent between delinquency and paying taxes.

[15] Our model assumes that changes in α do not impact $\tilde{z}$. This assumption is supported by Figure C.2, which shows that the share of back-taxes eventually collected by the government is stable over time. Moreover, to the extent that enforcement collects "low-hanging fruit" (such that $\tilde{z}$ on the remaining tax debt is smaller), this assumption errs on the side of making enforcement look more effective. We do not study the effects of policies that directly change $\tilde{z}$ (such as increasing property seizures) since we do not observe such policy variation in the data.

enforcement action improved welfare, i.e., whether $MVPF_\alpha > 0$.[16]

Importantly, while the welfare cost of enforcement $m'(\alpha)$ is unobserved, we can infer it using tax rate and enforcement elasticities. To do so, we consider a simultaneous change in taxes and enforcement that leaves compliance N^{Pay} unchanged.[17] This implies that we can write:

$$m'(\alpha) = u'(c) \times \frac{\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha)}{-\varepsilon_t^{\text{Pay}}}. \tag{4}$$

Intuitively, the marginal welfare loss from enforcement is equal to the utility cost associated with a tax rate increase that collects as much revenue as the enforcement intervention. Equation 4 shows that the marginal welfare loss from enforcement is smaller when taxes are more likely to distort compliance behavior (i.e., when $\varepsilon_t^{\text{Pay}}$ is more negative). For example, if $\varepsilon_t^{\text{Pay}} = -0.5$ such that a 10% tax rate increase raises tax payments by 5%, the welfare cost of enforcement in units of marginal utility, i.e., $\frac{m'(\alpha)}{u'(c)}$, is twice the revenue raised through enforcement. However, if $\varepsilon_t^{\text{Pay}} = -1.5$ such that tax rate increases are very distortionary, the welfare cost of enforcement of every additional peso collected would only be 0.66 pesos.

1.4 Relative Value of Tax Rate Increases and Enforcement

To evaluate whether the government should rely on tax hikes or enhanced enforcement to raise revenue, we need to consider a scenario in which both policies have similar effects on tax revenue. Suppose then that the government increased taxes and reduced enforcement to keep spending g constant. In Appendix A.3, we show that such a reform would improve welfare if:

$$MVPF_t - MVPF_\alpha \times \omega \frac{m'(\alpha)}{u'(c)} > 0. \tag{5}$$

Using Equations 3 and 4 implies that:

$$MVPF_\alpha \times \omega \frac{m'(\alpha)}{u'(c)} \;=\; \frac{v'(g)}{u'(c)} - \frac{\omega(1 - N^{\text{Pay}})}{(1-\tilde{z})N^{\text{Pay}}\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha) - a'(\alpha)} \times \frac{\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha)}{-\varepsilon_t^{\text{Pay}}}. \tag{6}$$

As with $MVPF_t$, the first term in this expression measures the welfare gain of using \$1 of consumption c to increase the provision of the public good g. The second term measures the welfare cost of tax payments, which, except for ω, is a function of empirical quantities. Substituting Equations 1 and 6 into Equation 5 shows that the welfare effects of this balanced-budget reform only depends on the welfare costs of raising revenue and does not depend on the value of public goods.

To develop intuition, consider a simplified version of Equation 5 that assumes that the ad-

[16] Since enforcement has both private and administrative costs, perfect compliance may not be an optimal outcome for the government. In Appendix A.4, we also characterize the level of enforcement that maximizes social welfare.

[17] This strategy is similar to that of Bertrand et al. (2010). Allcott and Taubinsky (2015) label this approach the equivalent price metric. See Appendix A.2 for details.

ministrative costs of enforcement are small relative to the effects of the intervention (i.e., that $\frac{a'(\alpha)}{\varepsilon_{\alpha}^{\text{Pay}}(Ht/\alpha)} \approx 0$):[18]

$$MVPF_t - MVPF_\alpha \times \omega \frac{m'(\alpha)}{u'(c)} = -\frac{1-\gamma\Delta c}{1+(1-\tilde{z})\varepsilon_t^{\text{Pay}}} + \frac{\omega(1-N^{\text{Pay}})}{N^{\text{Pay}}} \times \frac{1}{-(1-\tilde{z})\varepsilon_t^{\text{Pay}}}$$

A welfare-maximizing government would then prefer to raise taxes rather than increase enforcement as long as:

$$\varepsilon_t^{\text{Pay}} > -\frac{1}{(1-\tilde{z})} \frac{\omega(1-N^{\text{Pay}})}{(1-\gamma\Delta c)N^{\text{Pay}} + \omega(1-N^{\text{Pay}})}.$$

Focusing on static tax revenue (so that $\tilde{z} = 0$), absent liquidity constraints ($\Delta c = 0$), and assuming $\omega = 1$, this expression shows that taxes are preferable to enforcement as long as the tax elasticity of compliance $\varepsilon_t^{\text{Pay}} > -(1-N^{\text{Pay}})$, the rate of non-compliance. In our setting, since $N^{\text{Pay}} = 0.6$, tax rates would be preferable as long as $\varepsilon_t^{\text{Pay}} > -0.4$.[19] If we instead assume that the government values the utility of non-complying taxpayers half a much as complying taxpayers, tax rates would be only be preferable if $\varepsilon_t^{\text{Pay}} > -0.25$. The presence of liquidity constraints (i.e, $\Delta c < 0$) would further reduce the desirability of taxation. This result also implies that, by providing liquidity, the government can increase the $MVPF_t$ as well as the relative value of taxation over enforcement. We return to this discussion in Section 7, where we implement Equation 5 using our empirical estimates.

1.5 Limitations and Extensions

Our model simplifies the analysis to focus on the main forces in our empirical setting. We now discuss possible limitations and extensions to our conceptual framework.

1. Uncertainty. Our model assumes a discrete distribution for income shocks. Allowing for a range of possible incomes would not alter the main results. In this case, the average consumption drop across the range of possible incomes would measure the welfare costs from taxation.

2. Dynamics. Our model can also be extended to allow for multiple time periods, and the importance of liquidity constraints carries over to dynamic models with uncertain income.[20] Moreover, as in other sufficient statistic models (e.g., Chetty and Finkelstein, 2013), consumption changes are robust welfare measures in the presence of dynamic considerations.

3. Consumption commitments and housing choice. We assume housing consumption is initially fixed. Chetty (2004) shows that, when homeowners do not re-optimize their housing choice, the local value of γ is larger such that consumption commitments exacerbate the welfare costs of property-

[18] While this assumption biases this condition in favor of enforcement, it may be reasonable given the high cost-effectiveness of the kind of intervention we study.

[19] In developing country contexts where $\varepsilon_t^{\text{Pay}}$ is so negative that the government is on the wrong side of the "Laffer curve," as in Bergeron et al. (2023), this condition implies that enforcement is likely to be a preferable policy.

[20] For example, Deaton (1991) shows that even in periods when households are not constrained, precautionary savings produce behavior similar to that under liquidity constraints. This shows that interactions between property taxes and liquidity constraints continue to be important in dynamic settings.

tax-driven consumption drops. Following Chetty (2004), our empirical implementation considers a range of values for γ, including larger values that account for these forces.[21]

4. Housing market effects. Our model abstracts away from the effects of taxes on the supply of housing or on property values. While property taxes can impact the supply of housing, these effects are likely to be small in our setting. In Appendix F, we show that the complexity of the property tax system implies that increasing the property tax rate does not impact the supply of housing. Moreover, because property taxes are based on outdated cadastral values, the government's budget constraint is not affected by policy-driven changes in market values.

A potential concern is that tax and spending policies can be capitalized into property values and that this will affect inter-temporal consumption decisions. Brueckner (1982) models the capitalization of local taxes and public goods into property values. A key result in this literature is that at the efficient level of provision, increasing local public spending through property taxes leaves the value of housing unaffected. Our assumption of no capitalization effects therefore holds when $\text{MVPF}_t = 0$ (e.g., as in Equation 2).[22] Moreover, this assumption biases us against finding that $\text{MVPF}_t > 0$. When public goods are under-provided, increasing public good provision by taxing property would increase property values. Capitalization would increase the wealth of property owners.[23] Unconstrained households would then reduce savings, which would be reflected by a smaller consumption drop in our data. Constrained households would not change their consumption today but would have higher consumption in the future. Ignoring this positive impact on consumption would then lead us to overestimate the welfare cost of taxation.

5. Location decisions and tax competition. Unlike other models of property taxation (e.g., Agrawal et al., 2020), we do not consider migration or tax competition. These concerns are less important in developing countries where internal migration is generally lower (Bell et al. 2015) and in our specific setting because Mexico City offers unique amenities, which makes the possibility of tax-driven migration less relevant. Moreover, Brueckner (1982) shows that migration decisions are neutral to property tax changes in equilibrium.

6. Redistribution. We assume that the tax rate applies to properties of similar values. This assumption matches our setting, as the government relies on a partially progressive tax schedule with different tax rates for different value bands. If households responded to progressive property taxes by moving to properties of lesser value, our model could be extended by accounting for the resulting fiscal externality. However, as discussed above, we do not believe that this is the primary

[21] While the model could endogenize the decision to adjust housing, the ranges of property tax changes that we consider are unlikely to trigger such decisions from homeowners, who represent the vast majority of households.

[22] Because enforcement actions are directed at delinquent individuals, rather than at the market as a whole, Equation 3 is unlikely to be affected by capitalization effects.

[23] Bradbury et al. (2001) and Cellini et al. (2010) provide evidence that tax-financed increases in public good provision have positive effects on property values in the US. Gadenne (2017) finds that tax-financed public spending has significant impacts on the quality of public goods in Brazil. Gonzalez-Navarro and Quintana-Domeque (2016) show that randomly assigned street pavement increased property values in Mexico City. We ignore the case when public goods are over-provided, since property taxes reduce welfare when $\text{MVPF}_t < 0$ regardless of capitalization.

margin of adjustment in the context of developing countries. Future work can extend our model to consider whether a progressive property tax could be used to redistribute wealth.

7. Enforcement as reminders. It is possible that enforcement raises taxpayers' welfare by acting as a reminder, prompting them to take an action which they were planning to implement (see, e.g., Allcott and Kessler, 2019). This is unlikely to be the case in our setting, as taxpayers who do not pay in full at the beginning of the year receive six additional bimonthly tax bills, which also remind them of their outstanding liabilities. Our data on payment timing show that, while compliant taxpayers' payments are closely linked to the bimonthly deadlines associated with these reminders, we do not observe that these bills have a reminder effect for delinquent taxpayers (see Figure B.5). It is therefore unlikely that our additional enforcement message acts as a reminder.

8. Enforcement and non-pecuniary spillovers. A possibility that is absent in our model is that enforcement actions can have spillovers on the compliance of delinquent households that are not subject to enforcement. While such spillover effects may increase the revenue effects of enforcement actions, they would also carry welfare costs by increasing the private costs of not paying taxes. Appendix K estimates spillover effects of tax increases and enforcement letters and finds that, in our setting, the spillover effects of these policies are economically insignificant.

Finally, our baseline model does not allow for the rate of noncompliance to impact the utility of taxpayers. Besley et al. (2023) develop a model with such non-pecuniary spillovers to study the introduction of a poll tax in the UK that increased evasion by 300–500 percent. Appendix M incorporates social norms into our model, as in Besley et al. (2023). Adding social norms does not necessarily increase the relative desirability of enhanced enforcement over tax increases, or vice-versa.[24] While tax hikes increase noncompliance in our empirical setting, our effects are orders of magnitude smaller than those of the poll tax. These motives are therefore unlikely to play a central role in our context.

The model delivers expressions for the welfare effects of tax and enforcement policies and identifies the empirical estimates—including $\varepsilon_t^{\text{Pay}}$, Δc, and $\varepsilon_\alpha^{\text{Pay}}$—that are required to assess the welfare effects of these policies. The following sections describe how we estimate these parameters and Section 7 uses these estimates to provide actionable policy guidance.

2 Property Taxes in Mexico City

This section presents the property tax system in Mexico City.[25] We start by explaining the construction of the tax base and the tax rate schedule. We then discuss the main elements of the tax

[24] This is a result of the fact that the welfare impact of different policies depends on whether social norms generate a utility reward for compliant taxpayers that outweighs the utility cost for delinquent taxpayers.

[25] We abstract from political economy issues because Mexico City has had leftist governments since 1997 that have relied on political support from lower-income individuals, making it unlikely that pressure from wealthy individuals limits property taxation. Current government officials have expressed a desire to increase property taxes but are also sensitive to declining compliance rates and taxpayer hardship.

payment regulation and enforcement.[26]

2.1 Tax Base

The base for the property tax is the cadastral value V_{it} of property i in year t. V_{it} is the sum of the land and construction value, discounted for antiquity until the property is forty years old, whereupon the property value remains constant in age.[27] V_{it} is determined by the formula:

$$V_{it} = (A_{it}L_{it} + U_{it}M_{it})[1 - D_t \cdot (\mathbb{1}_{\{t-t_0 \leq 40\}}(t - t_0) + \mathbb{1}_{\{t-t_0>40\}}40)],$$

where A_{it} is the unit value of land in the neighborhood of property i, L_{it} is the total land area of the property in square meters, U_{it} is the unit value of construction in the neighborhood of property i, M_{it} is the total construction area of the property, D_t is a reduction applied per each year of antiquity, and t_0 is the year of construction of the property. Assessed property values in the cadaster correlate strongly with commercial values (Figure B.1, Panel B).

The distribution of property values is quite stable during the period of our study. The unit values of land and construction, A_{it} and U_{it}, are based on commercial values and were updated only once during the period of our study (between 2008 and 2009). This change does not affect our estimations, which exploit variation between 2009 and 2012. In theory, taxpayers can appeal the cadastral valuation proposed by the government and propose their own valuation. In practice, less than 0.2 percent of appeals are approved.[28] Finally, the age discount makes it possible for properties to change cadastral value bands over time. However, because the discount factor is so small ($D_t = 0.01$), very few properties drop to a lower cadastral value band during the five-year period that we study. We exclude properties with a change in cadastral value band between 2009 and 2012 from our analyses.[29]

[26] While housing property is also taxed indirectly (e.g., through transfer, capital gains, and inheritance taxes), these taxes do not interact with our variation. In contrast to the US, property taxes in Mexico do not determine neighborhood-level public goods.

[27] The registry of property transactions and the cadaster are held by two different levels of government (states and municipalities, respectively) and are not readily mergeable. For this reason, property transactions cannot be used to update cadastral values. While the tax base may depart from market values, note that this is also often the case in high-income countries as well. For instance, California's Proposition 13 generates large differences between assessed and market values. Similarly, Howard and Avenancio-Leon (2022) show that racial differences in assessment appeals drive large differences between assessments and market values across demographic groups.

[28] In 2010, 319,019 taxpayers filed appeals, but only 379 successfully obtained a reduction in their tax base. In 2011, 249 out of 177,681 taxpayer appeals were successful. In 2012, 162 out of 116,729 appeals were successful.

[29] In total, we exclude 284,686 out of more than 1.5 million properties, 87 percent of which registered a cadastral value band change due to an increase in the construction area, for an average cadastral value change of 42 percent. Changes in land area, special amenities (e.g., lifts), and value depreciation over time account for the remaining 13 percent of value-band changes. Our results are robust to including these properties. In addition, exemptions of 30 percent or more of the annual tax liability are available to single mothers with children and seniors with incomes below a specified threshold. Our results are robust to dropping the approximately 7 percent of properties that ever received these exemptions (Tables D.4 and E.3).

2.2 Tax Rates

Figure 2 shows that the property tax schedule is partly progressive. The schedule relies on 16 cadastral value bands: A to P.[30] For bands A–D, the tax is a band-specific lump-sum amount that increases over time with inflation. While the lump-sum amounts increase across bands A–D, they increase by less than the property value. For this reason, the average tax rate is decreasing in property values at the lower end of the value distribution. Properties in bands E–P face a progressive schedule, with marginal tax rates ranging from 7.5–16.9 basis points (a percent of a percent), which yield average tax rates that increase with property values. Using household survey data, we calculate that on average, property tax payments are between 1.5 and 6 percent of per capita household income, with higher values for poorer households (see Figure B.1, Panel C).[31]

While marginal tax rates change little over the years, the average tax rates in bands E–J are also affected by abatements, which are applied to the gross tax liability. Abatements vary over time and have large impacts on average tax rates: abatements vary between 65 percent (in band E) to 10 percent (in band J). The original purpose of abatements was to ensure that the mean tax liability increased gradually from band E to J. Following the 2008 financial crisis, the government decided to remove the abatements one cadastral value band at a time.[32]

The removal of abatements led to large and unexpected changes in mean tax rates over time and across value bands. Because our data cover the years 2008–2012, our analysis exploits three reform episodes: (1) the 2010 abatement removal for value band I, (2) the 2011 removal for band H, and (3) the 2012 removal for band G. Figure 2 shows that the largest rate changes between 2008 and 2012 were caused by removing these abatements. Among the three reform episodes, properties in band G saw the largest increase in taxes, and those in band I saw the smallest increase.

The government of Mexico City announced these rate increases every year when it published property tax rates for the following year. It is unlikely that the changes were anticipated by taxpayers, as they were not widely discussed in the media and each reform episode affected only a small subset of properties. We show below that there is no evidence for any behavioral change prior to the reforms. Figure B.4 shows that properties treated with the quasi-exogenous tax rate increases are well distributed across the city. The average share of treated properties within a zipcode is only about 3 percent. The responses to these reforms therefore constitute individual taxpayer responses

[30]The thresholds for these bands are constant over time, except between 2008 and 2009, when both the band thresholds and the tax rates were updated for inflation. In later years, only the tax rates are updated for inflation each year. Table B.2 shows the property tax schedule for 2009 as an example.

[31]Panel D in Figure B.1 shows the year-on-year growth rate of property tax payments. Given the small annual inflation adjustments to the liability, there is no anchoring of tax payments at the previous year's liability. Instead, the nominal payment amounts increase slightly each year.

[32]Table B.3 lists the abatement rates by value band and year. The government's intention was to remove all abatements but to do so gradually to minimize potential backlash or unrest. Our data do not cover the periods surrounding the abatement removals for bands J, E, and F. Properties in bands A–D never benefited from abatements since their tax liabilities are lump-sum amounts. Removing abatements was the administratively simplest way of raising tax rates. Figure B.1, Panel A, shows that property tax revenues dramatically increased after 2008 in Mexico City.

to tax rate changes rather than responses driven by public debate or general equilibrium changes in policy, attitudes, or perceptions.

2.3 Tax Payment and Enforcement

The legal liability for the property tax rests with the property owner. Property tax bills are delivered to the property and are addressed to the owner. At the beginning of the calendar year, taxpayers receive a bill for the yearly liability. To encourage early payment and increase compliance, the government offers early-bird and super-early-bird discounts if taxpayers pay their yearly liability in full before specific dates. The exact deadlines for the discounts and the discount rates vary over time (see Table B.4). While tax bills include the yearly tax liability, they can be paid in six bimonthly installments. Taxpayers are sent additional bills at the beginning of each bimonthly period with a reminder that payments are due by the last day of the period. Property tax bills can be paid in person at government offices, banks, and convenience stores.[33]

When taxpayers miss a payment, the government automatically updates the unpaid liabilities for monthly inflation and applies a surcharge for every month that a payment is late. Taxpayers who have not paid their yearly liability by April 30th of the following year are catalogued as delinquent taxpayers and face additional penalties and surcharges.[34]

The Ministry of Finance of Mexico City regularly conducts enforcement campaigns to encourage the payment of outstanding property tax debt and to sustain voluntary compliance. Enforcement interventions have varied over time. For instance, enforcement letters have varied in message content over the years (e.g., emphasizing sanctions or public goods provision or simply conveying a reminder). The delivery method for these messages (e.g., letter, phone call, or email) has also varied over time, as has the target group. In some years, all delinquent taxpayers were contacted, while in other years, enforcement has focused on smaller subsets of taxpayers with large debts. We use one of these enforcement interventions to estimate the effects of enforcement in Section 5.

While taxpayers who are unresponsive to administrative enforcement can be prosecuted, the government does not have the capacity to do this in a systematic way. In extreme cases, the government can seize a delinquent taxpayer's property. While this is rare, it does happen (see Table C.1, Panel C). The government can even pursue a jail sentence of up to ten years for tax delinquency.

[33]Figure B.3 shows a typical property tax bill. While the owner may not receive the tax bill if renters do not notify owners or if the cadaster is out of date, this is a minor concern for enforcement in Mexico City, where renters account for only 15 percent of households. Table B.5 shows descriptive statistics on property owners and renters. While it is possible (though not observable to us) that property tax compliance is lower for rented properties, this can explain only part of the delinquency rate, which is much higher (40 percent) than the rental rate.

[34]While the monthly surcharge varies over time, it is on average 1 percent of the outstanding liability for each month of delay. That is, if a taxpayer makes an overdue payment after 6 months, the government adds a 6 percent surcharge to the inflation-updated liability. Table C.1, Panel C, shows the additional fees paid by delinquent taxpayers who made outstanding tax payments in 2008 and 2009. The table shows that late payment fines are applied to almost all late-payers and that the fines represent a substantial fraction—between 15 and 30 percent—of the tax liability.

3 Administrative Tax Data

Our empirical analyses exploit three datasets on the universe of tax-liable residential properties in Mexico City from 2008 to 2012. First, the cadaster—or tax register—lists all properties with their unique property tax ID, post code, and property characteristics such as land area, construction area, land and construction value, and total property value.

Second, we use data from annual and bimonthly property tax bills for all properties. These bills include the property value, tax liability, bill issue date, and due date. Third, we use data from the universe of property tax payments. For each payment, we have data on the relevant tax bill and period, amount, date, and additional variables including inflation adjustments, surcharges, and penalties for late payment. We link the billing and payment data to the cadaster via the unique ID. Our main dataset consists of a balanced panel of 1.42 million properties.

The majority of our analyses study outcomes at the property-year level. The main outcome variables are the annual payment amount in current Mexican pesos (MXN, thousand) and the compliance share, defined as the ratio of tax payment to gross liability. Additional outcomes include dummies for zero, partial, and full payment of the net tax liability (net of any early-bird and super-early-bird discounts). Finally, we characterize payment timing with dummies indicating early (all-at-once), bimonthly (payment in installments), and late payment.

A salient feature of the data is that the distribution of properties is skewed toward low value bands. While the majority of properties fall into bands A–E, the distribution of cadastral values is slightly less skewed. Because of the progressive tax schedule, tax liabilities are more evenly distributed across value bands. However, because compliance is higher for low value bands, the distribution of tax payments is less evenly distributed (see Figure C.1 for details).

The data reveal interesting trends. While the gradual eliminations of abatements led to a rise in the mean tax liability, average tax payments—in absolute terms and as a share of the yearly liability—have decreased over time. The decrease in tax payments is partly driven by a rise in the share of properties making zero payments and a decline in the share of properties paying in full. These patterns showcase the importance of understanding how households respond to tax rate increases and whether enforcement efforts can influence the decaying compliance rate.[35]

4 The Elasticity of Tax Revenues to the Tax Rate

This section estimates the effects of tax rate changes on tax payment and compliance. We study three quasi-experimental reforms: the mean tax increases for properties in value bands I, H, and G in 2010, 2011, and 2012, respectively. We first present results from an RD estimation that exploits the sharp discontinuities in tax rate changes at thresholds between the treated value bands and

[35] Table C.1 describes additional property characteristics. For instance, the average property was built in 1985 and has a land area of 123 square meters and a construction area of 126 square meters.

bands below them. This approach yields estimates of short-term local average treatment effects for properties close to each threshold. We then estimate medium-term effects with a DiD design that uses properties in never-treated value bands as controls.

4.1 Short-Term Effects: Regression Discontinuity Estimation

Our RD estimation relies on discontinuous tax rate *changes* at the lower thresholds of cadastral value bands that experienced large tax rate increases. We focus on rate changes—rather than levels—since tax rate levels differ between value bands and there are small yearly inflation adjustments to rates in all bands. Finally, we use the band below as the counterfactual since properties in the band above the treated band were treated in the previous year.

Consider the properties in a treated band in year t as well as the properties in the band immediately below it. Let $\hat{V}_i = V_i - V_-$ denote the distance between the value of property i, V_i, and the lower limit of the treated band, V_-. Let $Y_{i,t}$ denote the outcome of interest for property i in period t. We estimate the effect of the tax rate increase on the year-on-year change in the outcome of interest as follows:

$$\Delta Y_{i,t} = \alpha + \beta T_i + f(\hat{V}_i) + g(\hat{V}_i)T_i + \epsilon_{i,t},$$

where T_i is an indicator for properties in the treated band, i.e., $T_i = \mathbb{1}_{\{V_i \geq V_-\}}$, f and g are continuous functions, and $\epsilon_{i,t}$ is an error term.

The validity of this approach relies on the assumption that taxpayers cannot manipulate their property valuation in response to a change in the tax rate. To validate this assumption, we test for a discontinuity in the distribution of the running variable around the treatment cut-off. We are unable to reject the null hypothesis of no manipulation of property values around the treatment cut-off for all three reforms using either the McCrary (2008) test (see Figure D.1, Column (A)) or the Bugni and Canay (2020) test (see Table D.1).[36] In addition, we test for discontinuities in property characteristics around the treatment thresholds and find no significant differences in the year of construction, land area, or construction area (see Figure D.1, Columns B–D). Finally, we note that in the unlikely case that taxpayers anticipated the tax rate increase and responded before the change went into effect, our estimates of the tax payment increase would be biased towards zero. We show below that there is no evidence of an anticipatory response.

Estimates

Figures 3 and 4 show the results for the three different reform episodes (rows) and four different outcomes (columns). Each panel plots the year-on-year (pre- vs post-reform) change for a given

[36] The Bugni and Canay (2020) test provides an alternative approach to testing for manipulation of cadastral values. This test examines the balance in the number of observations around the cut-off. Unlike the McCrary (2008) test, it does not rely on local density estimates.

outcome in 20 equally spaced cadastral value bins around the lower threshold of the treated band. Each graph reports a third-order polynomial fit along with 95 percent confidence intervals.

Column (A) of Figure 3 shows that the legislated tax rate increases were indeed applied as intended and generated a 9.1 basis point increase in the mean tax rate at the band threshold in 2010 and even larger increases of 12.1 and 18.0 basis points in 2011 and 2012. Column (B) shows that tax payments jumped substantially—between 450 and 600 MXN—in all reform episodes. However, payments increased by less than the mechanically expected increase, as compliance fell.[37] Figure 4 shows that the share of taxpayers paying their liability in full fell by 5.5 ppt in 2010, by 6.4 ppt in 2011, and by over 10 ppt in 2012. The compliance share also decreased in all reform episodes by 3.2–6.2 ppt.[38]

Table 2, Panel A, presents the implied tax rate elasticities for the outcomes considered in Figures 3–4.[39] Column (2) lists the elasticities of tax revenue to the tax rate. When the change in the mean tax rate is 9.1 basis points, the elasticity of tax revenues is 0.55, whereas the estimate is 0.31 when the tax rate increases by 18.0 basis points. While these estimates are consistent with the notion that larger tax rate changes also generate larger compliance responses, we cannot reject the null that these elasticities are equal to each other at conventional significance levels. Table 2 also reports tax compliance elasticities, which we use in our policy analysis in Section 7.

To demonstrate the robustness of our results, we compare our main estimates from the cubic polynomial regression with the results from local polynomial regressions with varying bandwidths and degrees of polynomial in Figure D.2. Table D.2 shows the results from specifying an optimal bandwidth in local linear regressions as in Calonico, Cattaneo and Titiunik (2014). The estimates are statistically indistinguishable from those in Table 2.

As an additional robustness test, we consider an alternative specification for our RD estimates. Because we study the effects of discontinuous tax changes around thresholds, we follow Lalive (2008), Lemieux and Milligan (2008), and Grembi, Nannicini and Troiano (2016) by estimating a differences-in-discontinuities model on our panel data:

$$\Delta Y_{i,t} = \alpha_0 + \beta_0 T_i + f_0(\hat{V}_i) + g_0(\hat{V}_i)T_i +, \\ [\alpha_1 + \beta_1 T_i + f_1(\hat{V}_i) + g_1(\hat{V}_i)T_i]D_t + \epsilon_{i,t},$$

where D_t is an indicator for the time period when the abatement is removed. The effect of the abatement removal, in excess of the effect of the smaller year-on-year tax rate changes, is given by β_1. The results displayed in Table D.3 show that the β_1 estimate from this equation is very similar

[37] The tax liability for a property at the value band threshold of 2.3 million MXN increased by 2100 MXN in 2010. Liabilities increased by close to 2350 MXN in 2011 and 2900 MXN in 2012.

[38] Note that these compliance drops are not due to changes in the tax base, as the latter is not updated during the study period and appeals against the tax liability are overwhelmingly unsuccessful, as mentioned in Section 2.1. Because administrative tax data do not record whether properties are occupied by renters or owners, we cannot explore heterogeneity along this margin.

[39] We compute the elasticity $\epsilon_{y,t} = \frac{\partial y}{\partial t}\frac{t}{y}$ using $\frac{\partial y}{\partial t}$ from the RD estimates and $\frac{t}{y}$ from outcome means at baseline.

to our main estimates, and we can generally reject the null hypothesis that $\beta_1 = \beta_0$.

These results show that while sharp increases in average tax rates have sizable effects on tax payments, taxpayers also respond by decreasing their compliance with the property tax. While the RD approach yields precise and highly credible estimates of the short-term responses to the tax increases, compliance in future years may depend on broader responses by taxpayers.[40]

4.2 Medium-Term Dynamics: Differences-in-Differences

An important question for policymakers is whether the effects of tax rate increases persist over time or are temporary. For instance, while liquidity-constrained taxpayers may temporarily decrease compliance after a tax rate increase, they may also make up for missed payments in later years. In this example, RD estimates would over-estimate the medium-term revenue elasticities.

We estimate medium-term effects using a DiD design that captures the evolution of compliance outcomes over time. This approach compares properties in the treated value band to properties in other high-value bands that never experienced a tax increase. Specifically, for a tax rate increase occurring in year t_0, we estimate:

$$Y_{it} = \alpha + DD_{it}\beta + \gamma_i + \delta_t + \epsilon_{it}, \tag{7}$$

where Y_{it} denotes the compliance outcome of property i in year t, α is a constant, DD_{it} is a dummy taking the value of 1 when property i belongs to the treated value band and $t \geq t_0$, δ_t and γ_i denote year and property fixed effects, and ϵ_{it} is the error term. For all years, we use properties in bands K and L as controls.[41] The identifying assumption is that absent the tax hikes, the outcomes for properties in the treatment and control groups would have trended in parallel. Under this assumption, the point estimate for β captures the causal effect of the tax rate change on compliance. Given the large share of zeros in our outcome variables, we estimate Equation 7 in levels. To obtain a relative effect, we scale the point estimates and standard errors by the mean outcome in the treatment group in the last pre-reform year.[42]

Estimates

Figures 5 and 6 capture dynamic responses to tax rate changes. The figures are structured like Figures 3 and 4: the rows pertain to the three different reform episodes, while the columns reflect the different outcome variables. In each graph, the vertical black line indicates the timing of

[40] Because the control bands are treated in $t+1$, we cannot use the RD approach to estimate dynamic responses.

[41] Among the properties in bands that are never treated with one of the large quasi-experimental tax rate increases, those in bands K and L are the closest in value to the properties in the treated bands. Figures E.2–E.4 show that properties in other untreated bands also exhibit similar pre-reform trends and yield similar results.

[42] The results are very similar when we estimate Equation 7 via pseudo-Poisson maximum likelihood (Santos Silva and Tenreyro 2006), which is suitable when outcome variables are highly skewed or have a large share of zeros (Brockmeyer and Hernandez 2022).

treatment, the red solid line represents the average outcome in the treated band, and the blue dotted line represents the average outcome in control bands K and L. In all panels, outcomes are scaled by the pre-reform group-specific mean.

The timing of the reforms and the length of our dataset mean that we can observe three post-reform periods for the 2010 reform, two post-periods for the 2011 reform, and one post-period for the 2012 reform. On the other hand, we observe the longest pre-reform period (four years) for the 2012 reform and the shortest (two years) for the 2010 reform. We detect no significant difference in pre-trends between the treatment and control groups for any reform episode and no evidence of an anticipated response to the reform. In contrast, we observe a precise and sharp deviation in trends in each of the reform years.

Consistent with the legislative changes, mean tax rates increased significantly after every reform event (Figure 5, Column (A)). As with the RD estimates, we find large increases in tax payments (Column (B)). However, decreases in the share of taxpayers paying in full and in the compliance share (Figure 6) show that compliance also fell significantly. The results are qualitatively similar across the three reform episodes, though the magnitude of the compliance drop is largest for the 2012 reform, which triggered the largest tax rate change. Following the 2012 reform, the full payment share fell by 30 percent, and the compliance share fell by 18 percent. For the 2010 reform, a 17 percent increase in the mean tax rate triggered a 11.6 percent increase in tax payment amounts. The payment response was moderated by a 4 percent reduction in the compliance share. In turn, the drop in the compliance share was partly driven by a 10 percent drop in the share of on-time payments-in-full.

One possible explanation for the smaller estimates for the 2010 and 2011 reforms is that these estimates capture effects over longer post-reform periods (two and three years, respectively), while the estimate of the 2012 tax change only captures the effect for a single post-reform period. However, we obtain similar compliance drops for the 2010 and 2011 tax changes when we estimate DiD effects using a single post-reform period (see Table E.1). These results raise the possibility that larger tax changes can trigger more-than-proportional compliance responses.[43]

Table 2, Panel B, summarizes the treatment effect estimates and the implied elasticities. The elasticity of tax revenues with respect to the tax rate is presented in Column (6). This elasticity ranges from 0.697 in response to the 17 percent increase in the tax rate affecting band I in 2010 to 0.489 in response to the 40 percent increase in the tax rate affecting band G in 2012.[44] While these numbers are very similar to the RD estimates, the DiD estimates are more precisely estimated. Using panel regressions that exploit variation in tax rates across the full range of value bands—and not just for those treated in our quasi-experiments—yields similar results: an average payment

[43]We also find that the tax rate hikes increased taxpayers' likelihood of making an appeal against their tax liability (Figure E.1), even though hardly any appeals lead to a change in the tax liability (see footnote 28).

[44]In Appendix E, we perform robustness tests of our DiD estimations. We confirm that the results are very similar when we use other value bands as the control group (Table E.2 and Figures E.2–E.4), dropping taxpayers benefiting from exemptions (Table E.3) and clustering the standard errors at different levels or bootstrapping them (Table E.4).

elasticity of 0.69 (SE=0.09) and a compliance elasticity of -0.29 (SE=0.10) (see Table E.5). These results suggest that our main estimates are broadly valid for property taxpayers in Mexico City.

The results of our DiD analysis show that tax rate increases lead to persistent changes in both tax payment and compliance behaviors. By using a different set of control properties than in the RD analysis, the DiD also bolsters the likelihood that we are measuring the causal effects of changes in tax rates.[45]

5 The Elasticity of Tax Revenues to Enforcement

Governments around the world have two options when seeking to raise tax revenue: increasing tax rates on taxpayers who are not delinquent or increasing the number of taxpayers by enforcing existing taxes on delinquent taxpayers. In the context of Mexico City, this trade-off is stark since 40 percent of taxpayers are delinquent. We study this trade-off by estimating the elasticity of tax payment to enforcement using a field experiment that we designed and evaluated in collaboration with the Ministry of Finance of Mexico City.

5.1 Field Experiment

The Ministry of Finance sent out enforcement letters to 80,000 delinquent taxpayers between July 28 and August 11, 2014, requesting that they pay their outstanding tax debt accumulated from bimester 4 of 2009 to bimester 3 of 2014. A control group of 10,000 delinquent taxpayers received no letter. The comparison of taxpayers who received a letter to those who did not receive any letter identifies the policy-relevant effect of enforcement efforts.[46] The mode of delivery, sample selection, and information provided in the letters corresponds to the ministry's typical practices. Therefore, our estimates can be viewed roughly as in-equilibrium effects. Our estimates also have general validity, as the cadastral value distribution among delinquent taxpayers is similar to that of the population (Table G.1). Each personalized letter lists the bimester(s) for which the tax payment is overdue, requests payment within 15 working days after receipt of the letter, and lists the institutions accepting payment (tax administration offices, bank branches, convenience stores).

The treatment group in our intervention was divided into eight groups of 10,000 taxpayers, each receiving a slightly different variant of the letter. Figure G.1 illustrates the experimental design, and Figure G.3 shows the text of the letters. The motivation for testing various treatments was to identify the strongest possible enforcement message and hence maximize the potential for enforcement to be effective—both in terms of raising revenue and in terms of improving welfare.

[45] A potential concern is that tax rate increases might reduce investment in real estate. We discuss this possibility and test for it empirically in Appendix F. We find no evidence of impacts on real estate investment.

[46] Studies that examine enforcement *design*, e.g. the optimal wording of messages, often compare status-quo enforcement messages to alternative designs (e.g., Atinyan and Asatryan, 2022). We are instead interested in studying whether the tax administration should send enforcement messages at all or whether it should focus on other policy changes, e.g. tax rate changes.

Half of the letters put additional emphasis on sanctions used to enforce the tax (referred to below as the sanctions treatment), while the other half emphasized the fact that property tax revenue is used to fund health services, education, and community infrastructure (referred to as the public goods treatment). However, even the public goods letter should be seen as an enforcement treatment, rather than as a nudge, reminder, or information treatment that could generate warm glow. Indeed, all messages were signed by government officials and mentioned that tax payment is a duty and that late payers would incur fines and surcharges.

5.2 Empirical Results

Figure 7 displays the effects of the enforcement intervention. The plots show trends in payment outcomes around the time of the intervention and distinguish the control group, the sanctions treatment, and the public goods treatment. As expected, the three groups exhibit linear trends in all outcomes prior to the intervention. The treatment groups start diverging in early August when the first letters are delivered.[47] The divergence accelerates sharply by mid-August. This timing coincides with the end of the 15-day deadline to respond to the letter starting after all the letters are delivered. In contrast, we do not see any trend changes for the control group.

The graphs display point estimates for β_1 and β_2 from the regression:

$$Y_i = \alpha + \beta_1 T1_i + \beta_2 T2_i + \epsilon_i,$$

where Y_i is the outcome for property i evaluated 40 days after the sending of all letters, α is a constant, $T1_i$ and $T2_i$ are dummies indicating the two mutually exclusive treatments (the sanctions treatment and the public goods treatment), and ϵ_i is the error term.

The results in Figure 7 show that the sanctions treatment generated a 9.4 ppt increase in the likelihood of making a payment toward outstanding tax debt and a 54 peso increase in the payment amount. Relative to the control group, the intervention close to tripled the payment likelihood and doubled overall payments. The public goods treatment had smaller but statistically significant effects. Comparing Panels B (any payment) and C (payment amount) suggests that the public goods treatment is relatively more regressive, generating payment by taxpayers with disproportionately smaller liabilities.[48] We find similar results when we control for property characteristics X_i (e.g., cadastral value and age of the property) or when we estimate treatment effects using a DiD framework (Table G.2).

Table 3 reports estimates of the treatment effects for all treatments in pair-wise comparisons. Panel A lists the effects on the likelihood of making a payment toward outstanding tax debt,

[47] It takes three to five days for letters to be delivered.

[48] We interpret the public goods treatment effect as being due to enforcement, since we compare the public goods letter to a non-letter control and since any type of letter from the Ministry of Finance is likely to give the impression of increased enforcement. Studies comparing public goods messages to neutral baseline messages from tax administrations mostly find no effect (Atinyan and Asatryan 2022).

and Panel B shows the effects on the payment amount. In addition to confirming the significant difference between the sanctions and public goods treatments (Column (1)), the table shows that the seniority of the enforcement officer matters: a fiscal attorney signature achieves a larger impact than a compliance officer signature (Column (2)). Furthermore, although one might expect a gender bias in taxpayers' response to male/female signatures, our evidence rejects this idea. In fact, male signatures have a smaller impact than gender-neutral signatures, and female signatures have a slightly larger impact than gender-neutral signatures (Columns (3) and (4)). While we consistently reject the null hypothesis that male signatures have larger effects, we can reject the null hypothesis that female signatures have smaller effects only when we control for property characteristics (Table G.2).

The size of the enforcement effects in Mexico City compares favorably to that of other compliance interventions. In a meta-analysis of tax compliance experiments, Atinyan and Asatryan (2022) find that deterrence nudges increase extensive-margin compliance on average by only 1.–2.5 percent. When focusing on comparable property tax compliance interventions, we find that our effects are slightly smaller than those in Weigel (2020), who shows that a door-to-door campaign in the Congo increased compliance (likelihood of payment) from 0.05 percent to 11 percent; similar to those in Okunogbe (2021), who finds that a detection and penalty intervention in Liberia increased compliance from 3 percent to 9 percent; and larger than those in Del Carpio (2014), who finds that a social norms intervention in Peru increased compliance from 29 percent to 34.5 percent.

The results of our field experiment show that enforcement actions are a cost-effective means of raising revenue for the government. However, while stricter enforcement may raise tax revenue, the welfare costs of enforcement due to the private costs it generates to delinquent taxpayers may exceed the revenue benefits. Section 7 uses our empirical estimates to implement our model and provides policy guidance on whether it is preferable to raise tax revenue through increased tax rates or through stricter enforcement.

6 Property Taxes and Liquidity Constraints

This section provides evidence that validates our model assumptions that liquidity constraints can exacerbate the welfare cost of taxing property. In Section 6.1, we use methods from the consumption literature to test for the presence of liquidity constraints and show that income shocks impact property tax payments. In Section 6.2, we present causal evidence that the taxpayers' decision to pay their liability in full as opposed to paying in installments is highly elastic to the tax rate, which suggests that liquidity constraints impact property tax payments.[49]

[49]In Appendix H, we present additional evidence from a third approach, which relies on taxpayer responses to the discount rate of promotional "early bird" campaigns, implemented annually by Mexico City's government. We use a discrete choice model that leverages daily variation in payment rates and annual variation in early bird discount rates to show that taxpayers have a high valuation for liquidity.

6.1 Excess Sensitivity to Income Shocks

The consumption literature identifies the presence of liquidity constraints through the extent to which household consumption responds to unanticipated temporary income shocks (see, e.g., Jappelli and Pistaferri, 2017). With longitudinal data, researchers could estimate the following relationship between household consumption and household income:

$$\log(\text{Consumption}_{it}) = \log(\text{Income}_{it})\beta + \gamma_i + \delta_t + \epsilon_{it}, \tag{8}$$

where γ_i and δ_t are household and time fixed effects and β is the elasticity of consumption with respect to temporary, idiosyncratic deviations of household income from its mean. Under the null of perfect insurance, we would expect to find a $\beta = 0$; a positive value of β shows that consumption is imperfectly insured, as it responds to temporary income shocks.

In the absence of longitudinal data, Attanasio (1999) shows that Equation 8 can be estimated using a pseudo-panel of household groups with shared characteristics. Specifically, the income elasticity of consumption β can be estimated by regressing the variance of household consumption on the variance of household income by age×education cohort as follows:

$$\text{var}_g[\log(\text{Consumption}_{it})] = \text{var}_g[\log(\text{Income}_{it})]\beta + \gamma_g + \epsilon_{gt}, \tag{9}$$

where γ_g is a cohort-specific fixed effect.[50]

To estimate this relationship between consumption and income shocks, we construct a pseudo-panel of head-of-household age×education cohorts, using data from the *Encuesta Nacional de Ingresos y Gastos de los Hogares* (ENIGH) household survey.[51] We use data from the 2008, 2010, 2012, and 2014 waves of the ENIGH for Mexico City and include heads of household born between 1948 and 1982. We create eight five-year period-of-birth categories and six education-level categories, for a total of 48 age×education cohorts. We then calculate the cohort-level variance of household log income and log consumption per capita. We also calculate the variances of two additional outcomes to examine how idiosyncratic income risk affects property tax compliance: the log of one plus the annual property tax payment per capita and a dummy for positive property tax payments.

Table 4 presents the results of estimating Equation (9) via OLS. Column (1) shows that in response to a 1 percent temporary drop in household per capita income, consumption drops by 0.617 percent, implying that Mexico City's households are able to insure only 38 percent of their consumption against income shocks. Column (2) shows that estimating the same equation by first differences instead of including cohort-specific fixed effects yields statistically indistinguishable

[50] The parameter β is the same across specifications 8 and 9 under the assumption that the covariance terms resulting from taking the variances from both sides of Equation 8 are constant over time, as in Attanasio and Székely (2004). In Appendix L, we interpret the results in this section using a structural model of the income process.

[51] While longitudinal data would be ideal for our purposes, the ENIGH, which is the only periodic household income and consumption survey in Mexico, is a repeated cross-section.

results.

Households are even less able to insure property tax payments than consumption against income shocks. Column (3) of Table 4 shows that a 1 percent temporary income shock reduces the annual proper tax payment by 0.837 percent. Similarly, column (5) shows that income shocks affect the extensive margin of payment, since a temporary income drop of 1 percent decreases the probability of payment by 0.076 percentage points. Estimating these regression by first differences in columns (4) and (6) yields similar coefficients.

6.2 Sensitivity of Payment Modality to the Tax Rate

We now examine how taxpayers' choice of payment modality responds to the tax rate. Taxpayers can pay their annual liability in installments or in full at the beginning of the year to take advantage of the early-bird discount. Liquidity constraints may lead households to pay in installments rather than all at once. By giving up the discount, these households effectively take out a loan from the government. Conditional on attempting to pay in installments, liquidity-constrained taxpayers may also be less likely to remain compliant.

We provide two pieces of evidence that taxes impact payment modality. First, we exploit variation generated by the non-monotonous shape of the tax rate schedule in a regression-kink-style analysis. As discussed above, the mean tax rate as a share of the property value is first decreasing and then increasing with property value. The lines with blue square markers in Panels A–C of Figure 8 plot the mean tax rate schedule, which features an inflection point in value band D.

Panel A in Figure 8 also shows that the likelihood of making a tax payment (including a partial payment) is negatively correlated with the mean tax rate. Payment compliance is first increasing and then decreasing in property values, with a peak in band D, where the tax rate is minimized. Panel B shows that conditional on the taxpayer making a payment, the likelihood of paying all at once (as opposed to paying in installments) is also negatively correlated with the mean tax rate, again with an inflection point in band D. Panel C shows that conditional on the taxpayer attempting to pay in installments, the likelihood of paying all six installments on time (as opposed to paying partially or late) is also negatively correlated with the mean tax rate, again with an inflection point in band D. This evidence is highly consistent with liquidity constraints shaping taxpayer behavior, as the inflection points in these three data series are otherwise hard to explain.

To provide a second piece of evidence that taxes impact payment modality, we now show that the quasi-experimental tax rate increases that we analyze in Section 4.2 are also associated with changes in payment modality. In Panels D–F of Figure 8, we implement the DiD design in Equation 7 using the sample of taxpayers who make at least a partial payment in each year. Consistent with the view that tax rate increases activate liquidity constraints, we find that treated taxpayers become more likely to pay their liability in installments instead of all at once. These taxpayers hence forgo the early-bird discounts. One interpretation of the discounts is that taxpayers who do not pay in

Table 1: Parameters for Welfare Calculations

A. Welfare Effects of Tax Increases

$MVPF_t$	$\varepsilon_t^{\text{Pay}}$	Δc^{Pay}	γ	$\tilde{z}$
$\frac{v'(g)}{u'(c)} - \frac{1-\gamma\Delta c^{\text{Pay}}}{1+(1-\tilde{z})\varepsilon_t^{\text{Pay}}}$	$\{-0.24, -0.46\}$	$\{0\%, -2.56\%\}$	$\{1, 3\}$	0.10

B. Welfare Effects of Enforcement

$MVPF_\alpha \times \omega\frac{m'(\alpha)}{u'(c)}$	$\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha)$	$\varepsilon_t^{\text{Pay}}$	N^{Pay}	ω
$\frac{v'(g)}{u'(c)} - \frac{\omega(1-N^{\text{Pay}})}{(1-\tilde{z})N^{\text{Pay}}\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha)-1} \times \frac{\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha)}{-\varepsilon_t^{\text{Pay}}}$	$\{16, 54\}$	$\{-0.24, -0.46\}$	$\{0.60, 0.75\}$	$\{0.5, 1\}$

Notes: This table displays the parameters that we use for the policy simulations: empirically estimated parameters (blue) and the calibrated or calculated parameters (red).

full before the deadline take out a loan from the government. The results suggest that taxpayers are more likely to avail themselves of these loans when tax rates increase.

Overall, our results confirm the importance of liquidity constraints in shaping consumer and taxpayer behavior in Mexico City.

7 Welfare Effects of Taxes and Enforcement

We now connect our empirical results with our conceptual framework to quantify the welfare effects of the policies most commonly used to raise property tax revenue.[52] Table 1 collects the ingredients needed to perform these calculations.

7.1 Tax Rates

Consider first the effects of raising tax rates on welfare. The formula for the $MVPF_t$ in Panel A of Table 1 shows that, absent liquidity constraints ($\Delta c^{\text{Pay}} = 0$), the welfare cost of raising tax rates is given by $\frac{1}{1+(1-\tilde{z})\varepsilon_t^{\text{Pay}}}$. To implement this formula, we use data to calibrate that the share of back-taxes that the government eventually collects in future years, $\tilde{z}$, equals 10%.[53] In the model, the elasticity $\varepsilon_t^{\text{Pay}}$ captures the decrease in compliance following a tax increase. This concept is best approximated by the compliance share elasticity, which weights drops in compliance by revenue. Our central estimate of this elasticity is $\varepsilon_t^{\text{Pay}} = -0.29$. For illustrative purposes, we consider the smallest and largest estimates of this elasticity: $\varepsilon_t^{\text{Pay}} \in \{-0.24, -0.46\}$. Assuming that $\varepsilon_t^{\text{Pay}} = -0.24$ implies that the welfare cost of an additional peso of tax revenue equals 1.28.[54] The welfare effect of using this tax revenue to provide public goods depends on the value of $\frac{v'(g)}{u'(c)}$. The blue line in

[52]Following the model in Section 1, we assume that the government aims to set tax and enforcement policies to maximize the well-being of its residents. In Appendix J, we alternatively assume that the government aims to maximize revenue, and we show that the current tax rates are far below the revenue-maximizing rates.

[53]Figure C.2 shows that the government recovered 10 percent of the outstanding debt between 2008 and 2012.

[54]This cost rises to 1.35 when $\varepsilon_t^{\text{Pay}} = -0.29$ and to 1.59 when $\varepsilon_t^{\text{Pay}} = -0.46$.

Panel A of Figure 9 plots the $MVPF_t$ for values of $\frac{v'(g)}{u'(c)} \in [1,3]$.[55] The $MVPF_t$ is increasing in the marginal value of public goods and, in this initial parameterization, indicates that tax hikes raise welfare as long as $\frac{v'(g)}{u'(c)} > 1.28$.

We now consider how liquidity constraints alter the welfare cost of raising tax rates. The formula for the $MVPF_t$ shows that when tax hikes lead to consumption drops, i.e., $\Delta c^{\text{Pay}} < 0$, the welfare cost of raising an additional peso of tax revenue is multiplied by $(1 - \gamma \Delta c^{\text{Pay}})$. To obtain an upper bound on the impact of liquidity constraints, we assume that consumption is entirely hand-to-mouth for liquidity-constrained consumers, so an increase in the property tax translates directly into an equally-sized change in consumption. The average share of property tax payments in per capita household income is 3.2% and we calibrate the share of liquidity-constrained households to 80%, which is the share of households that do not have a credit card.[56] The red-dotted line in Panel A of Figure 9 adjusts our $MVPF_t$ calculation using this calibrated consumption drop along with the assumption that $\gamma = 1$. Since the welfare cost of raising taxes increased by $-\gamma \Delta c^{\text{Pay}} = 3.2\% \times 80\% = 2.56\%$, this line lies slightly below the solid blue line. In this scenario, raising taxes to provide public goods raises welfare as long as $\frac{v'(g)}{u'(c)}$ is greater than 1.32. The dashed orange line shows the effect of instead assuming that $\gamma = 3$.[57] In this case, liquidity constraints increase the welfare cost of tax hikes by 9.6%, and the break-even value of $\frac{v'(g)}{u'(c)}$ rises to 1.4. For illustrative purposes, the green dot-dashed line plots the $MVPF_t$ assuming that $\varepsilon_t^{\text{Pay}} = -0.46$. This line shows that increasing taxes only raises welfare if the value of public goods is greater than 1.87. Overall, this figure shows that under our central estimates of $\varepsilon_t^{\text{Pay}}$, the government can increase welfare by raising property tax rates to provide public goods and that this conclusion is robust to accounting for liquidity constraints as well as a range of other parameter values.

Panel B of Figure 9 implements the optimal property tax from Equation 2.[58] For ease of interpretation, we plot only positive tax rates, and we top-code optimal tax rates at 250 basis points. We also plot the average tax rate (solid grey line) and the maximum tax rate (dashed grey line) for reference. Assuming no liquidity constraints and that $\varepsilon_t^{\text{Pay}} = -0.24$, the blue solid line shows that as long as the value of public goods exceeds 1.5, the optimal tax rate is greater than 250 basis points. Liquidity constraints (red dotted line), higher welfare costs of consumption declines (yellow dashed line), and larger compliance drops (green dot-dashed line) all work to reduce the optimal tax rate. While these lines illustrate how different forces influence optimal tax rates, the red dotted line is a reasonable case for practical purposes. This line shows optimal tax rates greater

[55] While estimates of the value of public goods from the United States imply $\frac{v'(g)}{u'(c)} \approx 1.5$ (Cellini et al., 2010; Suárez Serrato and Wingender, 2014), this value is likely larger in countries with a lower provision of public goods.

[56] This share is likely an upper bound. In Appendix L we present a structural model of income and consumption suggesting that at least 30% of households are liquidity constrained.

[57] This calibration serves two purposes. First, Chetty and Looney (2006) offer a similar calibration in a developing country context. A larger value of γ captures the possibility that consumption commitments amplify welfare costs (Chetty, 2004). Second, the dashed orange line can also be interpreted as setting $\gamma = 1$ and increasing the consumption drop for liquidity-constrained households to $3 \times -3.2\% = -9.6\%$.

[58] As with most sufficient statistic formulas (Chetty, 2009), we assume constant elasticities to implement Equations 2 and A.1. Equations 1 and 3 do not rely on this assumption, as they measure marginal effects of policy changes.

than 45 basis points whenever the value of public goods is greater than or equal to 1.5. This suggests that while liquidity constraints generally work to lower optimal tax rates, current property tax rates are likely below the optimal rates.[59]

7.2 Enforcement

Consider now the welfare effects of enforcement. To make these effects comparable to those of rate hikes, we multiply the $MVPF_\alpha$ by the change in tax liability that would be equivalent in its effect on compliance to a given enforcement treatment. Panel B of Table 1 shows that this equivalent tax liability is given by $\frac{\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha)}{-\varepsilon_t^{\text{Pay}}}$.[60] Assuming that $\varepsilon_t^{\text{Pay}} = -0.24$, the equivalent increases in tax liability are 66.67 MXN for the public goods treatment (i.e., $\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha) = 16$) and 225 MXN for the sanctions treatment (i.e., $\varepsilon_\alpha^{\text{Pay}}(Ht/\alpha) = 54$). Using the observed share of compliers $N^{\text{Pay}} = 60\%$ and assuming that society places an equal weight on the welfare of all households (i.e., $\omega = 1$), our calculation implies that an additional peso of tax revenue that is raised by the public goods treatment has a welfare cost of 3.49. In the case of the sanctions treatment, this welfare cost is slightly smaller and equals 3.20.[61] The blue and red dotted lines in Panel C of Figure 9 plot the $MVPF_\alpha$ under these two scenarios and show that raising tax revenue through these enforcement messages does not raise welfare for values of public goods smaller than 3. These lines show that, in our setting, even though enforcement interventions are effective in raising revenue, they are not a desirable policy tool for a welfare-maximizing government.

We now discuss alternative settings under which enhanced enforcement may increase welfare. First, consider a situation where rate hikes are not effective in raising revenue. Setting $\varepsilon_t^{\text{Pay}} = -0.46$ lowers the welfare cost of raising revenue through enforcement to 1.66 in the case of our sanctions treatment. The orange line in Panel C of Figure 9 plots the $MVPF_\alpha$ in this case. Enforcement can therefore be more effective in settings where rate hikes yield little additional revenue (e.g., Bergeron et al., 2023). Second, consider a case where the government under-values the welfare of non-compliant households. The green line in Panel C of Figure 9 plots the $MVPF_\alpha$ further assuming that $\omega = 0.5$. Because the welfare cost of enforcement are mechanically down-weighted by the government in this case, the $MVPF_\alpha$ is higher than in other cases. In the extreme where $\omega = 0$, the revenue from enforcement is "free money," so enforcement is always desirable in this extreme case.

While the green line helps visualize the role of the welfare costs of enforcement, there are at least two reasons why the government may not want to set $\omega < 1$. First, doing so would violate the principle of horizontal equity. Second, this assumption is particularly problematic in our setting

[59] As we note in Section 1, allowing for the capitalization of taxes and spending on property values would result in higher optimal tax rates.

[60] As we discuss in Section 1.4, this term results by expressing $m'(\alpha)$ as the utility cost of a tax that raises the same revenue as an increase in enforcement. Because in our setting enforcement can be interpreted in terms of the money spent on mailing enforcement letters, we assume that $a(\alpha) = \alpha$, so that $a'(\alpha) = 1$.

[61] If $\varepsilon_t^{\text{Pay}} = -0.29$, these values equal 2.90 and 2.65, respectively for the public goods and sanctions treatments.

since our results suggest that liquidity constraints may contribute to the observed non-compliance. Assuming that $\omega < 1$ may therefore penalize households that are simply unable to pay their taxes.

7.3 Comparison of Tax Rates and Enforcement

We now use our model to evaluate whether it is preferable to increase tax rates or tighten enforcement. We use Equation 5 to compute the welfare effect of a balanced-budget policy that increases taxes and reduces enforcement. Panel D of Figure 9 shows how the welfare effect of this policy varies with the tax elasticity $\varepsilon_t^{\text{Pay}}$. Since taxes are less distortionary when $\varepsilon_t^{\text{Pay}}$ is closer to zero, the welfare gains from this policy increase with $\varepsilon_t^{\text{Pay}}$. The dashed lines report the range of our estimates of $\varepsilon_t^{\text{Pay}}$ as well as the DiD estimate that pools data from all tax-rate changes. The blue solid line in Panel D shows that increasing tax rates and reducing the reliance on the public goods enforcement treatment raises welfare as long as $\varepsilon_t^{\text{Pay}}$ is greater than our smallest estimate of -0.46. The relative desirability of taxes over enforcement diminishes only slightly when we consider the sanctions treatment (red dotted line), when we allow for liquidity constraints (yellow dashed line), or when we allow for larger utility costs of consumption changes (green dashed line).[62] Panel D also shows the effect of under-valuing the welfare of delinquent households by varying ω. Lower values of ω mechanically lower the desirability of taxes over enforcement. At our central value of $\varepsilon_t^{\text{Pay}}$, assuming $\omega = 0.5$ (emerald dash-dot line) equalizes the welfare effects of both policy instruments in our setting.[63]

The results in this section show that the welfare costs of raising property tax rates are smaller than plausible magnitudes of the value of public goods. While liquidity constraints raise the welfare cost of taxation moderately, incorporating this force does not change the conclusion that tax hikes likely increase both revenue and welfare. In contrast, the private costs associated with enforcement likely exceed the value of public goods, so that enhanced enforcement is likely welfare reducing. While enforcement may be more effective in settings where tax hikes do not increase revenue or where the government does not value the welfare cost of non-compliance, we conclude that in our setting, a welfare-maximizing government would rather increase tax rates than enhance enforcement.

8 Conclusion

This paper studies the design of property taxes in developing countries in a holistic and detailed manner. Our work draws on administrative tax data from the universe of residential properties

[62] Figure A.4 considers additional scenarios for the consumption drop Δc^{Pay}. This graph considers setting the drop of consumption for liquidity constrained households equal to -4.2%, which is the 80th percentile of the ratio of property tax payments to income. It also shows the effect of setting the consumption drop of liquidity constrained households to -9.3%, which we estimate using a split-sample IV approach in Appendix I. The desirability of tax changes over enforcement is very similar for all scenarios. Tax changes are preferable to enforcement for most of the range of tax elasticity estimates.

[63] This line can also be interpreted as a robustness check on our estimation of the private cost of not paying taxes. At our central value of $\varepsilon_t^{\text{Pay}}$, taxes are more desirable than enforcement even if we over-estimated this cost by 100% of its true value. Figure A.3 plots the relative desirability of taxes over enforcement for additional values of ω.

in Mexico City and various quasi-experimental and experimental identification strategies. We examine the role of the two main tools used to raise property tax revenue—tax rate changes and enforcement—in addition to taking into account other tax system features such as payment schedules and payment modality.

We show that it is possible to raise property tax revenue through higher tax rates or through enforcement. Collectively, the variation that we study contributed to a 36 percent increase in total property tax revenue. We also find that taxpayers are liquidity constrained: consumption is only imperfectly insured against temporary unanticipated income shocks and an increase in the tax rate leads more taxpayers to pay in installments. Because our study is based in Mexico City, a setting that is similar to other developing countries in terms of the prevalence of household liquidity constraints and the level of administrative capacity, our findings carry broad relevance.

Our optimal tax model combines our empirical estimates of tax, enforcement, and consumption elasticities to quantify the optimal tax rate and compliance level. The model provides conditions determining when a welfare-maximizing government would prefer to increase tax rates or enhance enforcement. Since we calculate that our field experiment generated welfare losses, we conclude that it is not desirable to increase enforcement. In contrast, while liquidity constraints are present, they raise the welfare cost of tax increases only modestly. This implies that current tax rates are still below their optimal level. Our model shows that the relative desirability of tax rates versus enforcement depends on key empirical quantities that may vary across settings. Accurate measures of how policy changes affect both revenue and compliance costs would allow governments to improve property tax design.

Overall, our results reveal that details of property tax systems can have important impacts on taxpayer welfare and revenue collection. In particular, the optimal design and administration of loans for liquidity-constrained taxpayers—including terms of eligibility, interest rates, and payment schedules—is an important avenue for future research. Future work continuing to build on the tax systems paradigm will likely yield important insights in other settings (e.g., Okunogbe, 2021; Bergeron et al., 2023) and illuminate the roles of tax fairness and equity (e.g., Best et al., 2020).

References

Agrawal, David R., William H. Hoyt, and John D. Wilson, "Local Policy Choice: Theory and Empirics," 2020. Working Paper. 5, 12

Allcott, Hunt and Dmitry Taubinsky, "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market," *American Economic Review*, August 2015, *105* (8), 2501–38. 10, 46

_ **and Judd B. Kessler**, "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons," *American Economic Journal: Applied Economics*, January 2019, *11* (1), 236–76. 3, 6, 13, 47

_ **, Daniel Cohen, William Morrison, and Dmitry Taubinsky**, "When do "Nudges" Increase Welfare?," *NBER Working Paper 30740*, 2022. 3

Andreoni, James, "IRS as loan shark tax compliance with borrowing constraints," *Journal of Public Economics*, 1992, *49* (1), 35 – 46. 49

Angrist, Joshua D. and Alan B. Krueger, "The Effect of Age at School Entry on Educational Attainment: An Application of Instrumental Variables with Moments from Two Samples," *Journal of the American Statistical Association*, 1992, *87* (418), 328–336. 89

Athey, Susan and Guido W Imbens, "The econometrics of randomized experiments," in "Handbook of economic field experiments," Vol. 1, Elsevier, 2017, pp. 73–140. 98

Atinyan, Armenak and Zareh Asatryan, "Nudging for Tax Compliance: A Meta Analysis," 2022. Working Paper. 22, 23, 24

Atkinson, A. B. and N. H. Stern, "Pigou, Taxation and Public Goods12," *The Review of Economic Studies*, 01 1974, *41* (1), 119–128. 8

Attanasio, Orazio P, "Consumption," *Handbook of Macroeconomics*, 1999, *1*, 741–812. 25

— **and Miguel Székely**, "Wage shocks and consumption variability in Mexico during the 1990s," *Journal of Development Economics*, 2004, *73* (1), 1–25. 25

Balan, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan Weigel, "Local Elites as Tax Collectors: Experimental Evidence from the DRC," *American Economic Review*, 2022, *112* (3). 5

Basri, M. Chatib, Mayara Felix, Rema Hanna, and Benjamin A. Olken, "Tax Administration versus Tax Rates: Evidence from Corporate Taxation in Indonesia," *American Economic Review*, December 2021, *111* (12), 3827–71. 4, 8

Bell, Martin, Elin Charles-Edwards, Philipp Ueffing, John Stillwell, Marek Kupiszewski, and Dorota Kupiszewska, "Internal migration and development: comparing migration intensities around the world," *Population and Development Review*, 2015, *41* (1), 33–58. 12

Bergeron, Augustin, Gabriel Tourek, and Jonathan Weigel, "The State Capacity Ceiling on Tax Rates: Evidence from Randomized Tax Abatements in the DRC," 2023. Working Paper. 5, 8, 11, 29, 31

Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman, "What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment*," *The Quarterly Journal of Economics*, 02 2010, *125* (1), 263–306. 10, 46

Besley, Timothy and Torsten Persson, "Taxation and Development," in Alan J Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, eds., *Handbook of Public Economics*, Vol. 5 2013. 4

— **, Anders Jensen, and Torsten Persson**, "Norms, Enforcement, and Tax Evasion," *Review of Economics and Statistics*, 2023, *forthcoming.* 13, 95, 104

Best, Michael, Anne Brockmeyer, Henrik Kleven, Johannes Spinnewijn, and Mazhar Waseem, "Production vs Revenue Efficiency with Low Tax Compliance: Theory and Evidence from Pakistan," *Journal of Political Economy*, 2015, *123* (6), 1311–1355. 4

— **, Francois Gerard, Evan Kresch, and Joana Naritomi**, "Greener on the Other Side? Spatial Discontinuities in Property Tax Rates and their Effects on Tax Morale," 2020. Working Paper. 5, 31, 95

Blundell, Richard and Ian Preston, "Consumption inequality and income uncertainty," *The Quarterly Journal of Economics*, 1998, *113* (2), 603–640. 101

— **and Monica Costas Dias**, "Alternative approaches to evaluation in empirical microeconomics," *Journal of Human Resources*, 2009, *44* (3), 565–640. 97

Bradbury, Katharine L., Christopher J. Mayer, and Karl E. Case, "Property tax limits, local fiscal behavior, and property values: evidence from Massachusetts under Proposition 212," *Journal of Public Economics*, 2001, *80* (2), 287 – 311. 12

Brockmeyer, Anne and Marco Hernandez, "Taxation, Information and Withholding: Evidence from Costa Rica," 2022. Working Paper. 20

— **, Spencer Smith, Marco Hernandez, and Stewart Kettle**, "Casting a Wider Tax Net: Experimental Evidence from Costa Rica," *American Economic Journal: Economic Policy*, 2019, *11* (3), 55–87. 4

Brueckner, Jan K., "A test for allocative efficiency in the local public sector," *Journal of Public Economics*, 1982, *19* (3), 311 – 331. 12

Bugni, Federico A and Ivan A Canay, "Testing Continuity of a Density via g-order statistics in the Regression Discontinuity Design," *Journal of Econometrics*, 2020. 18, 64

Busso, Matias, Jesse Gregory, and Patrick Kline, "Assessing the Incidence and Efficiency of a Prominent Place Based Policy," *American Economic Review*, April 2013, *103* (2), 897–947. 6

Butera, Luigi, Robert Metcalfe, William Morrison, and Dmitry Taubinsky, "Measuring the Welfare Effects of Shame and Pride," *American Economic Review*, January 2022, *112* (1), 122–68. 6

Cabral, Marika and Caroline Hoxby, "The Hated Property Tax: Salience, Tax Rates, and Tax Revolts," Working Paper 18514, National Bureau of Economic Research November 2023. 5

Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik, "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 2014, *82* (6), 2295–2326. 19, 64

Caplin, Andrew, *Fear as a Policy Instrument*, Russell Sage, 2003. 6, 47

Card, David and Brian P. McCall, "Is Workers' Compensation Covering Uninsured Medical Costs? Evidence from the "Monday Effect"," *Industrial and Labor Relations Review*, 1996, *49* (4), 690–706. 89

Carpio, Lucia Del, "Are the Neighbors Cheating? Evidence from Social Norm Experiment on Property Taxes in Peru," 2014. Mimeo. 24

Cellini, Stephanie Riegg, Fernando Ferreira, and Jesse Rothstein, "The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design*," *The Quarterly Journal of Economics*, 02 2010, *125* (1), 215–261. 12, 28

Chamberlain, Gary, "Chapter 22 Panel data," in "in," Vol. 2 of Handbook of Econometrics, Elsevier, 1984, pp. 1247 – 1318. 92

Chetty, Raj, "Consumption Commitments, Unemployment Durations, and Local Risk Aversion," Working Paper 10211, National Bureau of Economic Research January 2004. 11, 12, 28, 90

_, "A general formula for the optimal level of social insurance," *Journal of Public Economics*, 2006, *90*, 1879–1901. 8

_, "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 2009, *1* (1), 451–488. 28

_ **and Adam Looney**, "Consumption smoothing and the welfare consequences of social insurance in developing economies," *Journal of Public Economics*, December 2006, *90* (12), 2351–2356. 28

_ **and Amy Finkelstein**, "Chapter 3 - Social Insurance: Connecting Theory to Data," in Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, eds., *handbook of public economics, vol. 5*, Vol. 5 of Handbook of Public Economics, Elsevier, 2013, pp. 111 – 193. 11

Chirico, Michael, Robert Inman, Charles Loeffler, John MacDonald, and Holger Sieg, "Deterring Property Tax Delinquency in Philadelphia: An Experimental Evaluation of Nudge Strategies," *National Tax Journal*, 2019, *72* (3), 479–506. 5

Clausing, Kimberly A., "Corporate tax revenues in OECD countries," *International Tax and Public Finance*, 2007, *14(2)*, 115â133. 92

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, 1991, *59* (5), 1221–1248. 11

Demirguc-Kunt, Asli, Leora Klapper, Singer Dorothee, Saniya Ansar, and Jake Hess, "The Global Findex Database: Measuring Financial Inclusion and the Fintech Revolution," Technical Report, World Bank 2017. 4

Devereux, Michael, "Developments in the Taxation of Corporate Profit in the OECD Since 1965: Rates, Bases, and Revenues," *Oxford University Centre of Business Taxation Working Paper No. 07/04*, 2007. 92

Dusek, Libor, Nicolas Pardo, and Christian Traxler, "Salience and Timely Compliance: Evidence from Speeding Tickets," *Journal of Policy Analysis and Management*, 2022, *41* (2), 426–449. 4

Dzansi, James, Anders Jensen, David Lagakos, and Henry Telli, "Technology and Tax Capacity: Evidence from Local Governments in Ghana," *NBER Working Paper 29923*, 2022. 5

Fang, Hanming and Dan Silverman, "On the compassion of time-limited welfare programs," *Journal of Public Economics*, 2004, *88* (7), 1445 – 1470. 84

Gadenne, Lucie, "Tax me, but spend wisely? Sources of public finance and government accountability," *American Economic Journal: Applied Economics*, 2017, pp. 274–314. 12

Glaeser, Edward L., “Paternalism and Psychology,” *University of Chicago Law Review*, 2006, *73* (1), 133–156. 6, 47

Gonzalez-Navarro, Marco and Climent Quintana-Domeque, “Paving Streets for the Poor: Experimental Analysis of Infrastructure Effects,” *The Review of Economics and Statistics*, 2016, *98* (2), 254–267. 12

Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano, “Do fiscal rules matter?,” *American Economic Journal: Applied Economics*, 2016, pp. 1–30. 19

Gruber, Jonathan, “The Consumption Smoothing Benefits of Unemployment Insurance,” *The American Economic Review*, 1997, *87* (1), 192–205. 89, 90

Hall, Robert E, “Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence,” *Journal of political economy*, 1978, *86* (6), 971–987. 102

Haughwout, Andrew, Robert Inman, Steven Craig, and Thomas Luce, “Local revenue hills: evidence from four US cities,” *Review of Economics and Statistics*, 2004, *86* (2), 570–585. 93

Hebous, Shafik and Michael Keen, “Pareto-Improving Minimum Corporate Taxation,” *IMF Working Papers*, 2021, *2021* (250). 7

Hendren, Nathaniel, “The Policy Elasticity,” *Tax Policy and the Economy*, 2016, *30* (1), 51–89. 8

_ and Ben Sprung-Keyser, “A Unified Welfare Analysis of Government Policies*,” *The Quarterly Journal of Economics*, 03 2020, *135* (3), 1209–1318. 8

Hotz, V. Joseph and Robert A. Miller, “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *The Review of Economic Studies*, 1993, *60* (3), 497–529. 84

Howard, Troup and Carlos Avenancio-Leon, “The Assessment Gap: Racial Inequalities in Property Taxation,” *Quarterly Journal of Economics*, 2022, *137* (3), 1383–1434. 14

International Monetary Fund, “World Revenue Longitudinal Data,” 2019. 38

Jack, William and Tavneet Suri, “Risk Sharing and Transactions Costs: Evidence from Kenya’s Mobile Money Revolution,” *American Economic Review*, January 2014, *104* (1), 183–223. 4

Jappelli, Tullio and Luigi Pistaferri, *The economics of consumption: theory and evidence*, Oxford University Press, 2017. 25

Kawano, Laura and Joel Slemrod, “How do corporate tax bases change when corporate tax rates change? With implications for the tax rate elasticity of corporate tax revenues,” *International Tax and Public Finance*, 2015, *23(3)*, 401â433. 92

Keen, Michael and Joel Slemrod, “Optimal Tax Administration,” *Journal of Public Economics*, July 2017, *152*, 133–142. 1, 4, 5, 8, 48

Kessler, Ryan, “Does Punishment Compel Payment? Driver’s License Suspensions and Fine Delinquency,” 2020. Working Paper. 4

Khan, Adnan Q., Asim I Khwaja, and Benjamin A. Olken, “Tax farming redux: Experimental evidence on performance pay for tax collectors,” *The Quarterly Journal of Economics*, 2016, *131* (1), 219–271. 5

_ , _ , and _ , “Making Moves Matter: Experimental Evidence on Incentivizing Bureaucrats through Performance-Based Postings,” *American Economic Review*, January 2019, *109* (1), 237–70. 5

Kroft, Kory and Matthew J. Notowidigdo, “Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence,” *The Review of Economic Studies*, 02 2016, *83* (3), 1092–1124. 89, 90

Lalive, Rafael, “How do extended benefits affect unemployment duration? A regression discontinuity approach,” *Journal of econometrics*, 2008, *142* (2), 785–806. 19

Lemieux, Thomas and Kevin Milligan, “Incentive effects of social assistance: A regression discontinuity approach,” *Journal of Econometrics*, 2008, *142* (2), 807–828. 19

Loewenstein, George and Ted O’Donoghue, “" We can do this the easy way or the hard way": Negative emotions, self-regulation, and the law,” *The University of Chicago Law Review*, 2006, *73* (1), 183–206. 6, 47

McCrary, Justin, “Manipulation of the running variable in the regression discontinuity design: A density test,” *Journal of Econometrics*, 2008, *142* (2), 698–714. 18, 66

Meiselman, Ben S., "Ghostbusting in Detroit: Evidence on nonfilers from a controlled field experiment," *Journal of Public Economics*, 2018, *158*, 180 – 193. 4, 6

Morduch, Jonathan, "Income Smoothing and Consumption Smoothing," *Journal of Economic Perspectives*, September 1995, *9* (3), 103–114. 4

__ **and Dean Karlan**, "Access to Finance," *Handbook of Development Economics*, 06 2009, *5.* 4

Mullainathan, Sendhil, Joshua Schwartzstein, and William J. Congdon, "A Reduced-Form Approach to Behavioral Public Finance," *Annual Review of Economics*, 2012, *4* (1), 511–540. 46

Naritomi, Joana, "Consumers as Tax Auditors," *American Economic Review*, 2019, *109* (9), 3031–72. 4

Okunogbe, Oyebola, "Becoming Legible to the State: The Role of Detection and Enforcement Capacity on Tax Compliance," 2021. Working Paper. 5, 24, 31

Organ, Paul R., Alex Ruda, Joel Slemrod, and Alex Turk, "Incentive effects of the IRSâ passport certification and revocation process," *Journal of Public Economics*, 2022, *208*, 104625. 4, 6

Oster, Emily, "Unobservable selection and coefficient stability: Theory and evidence," *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204. 99

Perez-Truglia, Ricardo and Ugo Troiano, "Shaming tax delinquents," *Journal of Public Economics*, 2018, *167*, 120 – 137. 4

Pomeranz, Dina, "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax," *American Economic Review*, 2015, *105* (8), 2539–2569. 4

__ **and José Vila-Belda**, "Taking State-Capacity Research to the Field: Insights from Collaborations with Tax Authorities," *Annual Review of Economics*, 2019, *11* (1), 755–781. 4

Silva, Joao Santos and Silvana Tenreyro, "The Log of Gravity," *The Review of Economics and Statistics*, 2006, *88* (4), 641–658. 20

Singh, Divya, "Do Property Tax Incentives for New Construction Spur Gentrification? Evidence from New York City," 2020. Working Paper. 78

Singhal, Monica and Erzo Luttmer, "Tax Morale," *Journal of Economic Perspectives*, 2014, *28* (4), 149–168. 6

Slemrod, Joel, "Buenas notches: lines and notches in tax system design," *eJTR*, 2013, *11*, 259. 84

__ , "Tax Compliance and Enforcement," *Journal of Economic Literature*, December 2019, *57* (4), 904–54. 4

__ **and Christian Gillitzer**, *Tax Systems*, MIT Press, 2013. Zeuthen Lecture Series. 4

__ **and Shlomo Yitzhaki**, "Integrating Expenditure and Tax Decisions:The Marginal Cost of Funds and the Marginal Benefit of Projects," *National Tax Journal*, June 2001, *54* (2), 189–202. 8

Suárez Serrato, Juan Carlos and Owen Zidar, "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms," *American Economic Review*, 2016, *106* (9), 2582–2624. 92

__ **and** __ , "The structure of state corporate taxation and its impact on state tax revenues and economic activity," *Journal of Public Economics*, 2018, *167*, 158 – 176. 92

__ **and Philippe Wingender**, "Estimating the Incidence of Government Spending," December 2014. Mimeo. 28

Townsend, Robert M., "Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies," *Journal of Economic Perspectives*, September 1995, *9* (3), 83–102. 4

Train, Kenneth, *Discrete Choice Methods with Simulation* 2009. 85

Versprille, Allyson, "Growing Collection Woes May Haunt IRS Years After Pandemic Ends," *Bloomberg*, May 2020. 4

Weigel, Jonathan, "The Participation Dividend of Taxation: How Citizens in Congo Engage More with the State When it Tries to Tax Them," *Quarterly Journal of Economics*, 2020, *Forthcoming.* 5, 24

Wong, Francis, "Mad as Hell: Property Taxes and Financial Distress," Working Paper 2020. 5

World Bank, "World Development Indicators," 2019. 38

Tables

Table 2: The Effect of Tax Rates on Tax Payment

	A. Regression Discontinuity				B. Difference-in-Differences			
	Mean Tax Rate (1)	Payment Amount (2)	Payment in Full (3)	Compliance Share (4)	Mean Tax Rate (5)	Payment Amount (6)	Payment in Full (7)	Compliance Share (8)
I. Estimates for the 2009-2010 Treatment								
Level Effect	9.127***	.584**	-5.483***	-3.208	8.319***	.677***	-3.772***	-1.963***
	(.059)	(.240)	(2.122)	(2.043)	(.000)	(.099)	(.623)	(.527))
Mean at Baseline (Treated)	50.112	5.836	36.626	47.881	50.112	5.836	36.626	47.881
Relative Effect	.182	.100	-.150	-.067	.166	.116	-.103	-041
	(.001)	(.041)	(.058)	(.043)	(.000)	(.017)	(.017)	(.011)
Implied Elasticity		.55	-.822	-.368		.697	-.617	-.244
		(.226)	(.318)	(.234)		(.101)	(.104)	(.067)
II. Estimates for the 2010-2011 Treatment								
Level Effect	12.147***	.575***	-6.433***	-4.939***	11.011***	.800***	-5.226***	-3.323***
	(.031)	(.131)	(1.523)	(1.296)	(.000)	(.095)	(.526)	(.475)
Mean at Baseline (Treated)	47.461	4.734	35.072	47.478	47.461	4.734	35.072	47.478
Relative Effect	.256	.121	-.183	-.104	.232	.169	-.149	-.070
	(.000)	(.027)	(.043)	(.027)	(.000)	(.017)	(.017)	(.011)
Implied Elasticity		.475	-.717	-.407		.728	-.642	-.300
		(.108)	(.170)	(.107)		(.085)	(.064)	(.043)
III. Estimates for the 2011-2012 Treatment								
Level Effect	18.002***	.452***	-10.949***	-6.228***	16.465***	.644***	-11.391***	-8.169***
	(.024)	(.085)	(1.387)	(1.185)	(.000)	(.095)	(.532)	(.449)
Mean at Baseline (Treated)	41.06	3.287	37.969	44.885	41.06	3.287	37.969	44.885
Relative Effect	.438	.137	-.288	-.138	.401	.196	-.300	-.182
	(.000)	(.025)	(.036)	(.026)	(.000)	(.029)	(.014)	(.010)
Implied Elasticity		.314	-.658	-.316		.489	-.747	-.455
		(.059)	(.083)	(.060)		(.073)	(.035)	(.025)

Notes: This table reports results from the RD and DiD estimations discussed in Section 4. Each year, properties in a specific value band are treated with a large tax rate increase. The treated value bands are I, H, and G in the years 2010, 2011, and 2012, respectively. In the RD, we compare these properties to properties just below the lower threshold of the treated value band. The estimating equation is $\Delta Y_{i,t} = \alpha + \beta T_i + f(\hat{V}_i) + g(\hat{V}_i)T_i + \epsilon_{i,t}$, where $\hat{V}_i$ denotes the distance between the value of property i and the lower limit of the treated band, T_i indicates properties in the treated band, and f and g are third-order polynomial functions. Standard errors are robust to heteroskedasticity and clustered at the postcode level. In the DiD, we compare treated properties to properties in value bands K and L. The estimating equation is $Y_{it} = \alpha + DD_{it}\beta + \gamma_i + \delta_t + \epsilon_{it}$, where DD_{it} indicates treated properties in post-reform years and δ_t and γ_i denote year and property fixed effects. We estimate this equation in levels and transform the point estimates into relative effects, scaling them by the treatment group mean in the last pre-reform year. Standard errors are robust to heteroskedasticity and clustered at the property level. The number of {treated properties, control properties} in the RD is {6763, 11101} in 2010, {11039, 17055} in 2011, and {16890, 31948} in 2012; in the DiD (which uses a balanced sample) it is {5747, 6510} in 2010, {9661, 6510} in 2011, and {15227, 6510} in 2012. Outcomes: payment amount is in MXN thousands, full payment is in percentage points, compliance share is tax payment over liability X 100. Figures 3–6 present the RD and DiD estimations graphically. Appendices D and E present various identification and robustness tests. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 3: The Effect of Enforcement Letters on Tax Payment – Experimental Estimates

1: Letter Content		2: Sender Position		3: Male Sender		4: Female Sender	
Sanctions	Public Good	Fiscal Attorney	Compliance Officer	Male	Neutral	Female	Neutral
Outcome A: Any Payment							
9.364***	4.858***	7.598***	6.624***	7.038***	8.15***	6.772***	6.477***
(.29)	(.274)	(.284)	(.281)	(.332)	(.337)	(.329)	(.329)
.000		.000		.002		.390	
Outcome B: Payment Amount (MX Pesos)							
54.226***	16.051***	42.416***	27.861***	39.707***	45.124***	29.025***	26.698***
(2.593)	(2.297)	(2.496)	(2.407)	(2.979)	(2.942)	(2.807)	(2.811)
.000		.000		.089		.422	

Notes: This table evaluates the effect of enforcement letters, as discussed in Section 5, on payment of outstanding tax debt, contrasting the different treatments summarized in Figure G.1. Each of Columns (1)-(4) and each panel corresponds to one regression. We estimate $Y_i = \alpha + \beta_1 T1_i + \beta_2 T2_i + \epsilon_i$, where Y_i is the outcome for property i evaluated 40 days after all letters were sent, α is a constant, $T1_i$ and $T2_i$ are dummies indicating the two mutually exclusive treatments (i.e., the sanctions treatment and the public goods treatment), and ϵ_i is the error term. The outcome is any payment in Panel A and the payment amount in Panel B. The outcomes are cumulative over time for each property. The bottom line of each panel reports the p-values from a Wald test of significant differences between each pair of treatment estimates. Figure 7 presents accompanying non-parametric evidence for the effect of the sanctions treatment and public goods treatment. Since the treatment and control groups exhibit slightly different trends prior to the intervention, as shown in Figure G.2, we display the data and run our estimations on detrended data. To do so, we estimate the following regression on the pre-intervention data: $Y_{igt} = \mu_g \cdot t + \alpha_i + \lambda_t + \epsilon_{igt}$, where t indicates days and g treatment groups. We then subtract the trend $\mu_g \cdot t$ from each treatment group. In all estimations, weekends are excluded from the sample. Payment amounts are winsorized at the 99th percentile. Standard errors in parentheses are robust to heteroskedasticity and are clustered at the property level. Table G.2 shows that the results are robust to including property characteristics as controls and estimating the treatment effects via difference-in-differences. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 4: Income, Consumption, and Property Tax Payment Variability in Mexico City

	$\text{var}_g[\log(\text{Consumption}_{it})]$		$\text{var}_g[\log(1+\text{Payment}_{it})]$		$\text{var}_g[\mathbb{1}_{\{\text{Payment}_{it}>0\}}]$	
	Levels	Diffs.	Levels	Diffs.	Levels	Diffs.
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{var}_g[\log(\text{Income}_{it})]$	.685 ***	.588 ***	.837 ***	.96 ***	.076 *	.091 **
	(.181)	(.191)	(.307)	(.347)	(.042)	(.041)
Outcome mean	.497	.016	2.026	-.012	.124	-.008
Group dummies	Y	N	Y	N	Y	N
N	167	115	167	115	167	115
R^2	.351	.337	.09	.112	.047	.069

Notes: This table presents the results of estimating Equation (9) using the ENIGH households survey for Mexico City. Income, consumption and property tax payments are measured per capita at the household level. The data is divided into 192 bins (8 year-of-birth categories × 6 education categories × 4 survey waves). Of these, 25 bins are dropped as they do not contain any observations. Standard errors are clustered by bin. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. This table is discussed in Section 6.1.

Figures

Figure 1: Ratio of Tax Revenue to GDP in High-Income vs Lower-Income Countries

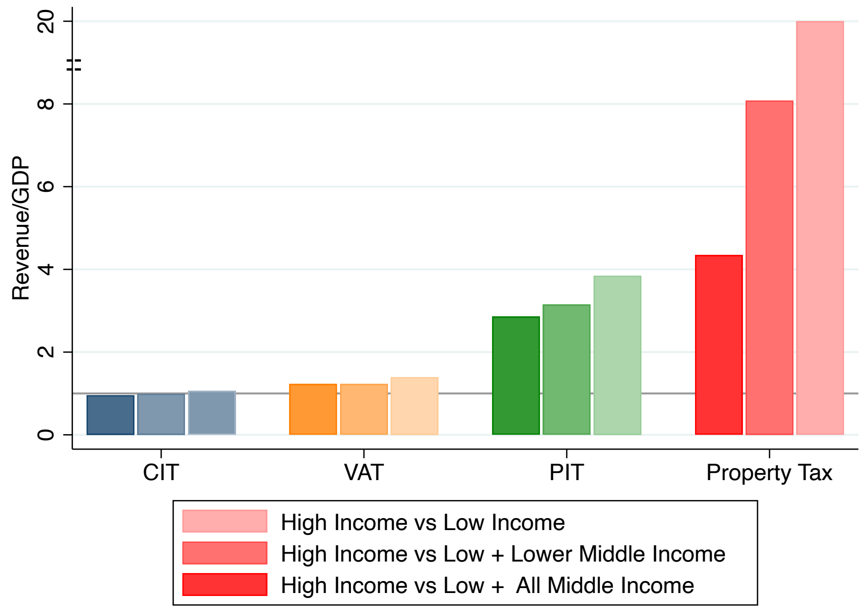

Notes: This figure shows the ratio of tax revenue as a share of GDP in high-income versus lower-income countries, as discussed in the introduction, for corporate income tax (CIT), value-added tax (VAT), personal income tax (PIT), and property taxes in 2017. The data are from the International Monetary Fund (2019) World Revenue Data and the World Bank (2019) World Development Indicators. The country income classification follows that of the World Bank Atlas methodology.

Figure 2: Variation in Tax Rates over Time and across Cadastral Value Bands, 2008–2012

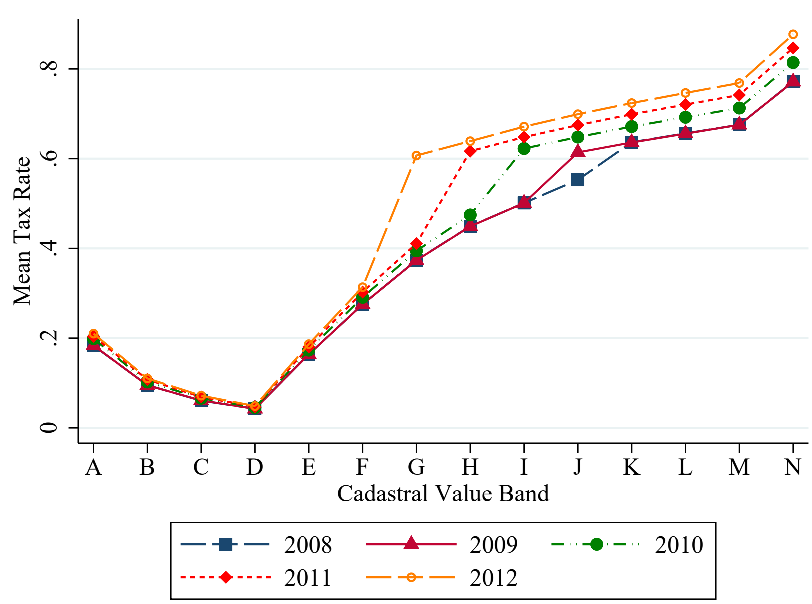

Notes: This figure, discussed in the introduction, shows the mean tax rate in percentage points by cadastral value band and year. We construct this figure using administrative tax bills.

Figure 3: The Effect of Tax Rates on Tax Payment – Regression Discontinuity Estimates

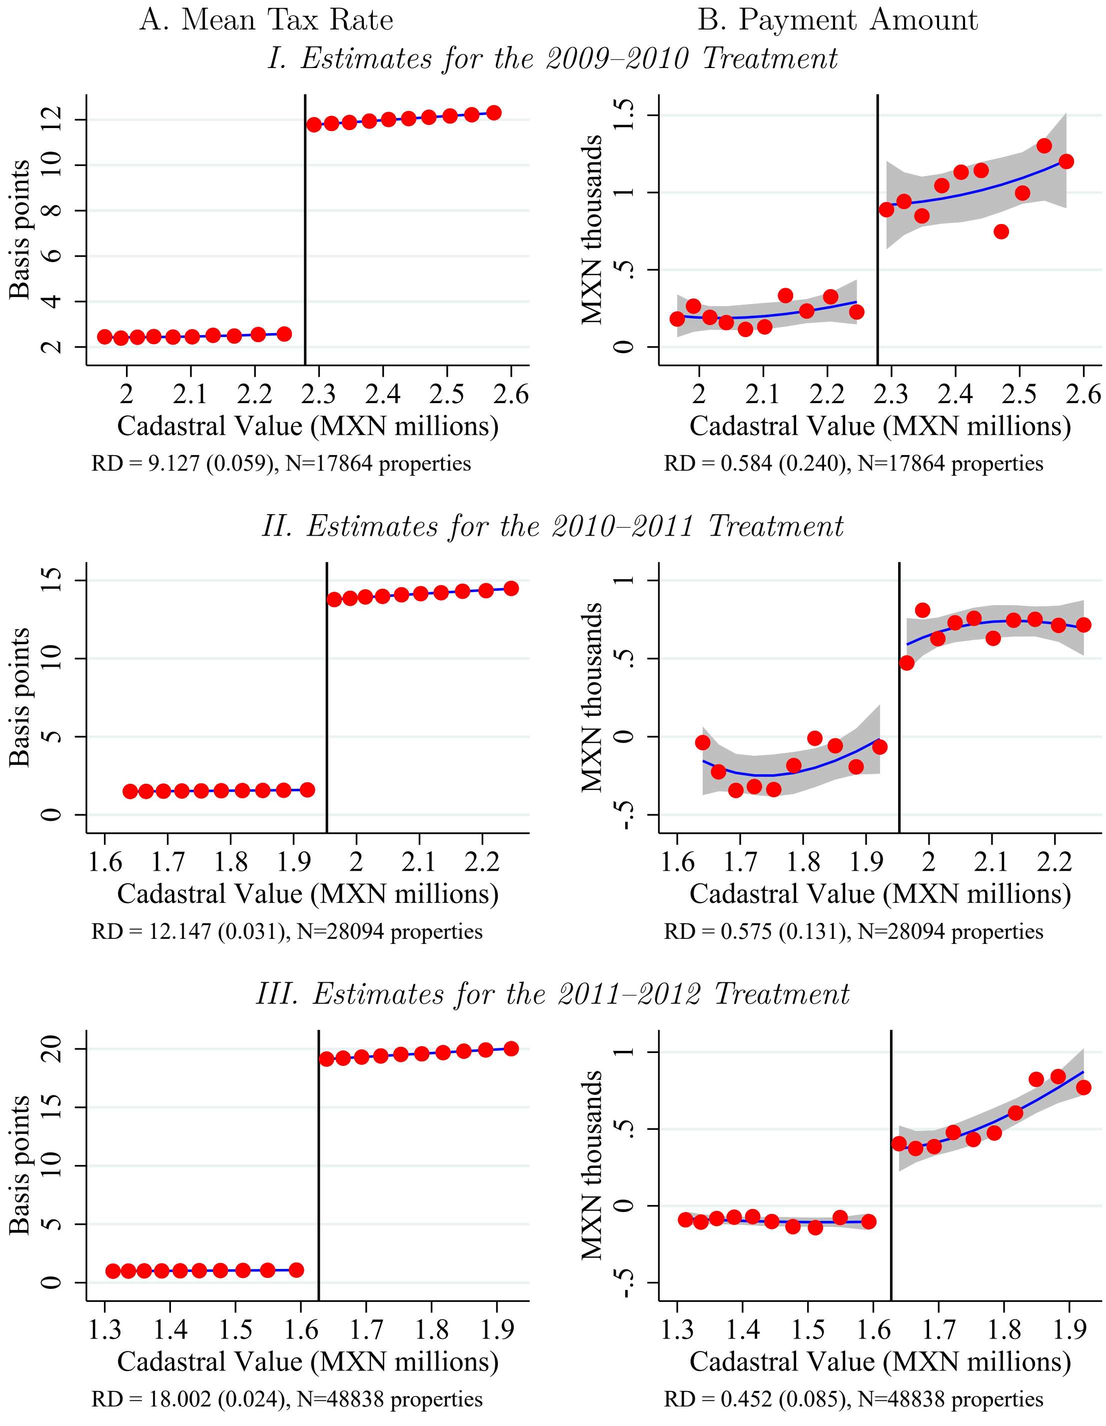

Notes: These graphs implement the RD estimation from Section 4.1. The red dots represent the mean outcome ($\Delta Y_{i,t}$)in equally spaced cadastral value bins. The solid blue lines (gray areas) depict a fitted third-order polynomial (the corresponding 95 percent confidence intervals). The vertical black lines mark the thresholds between the control and treatment bands. Properties to the right of the threshold are treated with a tax rate increase. We use the band below as the counterfactual since properties in the band above the treated band are treated in the previous year. The treated value bands are I, H and G in the years 2010, 2011 and 2012, respectively. The notes display the estimate for β from $\Delta Y_{i,t} = \alpha + \beta T_i + f(\hat{V}_i) + g(\hat{V}_i)T_i + \epsilon_{i,t}$, where $\hat{V}_i$ denotes the distance between the value of property i and the lower limit of the treated band, T_i indicates properties in the treated band, and f and g are third-order polynomial functions. Standard errors are robust to heteroskedasticity and clustered at the postcode level. Table 2 summarizes the estimates and implied elasticities. Table D.2 shows the robustness to local linear regressions with optimal bandwidth. Figure D.2 shows the robustness to varying bandwidths and degrees of polynomial.

Figure 4: The Effect of Tax Rates on Compliance – Regression Discontinuity Estimates

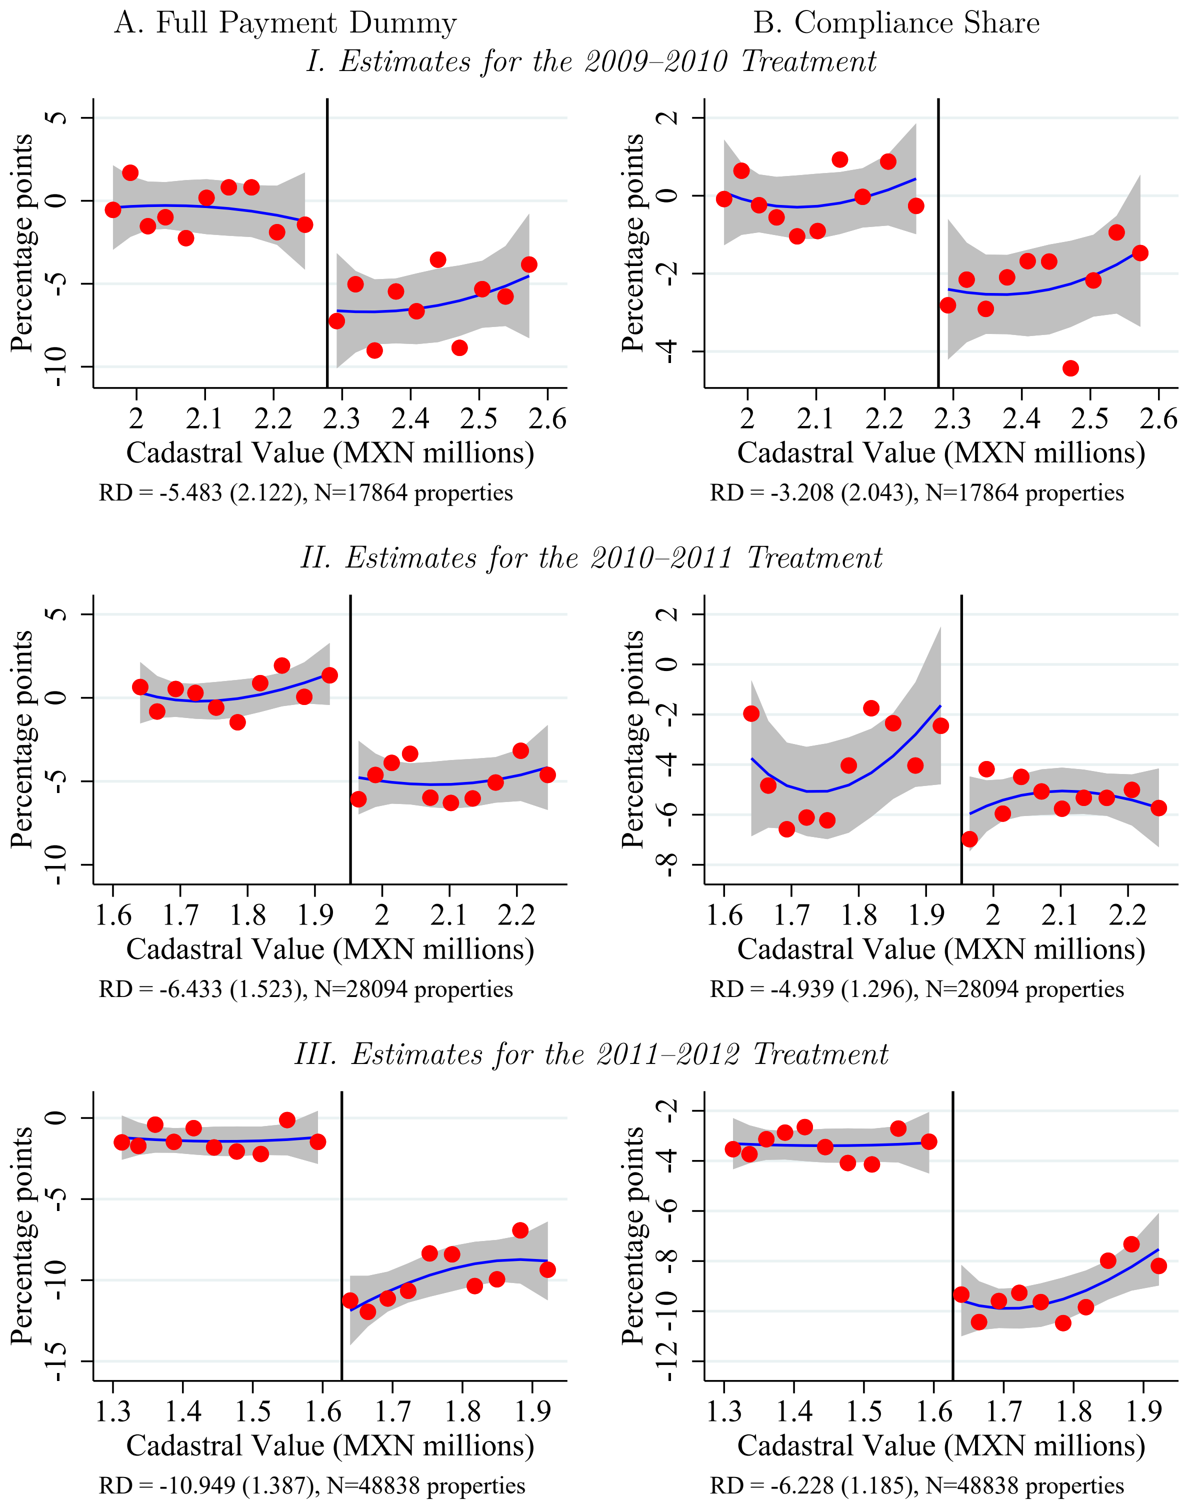

Notes: This figure is identical to Figure 3 but displays the results for different outcomes: a dummy indicating that taxpayers paid their liability fully and on time and the compliance share, defined as the tax payment divided by the liability. In each panel, the y-axis is in differences, showing $\Delta Y_{i,t}$.

Figure 5: The Effect of Tax Rates on Tax Payment – Difference-in-Difference Estimates

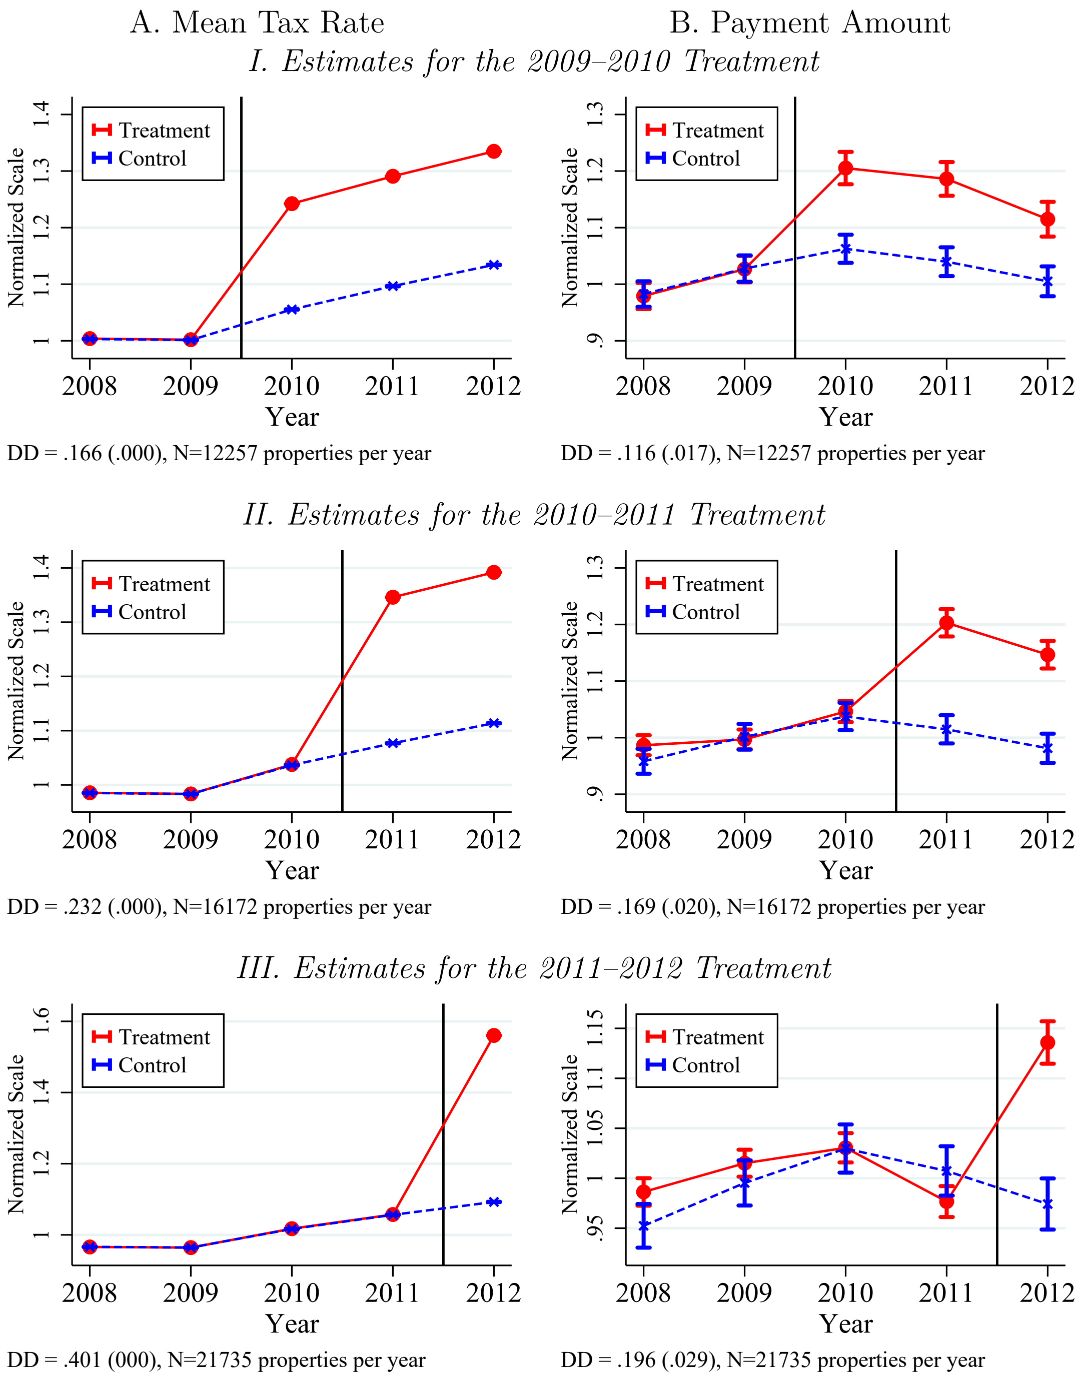

Notes: These graphs implement the DiD estimation from Section 4.2. Treatment and control group outcomes are normalized by their pre-reform mean. The vertical black lines mark the treatment timing. The notes display the estimate for β from $Y_{it} = \alpha + DD_{it}\beta + \gamma_i + \delta_t + \epsilon_{it}$, where DD_{it} indicates treated properties in post-reform years and δ_t and γ_i denote year and property fixed effects. We estimate this equation in levels and transform the point estimates into relative effects, scaling them by the treatment group mean in the last pre-reform year. Standard errors are robust to heteroskedasticity and clustered at the property level. The treated value bands are I, H and G in the years 2010, 2011 and 2012, respectively. The control group is composed of properties in bands K and L. Table 2 summarizes the estimates and implied elasticities. Appendix E presents various robustness tests.

Figure 6: The Effect of Tax Rates on Compliance – Difference-in-Difference Estimates

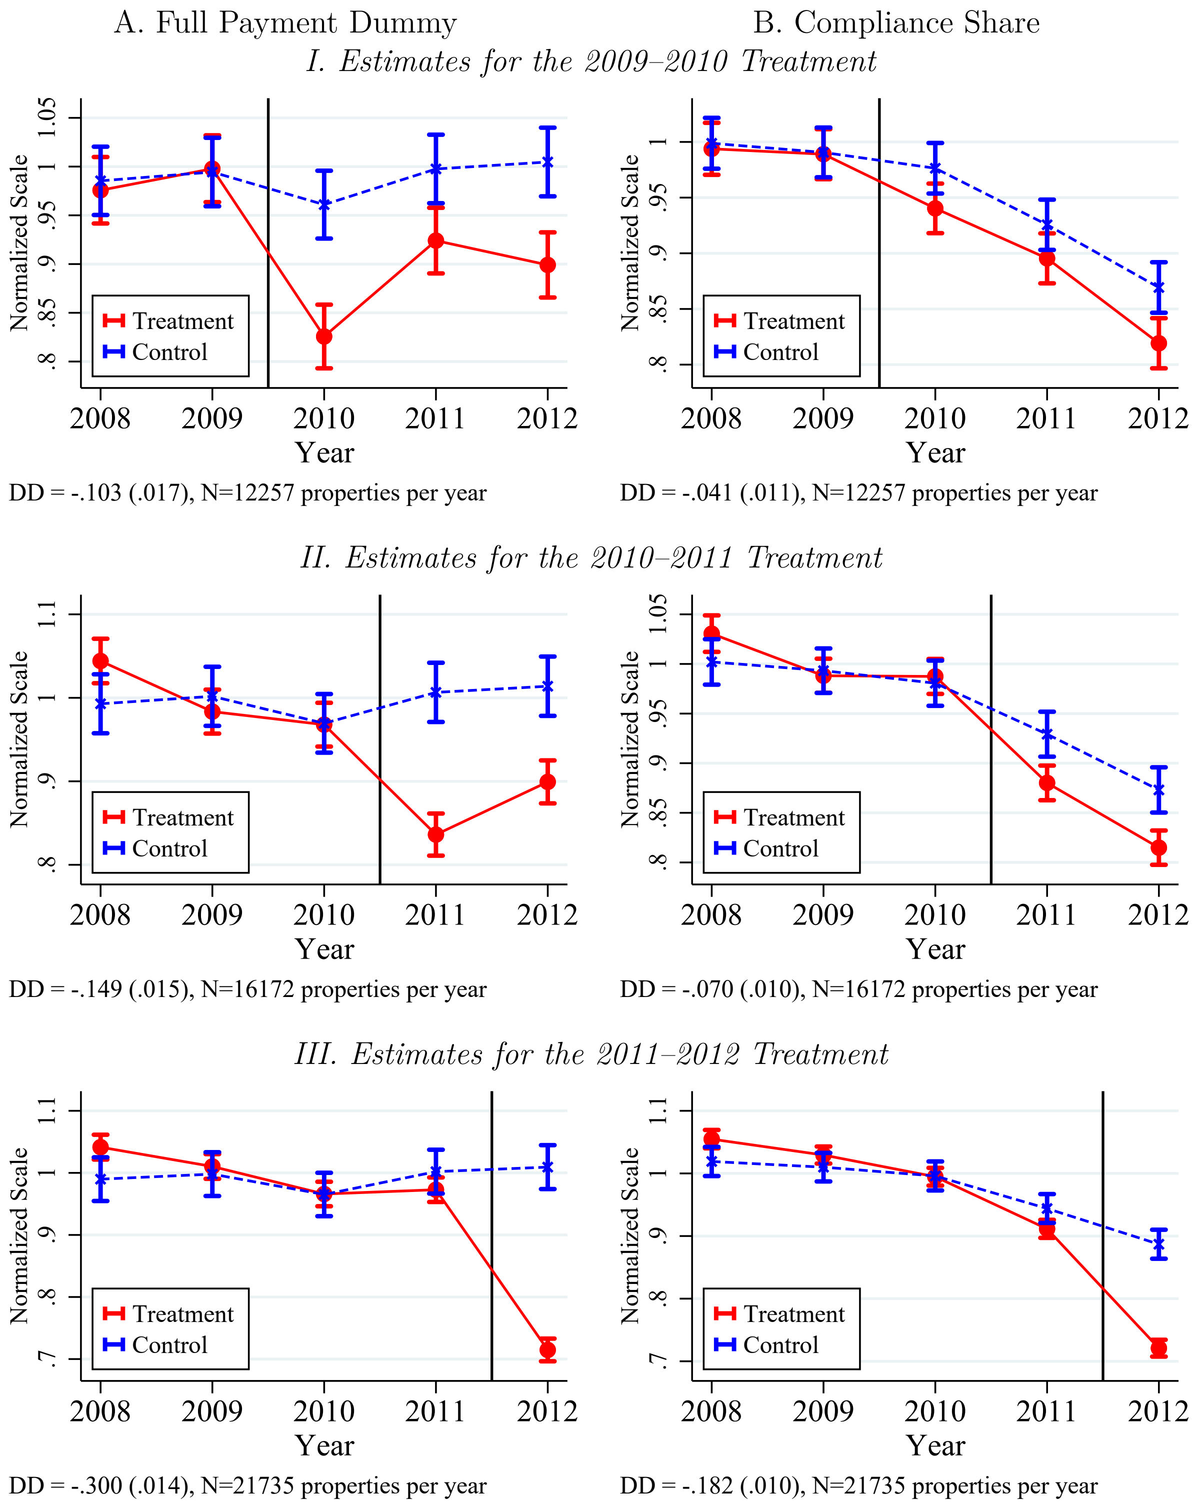

Notes: This figure is identical to Figure 5 but displays the results for different outcomes: a dummy indicating that taxpayers paid their liability fully and on time and the compliance share, defined as the tax payment divided by the liability.

Figure 7: The Effect of Enforcement Letters on Tax Payment

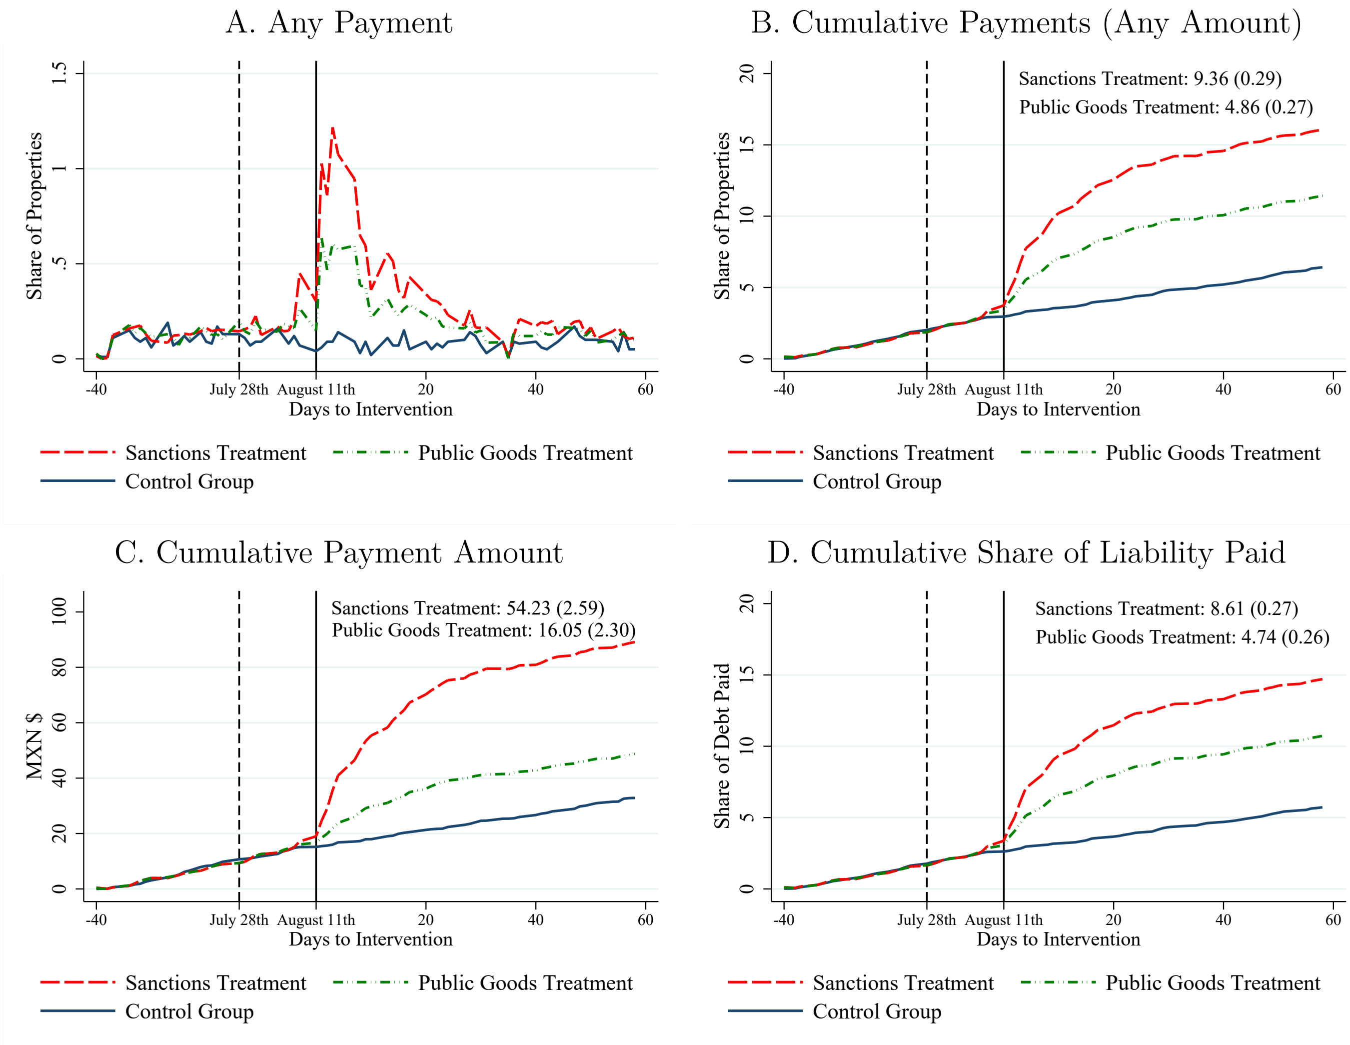

Notes: This figure displays taxpayers' response to enforcement letters, as discussed in Section 5. Panel A shows the share of properties for which a payment was made on any given day around the time of the enforcement intervention. Panel B shows the cumulative share of properties for which a payment was made, Panel C shows the average cumulative payment amount, and Panel D displays the share of the outstanding liability paid. We consider payments made between July and November 2014 against outstanding debt for the period from bimester 4 of 2009 to bimester 3 of 2014. The period during which the letters were sent—July 28 to August 11, 2014—is represented by the vertical lines. Panels B–D display the point estimates β_1 and β_2 from the OLS regression $Y_i = \alpha + \beta_1 T1_i + \beta_2 T2_i + \epsilon_i$, where Y_i is the outcome for property i evaluated 40 days after all letters were sent and $T1_i$ and $T2_i$ are dummies for the sanctions treatment and the public goods treatment, respectively. Since the treatment and control groups exhibit slightly different trends prior to the intervention, as shown in Figure G.2, we display here and run our estimations on detrended data. To do that, we run the following regression on the pre-intervention data: $Y_{igt} = \mu_g \cdot t + \alpha_i + \lambda_t + \epsilon_{igt}$, where t indicates days and g treatment groups. We then subtract the trend $\mu_g \cdot t$ from each treatment group. This is reasonable because the pre-intervention trend is indeed almost perfectly linear and the control group trend continues linearly after the intervention. In all estimations, weekends are excluded from the sample. Payment amounts are winzorized at the 99th percentile. Standard errors are robust to heteroskedasticity and are clustered at the property level. Table 3 presents regression estimates evaluating all treatment arms of the intervention, and Table G.2 shows the robustness of these results to controlling for property characteristics in the estimation and to estimation via difference-in-difference.

Figure 8: The Effect of Tax Rates on Tax Payment Modality

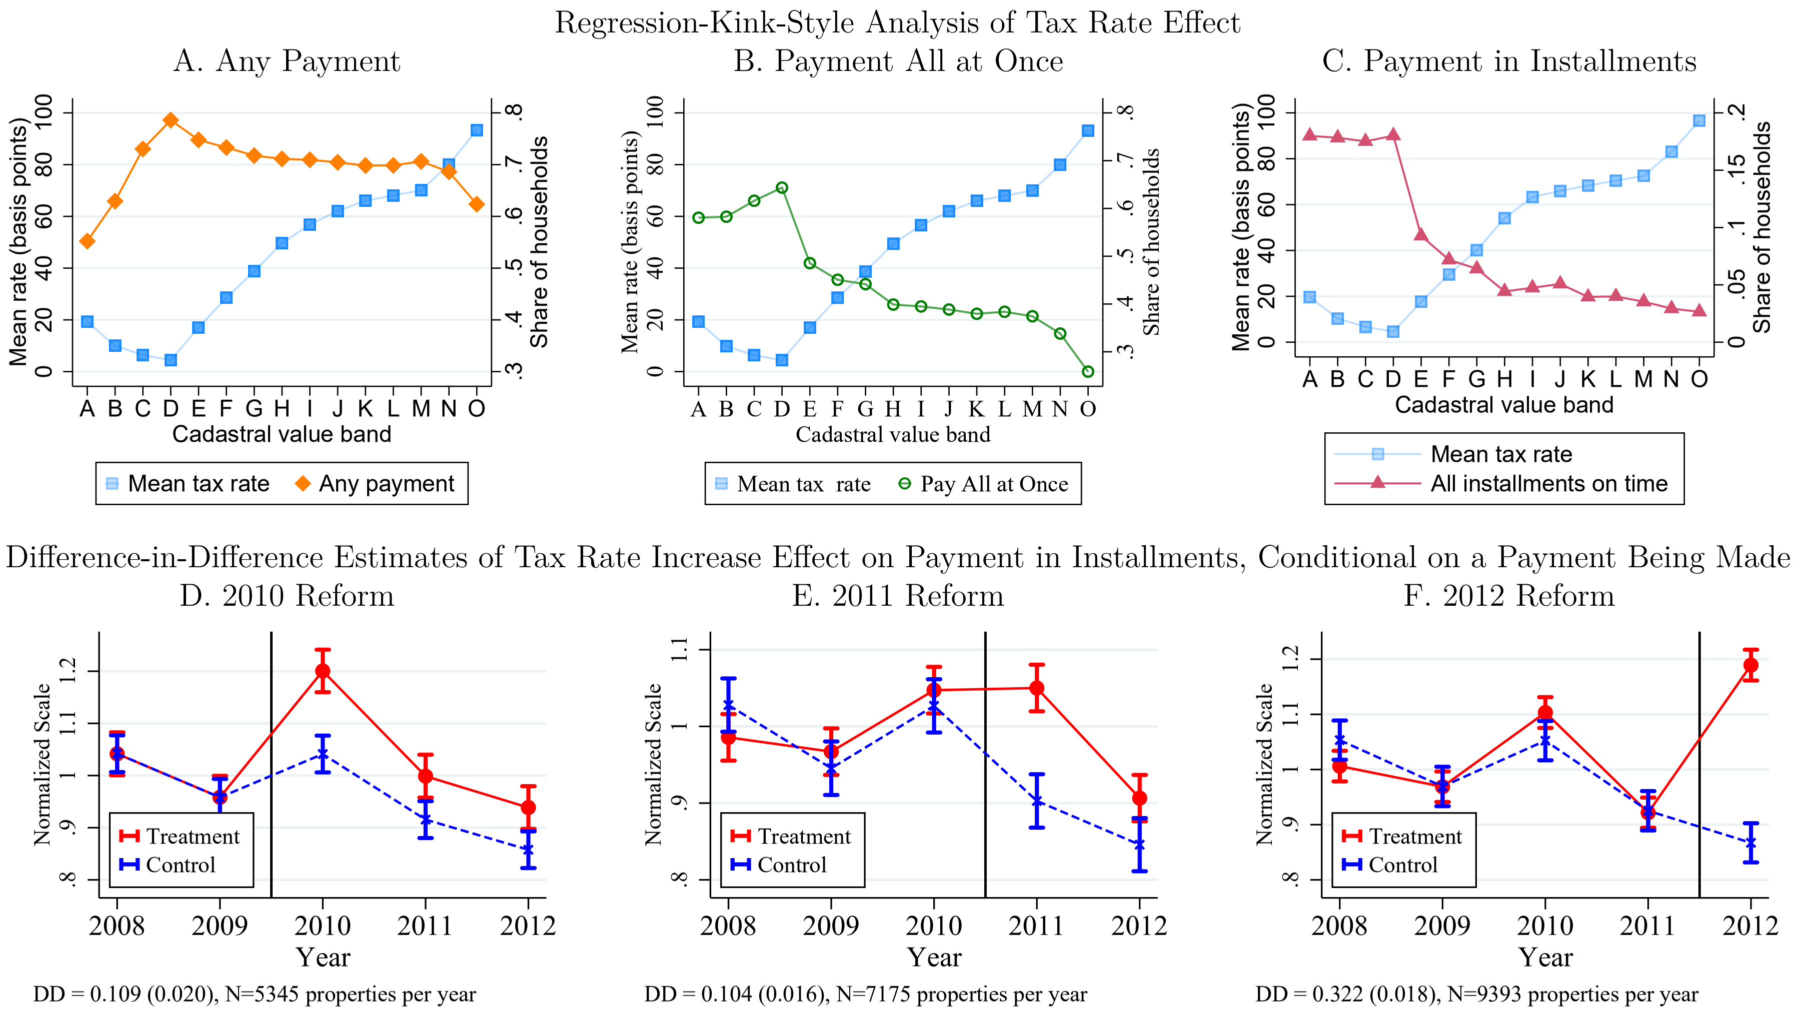

Notes: These graphs examine the effect of tax rates on payment modality, as discussed in Section 6.2. Panels A–C display the mean tax rate and households' payment modality choice by cadastral value bands. In all panels, the blue squares indicate the mean tax rate. In addition, Panel A shows the share of households that made any payment towards their tax liability, regardless of the payment modality. For households that made a payment, Panel B shows the share that paid their liability in full all at once (as opposed to attempting to pay in installments). For households that attempted to pay in installments, Panel C displays the share that made all six installment payments on time, hence remaining compliant (as opposed to households that paid only partially or late). Panels D–F are identical to Figures 5 and 6, with the sample restricted to taxpayers that register a payment (including partial payments) every year. The outcome is a dummy indicating whether the taxpayer paid in installments as opposed to paying all at once.

Figure 9: Welfare Analysis

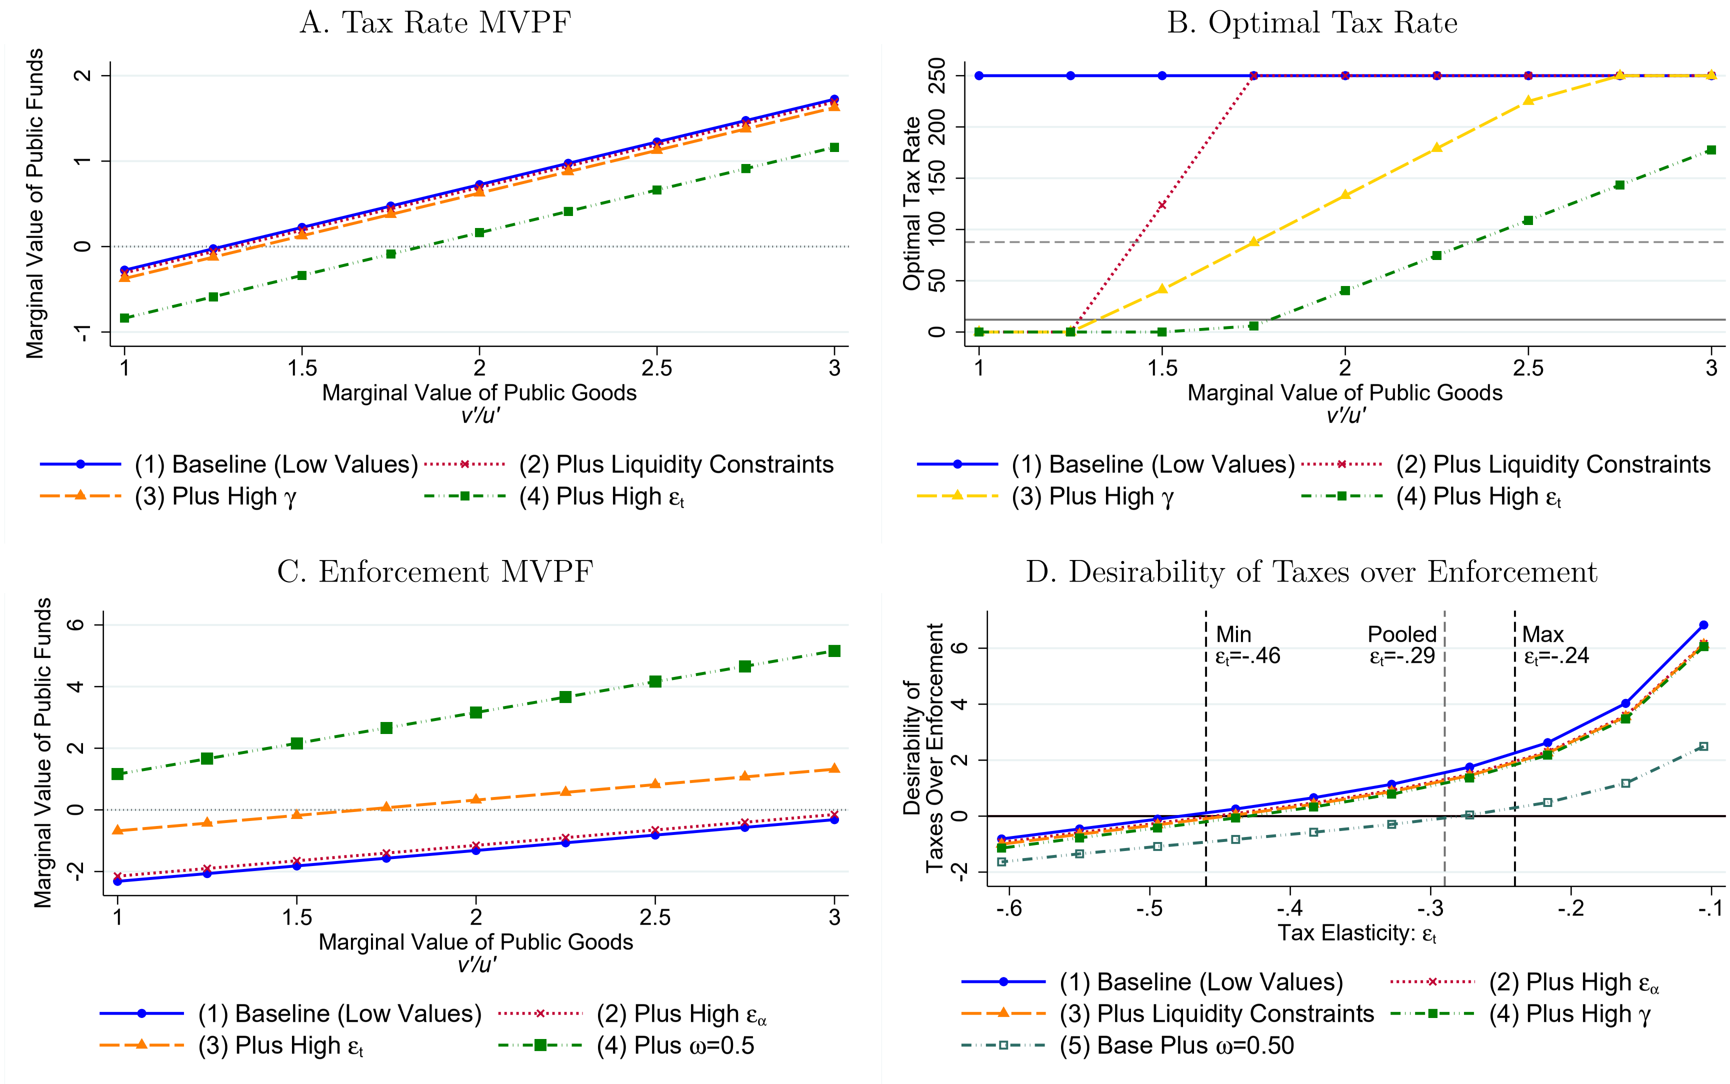

Notes: This figure displays results from the welfare analysis discussed in Section 7. Panel A plots the tax rate $MVPF_t$. Panel B plots the welfare-maximizing tax rate. The optimal tax rate is top-coded at 250 basis points. The horizontal solid line in Panel B represents the observed average tax rate between 2008 and 2012, while the dashed line represents the max tax rate in the same period. Panel C plots the enforcement $MVPF_\alpha$. Panel D plots the welfare impact of a balanced-budget policy that increases taxes and reduces enforcement, as in Equation 5. The dashed lines represent estimates of $\varepsilon_t^{\text{Pay}}$ from Table 2 and Table E.5.

Online Appendix: Not for Publication

This appendix contains additional information and analyses. Appendix A provides additional model results. Appendix B includes additional contextual information on property taxes in Mexico City. Appendix C presents summary statistics on the data that we use. We present additional details for the regression discontinuity analysis in Appendix D, for the difference-in-difference analysis in Appendix E, for the real response estimations in Appendix F, for the field experiment in Appendix G, for the analysis of payment modality and timing in Appendix H. In addition, we present an instrumental variable estimation in Appendix I, a calculation of the revenue-maximizing rate in Appendix J, an analysis of spillovers in Appendix K, a structural model of the income process in Appendix L, and extension of our model to include social norms in Appendix M.

A Model Appendix

This section expands on Section 1 by presenting additional derivations and results.

A.1 Approximating Marginal Utility

For a given individual, we approximate marginal utility with a first-order Taylor expansion:

$$u'(c) \approx u'(\bar{c}) + u''(\bar{c}) \times (c - \bar{c}) = u'(\bar{c})[1 - \gamma \times \Delta c],$$

where Δc is the percentage change in consumption (i.e., $\Delta c \leq 0$) and $\gamma = -\frac{u''(\bar{c})\bar{c}}{u'(\bar{c})}$ is the coefficient of relative risk aversion and captures the curvature of utility.

We now approximate the average marginal utility. Letting $\bar{c} = \bar{c}_0^{\text{Pay}} \pi_0^{\text{Pay}} + \bar{c}_s^{\text{Pay}} \pi_s^{\text{Pay}}$ be the average consumption across the two types of households, we express the average marginal utility as:

$$\begin{aligned}
\pi_s^{\text{Pay}} u'(c_s^{\text{Pay}}) + \pi_0^{\text{Pay}} u'(c_0^{\text{Pay}}) &\approx \pi_0^{\text{Pay}}[u'(\bar{c}) + u''(\bar{c})(c_0^{\text{Pay}} - \bar{c})] + \pi_s^{\text{Pay}}[u'(\bar{c}) + u''(\bar{c})(c_s^{\text{Pay}} - \bar{c})] \\
&= u'(\bar{c}) + u''(\bar{c})(\pi_0^{\text{Pay}} c_0^{\text{Pay}} + \pi_s^{\text{Pay}} c_s^{\text{Pay}} - \bar{c}) \\
&= u'(\bar{c})[1 - \gamma \underbrace{(\pi_{0,c}^{\text{Pay}} \Delta c_0^{\text{Pay}} + \pi_{s,c}^{\text{Pay}} \Delta c_s^{\text{Pay}})}_{\equiv \Delta c}] \\
&= u'(\bar{c})[1 - \gamma \Delta c],
\end{aligned}$$

where $\pi_{0,c}^{\text{Pay}} = \frac{\bar{c}_0 \pi_0^{\text{Pay}}}{\bar{c}_0 \pi_0^{\text{Pay}} + \bar{c}_s \pi_s^{\text{Pay}}}$ is the consumption share of liquidity-constrained households. The term Δc represents the average drop in consumption across taxpayers with different values of s. Similarly, in Equation 2, we rely on the average tax semi-elasticity of consumption: $\eta_t^c = \pi_{s,c}^{\text{Pay}} \eta_{t,s}^c + \pi_{0,s}^{\text{Pay}} \eta_{t,0}^c$.

A.2 Measuring $m'(\alpha)$

One drawback of Equation 3 is that we do not directly observe the welfare cost of additional enforcement, $m'(\alpha)$. Building on previous approaches (e.g., Bertrand et al., 2010; Mullainathan et al., 2012; Allcott and Taubinsky, 2015), we express $m'(\alpha)$ as the utility cost of a tax with the same impact on N^{Pay}. First, note that because $N^{\text{Pay}} = \mathbb{P}r(V^{\text{Pay}} > V^{\text{Delinquent}} + \varepsilon_i)$, it follows that

$$\frac{\partial N^{\text{Pay}}}{\partial V^{\text{Pay}}} = -\frac{\partial N^{\text{Pay}}}{\partial V^{\text{Delinquent}}}.$$

Let Δt be a tax increase such that the combined effect of the tax and the marginal enforcement action leaves N^{Pay} unaffected. We then have:

$$
\begin{aligned}
0 &= dN^{\text{Pay}} = \frac{\partial N^{\text{Pay}}}{\partial V^{\text{Pay}}}\frac{\partial V^{\text{Pay}}}{\partial t}\Delta t + \frac{\partial N^{\text{Pay}}}{\partial V^{\text{Delinquent}}}\frac{\partial V^{\text{Delinquent}}}{\partial \alpha} \\
&= \underbrace{\frac{\partial N^{\text{Pay}}}{\partial V^{\text{Pay}}}}_{>0}\underbrace{\left(\frac{\partial V^{\text{Pay}}}{\partial t}\Delta t - \frac{\partial V^{\text{Delinquent}}}{\partial \alpha}\right)}_{=0}.
\end{aligned}
$$

Because the first term is non-zero, the second term being equal to zero implies that:

$$
\Delta t = \frac{\frac{\partial V^{\text{Delinquent}}}{\partial \alpha}}{\frac{\partial V^{\text{Pay}}}{\partial t}} = \frac{-m'(\alpha)}{-u'(c)H} = \frac{m'(\alpha)}{u'(c)H}.
$$

Because this joint tax and enforcement change is such that N^{Pay} is unaffected, we can write:

$$
\begin{aligned}
0 &= \Delta t\frac{\partial N^{\text{Pay}}}{\partial t} + \frac{\partial N^{\text{Pay}}}{\partial \alpha} \\
0 &= \frac{m'(\alpha)}{u'(c)H}\frac{\varepsilon_t^{\text{Pay}}}{t} + \frac{\varepsilon_\alpha^{\text{Pay}}}{\alpha} \\
m'(\alpha) &= u'(c)\left(\frac{Ht}{\alpha}\right)\left(\frac{\varepsilon_\alpha^{\text{Pay}}}{-\varepsilon_t^{\text{Pay}}}\right),
\end{aligned}
$$

where the second line substitutes for Δt and transforms the expression into terms of elasticities and the third line solves for $m'(\alpha)$. This expression shows that we can measure $m'(\alpha)$ as a multiple of marginal utility that depends on the relative effects of taxes and enforcement on compliance.

This expression also shows that the welfare cost of enforcement is increasing in $\varepsilon_\alpha^{\text{Pay}}$. This makes sense. If a given enforcement action has a large effect on payment, the equivalent tax increase would have to be greater to result in the same effect on compliance. However, while a larger value of $\varepsilon_\alpha^{\text{Pay}}$ implies that enforcement raises more revenue, it also implies that enforcement is relatively less attractive from a welfare perspective.

It is worth pointing out that our approach to measuring $m'(\alpha)$ relies on the assumption that our enforcement treatment takes the form of a psychological tax (e.g., as in Caplin, 2003; Loewenstein and O'Donoghue, 2006; Glaeser, 2006). In addition to the case of the psychological tax, Allcott and Kessler (2019) consider that nudges may provide information that impacts decision-making of households, as well as the case where nudges provide a moral subsidy. Since taxpayers received six reminder bills prior to receiving the enforcement message, it is unlikely that the enforcement message provided any new information about the tax liability. Alternatively, the message may be interpreted as providing information that may impact beliefs about $\tilde{z}$. However, since stronger enforcement actions apply only to a small subset of delinquent taxpayers, it is unlikely that the treatment improved beliefs about $\tilde{z}$. In fact, even if taxpayers updated their beliefs, they might have updated to a $\tilde{z}$ that was higher than its actual value. If this were the case, household responses to the enforcement letter would carry further welfare losses arising from acting sub-optimally based on these erroneous beliefs. Accounting for this possibility would further reduce the desirability of enforcement messages as a policy tool. Finally, while other types of nudges may activate a "warm glow" effect from paying taxes, even our public goods treatment was primarily an enforcement message. For this reason, we do not believe that the enforcement actions that we study are likely to generate a moral subsidy.

A.3 Comparing $MVPF_\alpha$ and $MVPF_t$

We now consider the welfare effect of increasing taxes and reducing enforcement while keeping government expenditure constant. The welfare impact of this policy experiment is given by:

$$dW = W_t + W_\alpha \cdot \left.\frac{d\alpha}{dt}\right|_{dg=0}.$$

From the government budget constraint, we have that:

$$\left.\frac{d\alpha}{dt}\right|_{dg=0} = -\frac{HN(1+(1-\tilde{z})\varepsilon_t^{\text{Pay}})}{\frac{tHN}{\alpha}(1-\tilde{z})\varepsilon_\alpha^{\text{Pay}} - 1}.$$

From Equations 1 and 3, we can write:

$$\begin{aligned} W_t &= MVPF_t \times u'(c)\left(HN(1+(1-\tilde{z})\varepsilon_t^{\text{Pay}})\right) \quad \text{and} \\ W_\alpha &= MVPF_\alpha \times \omega m'(\alpha)\left(\frac{tHN}{\alpha}(1-\tilde{z})\varepsilon_\alpha^{\text{Pay}} - 1\right). \end{aligned}$$

Combining these four expressions, we then have that:

$$dW = u'(c)\left(HN(1+(1-\tilde{z})\varepsilon_t^{\text{Pay}})\right) \times \left[MVPF_t - MVPF_\alpha \times \frac{\omega m'(\alpha)}{u'(c)}\right].$$

Recalling from above that $\frac{m'(\alpha)}{u'(c)} = \left(\frac{Ht}{\alpha}\right)\left(\frac{\varepsilon_\alpha^{\text{Pay}}}{-\varepsilon_t^{\text{Pay}}}\right)$, we have:

$$dW \propto MVPF_t - MVPF_\alpha \times \omega\left(\frac{Ht}{\alpha}\right)\left(\frac{\varepsilon_\alpha^{\text{Pay}}}{-\varepsilon_t^{\text{Pay}}}\right).$$

We use this expression to plot Panel C of Figure 9.

A.4 Optimal Enforcement

Setting $MVPF_\alpha = 0$ yields the following condition for the optimal enforcement level α :

$$\varepsilon_\alpha^{\text{Pay}} = \frac{(1-N^{\text{Pay}})\frac{\alpha\omega m'(\alpha)}{v'(g)} + \alpha}{(1-\tilde{z})N^{\text{Pay}}Ht}.$$

At the optimum α, the government equates the enforcement elasticity $\varepsilon_\alpha^{\text{Pay}}$ to the ratio of enforcement costs to tax revenue. Importantly, this ratio accounts for both public enforcement expenditures α and the private disutility from enforcement $m(\alpha)$ that is incurred by the $1-N^{\text{Pay}}$ delinquent taxpayers.[64]

[64] Keen and Slemrod (2017) call this term the adjusted marginal cost-revenue ratio. The term $\alpha m'(\alpha)$ can be viewed as a first-order approximation of $m(\alpha)$. Private enforcement costs are discounted by the value of public goods

From a full fairness perspective that ignores the welfare costs to non-compliant taxpayers, i.e., $\omega = 0$, this equation simplifies to a cost-benefit analysis that compares the effect of enforcement on revenue to the administrative cost of enforcement.

For a given enforcement level α, this expression defines the equilibrium compliance rate that equates the marginal costs and benefits of enforcement:

$$N_{\alpha}^{\text{Pay}} = \frac{1 + \frac{\omega m'(\alpha)}{v'(g)}}{(1-\tilde{z})\frac{\varepsilon_{\alpha}^{\text{Pay}} Ht}{\alpha} + \frac{\omega m'(\alpha)}{v'(g)}}. \tag{A.1}$$

This expression is useful from a policy perspective, as it tells the government whether the marginal benefit exceeds the marginal cost of enforcement. This is the case when N_{α}^{Pay}—which combines the administrative costs, private costs, and effectiveness of enforcement—is smaller than the observed compliance rate N^{Pay}. In this situation, the government can raise welfare by increasing enforcement (Figure A.1 provides a graphical intuition of this result).

Figure A.2 plots the compliance rate that equates the marginal cost and benefit of enforcement. As we discuss in Section 1, the government can increase welfare by increasing enforcement when $N_{\alpha}^{\text{Pay}} < N^{\text{Pay}}$. Conversely, it would be preferable to reduce enforcement when $N_{\alpha}^{\text{Pay}} > N^{\text{Pay}}$. Regardless of the value of public goods, the observed compliance rate of $N^{\text{Pay}} = 60\%$ lies below N_{α}^{Pay} in both our baseline parameterization and the case with higher $\varepsilon_{\alpha}^{\text{Pay}}$. This result implies that the government is over-relying on enforcement as a means to raise tax revenue. While the equilibrium compliance rate is not very sensitive to the enforcement elasticity, a larger tax elasticity shifts the curve N_{α}^{Pay} down. The green dot-dashed line shows that, under the assumptions that $\varepsilon_{t}^{\text{Pay}} = -0.46$ and $\frac{v'(g)}{u'(c)} = 1.75$, current compliance rates are close to optimal. Current enforcement levels could then be rationalized by a government that assumed that tax hikes have small effects on tax revenue (i.e., that $\varepsilon_{t}^{\text{Pay}}$ is large and negative). This insight showcases the importance of studying the welfare effects of enforcement by taking into account the efficiency of alternative policies available to the government.

A.5 Government Provision of Liquidity

Assume now that the government allows households that pay property taxes to borrow up to the amount of the property taxes at interest rate r.[65] We can interpret this rate of return as incorporating a risk adjustment for the possibility that households do not pay back the loan. Because the government can eventually seize the asset, this collateral implies that this adjustment is low.

The provision of liquidity to constrained taxpayers lowers the welfare cost of taxation since consumption would be less affected. Specifically, the change in consumption for constrained households is now $\eta_{t,l}^{c} \times t$, where it is plausible to assume that $\eta_{t,l}^{c} \approx \eta_{t,s}^{c} < \eta_{t,0}^{c}$. Therefore, when the government provides liquidity, $MVPF_t$ is greater, since the effect on consumption is smaller. The provision of liquidity to constrained taxpayers also means that enforcement becomes relatively less desirable since liquidity increases the value of $MVPF_t$.

The assumption that the government charges a risk-adjusted interest rate implies that the government's budget constraint is not affected by providing liquidity. Departing from this assumption, it is also possible to study the optimal provision of liquidity. As in Andreoni (1992), the government

$v'(g)$ since an increase in α requires lowering g.

[65] We assume that only households that pay property taxes and have no savings may decide to take out a loan.

may have incentives to act as a "loan shark." The government's budget constraint is now:

$$g + a(\alpha) = tHN^{\text{pay}} + zH(1 - N^{\text{pay}}) + (\rho - r) \times \pi_l^{\text{Pay}} N^{\text{pay}} tH,$$

where the last term is the revenue from charging interest ρ on the taxes of the share of taxpayers π_l^{Pay} who obtain a loan from the government.

This implies $\frac{dV^{\text{Pay}}}{d\rho} = -u'(c^{\text{Pay}})H\pi_l^{\text{Pay}}$ and $\frac{dV_l^{\text{Delinquent}}}{d\rho} = 0$. The effect of increasing ρ on welfare is then:

$$-N^{\text{Pay}}\pi_l^{\text{Pay}}u'(c_l^{\text{Pay}})H + v'(g) \times \left\{(t - z)H\frac{\partial N^{\text{Pay}}}{\partial \rho} + \pi_l^{\text{Pay}}N^{\text{pay}}tH + (\rho - r) \times tH\left[\pi_l^{\text{Pay}}\frac{\partial N^{\text{Pay}}}{\partial \rho} + \frac{\partial \pi_l^{\text{Pay}}}{\partial \rho}N^{\text{pay}}\right]\right\}.$$

The MVPF for ρ is then:

$$MVPF_\rho = \frac{v'(g)}{u'(\bar{c})} - \frac{\rho(1 - \gamma\pi_{l,\rho}^{\text{Pay}}\Delta c_{l,\rho}^{\text{Pay}})}{\frac{t-z}{\pi_l^{\text{Pay}}}\varepsilon_\rho^{\text{Pay}} + t\rho + (\rho - r) \times t[\varepsilon_\rho^{\text{Pay}} + \varepsilon_\rho^{\pi_l}]}.$$

The optimal value of ρ solves this expression when set equal to zero. From this expression, it follows that the government might set $\rho > r$ and therefore act as a loan shark if the value of providing public goods through loans exceeds the welfare cost of raising revenue in this way.

Figure A.1: N_α Equates the Marginal Cost (MC) and Benefit (MB) From Enforcement

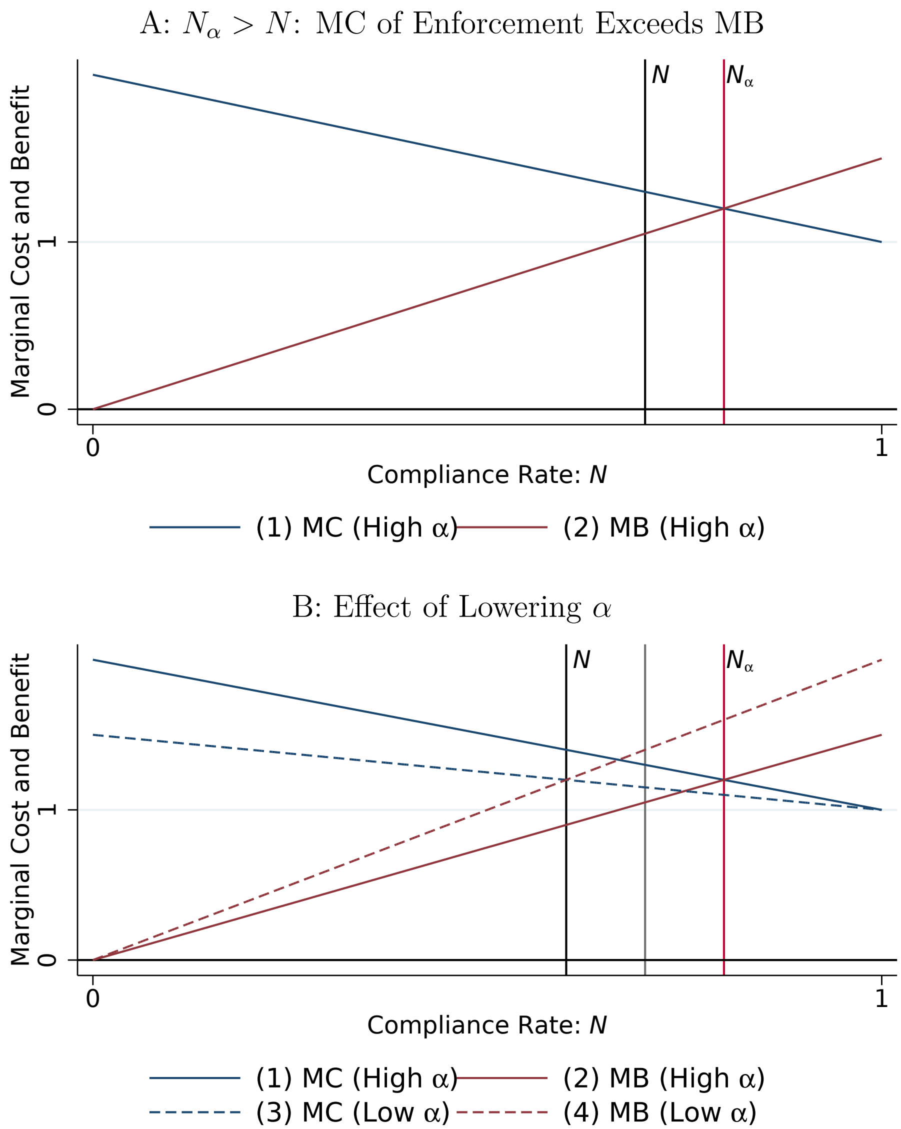

Notes: This figure plots the marginal cost and benefit of enforcement as well as the implied N_α. To simplify exposition, we present the case when $\omega = 1$. The marginal cost of enforcement (plotted in blue) is given by $(1-N)\frac{m'(\alpha)}{v'(g)} + 1$, which includes the private and public costs of enforcement. The MC of enforcement is a decreasing function of N and equals 1 when $N = 1$ (this is because we assume $a(\alpha) = \alpha$, so that $\frac{\partial a(\alpha)}{\partial \alpha} = 1$). The marginal benefit of enforcement (plotted in red) is given by $(1-\tilde{z})\frac{Ht}{\alpha}\varepsilon_\alpha^{\text{Pay}} N$, which starts at the origin and increases with N. In Panel A, the MC equals the MB at N_α. This figure assumes that the observed compliance rate $N < N_\alpha$. At the observed compliance rate N, the MC of enforcement exceeds the MB, such that welfare would be increased by lowering α. Panel B shows the effect of lowering α. For lower values of α, the term $m'(\alpha)$ is smaller, resulting in a flatter MC curve (shown with the dashed line). For lower values of α, the elasticity $\varepsilon_\alpha^{\text{Pay}}$ is higher (i.e., initial enforcement efforts are more effective), resulting in a steeper MB curve (shown with the dashed line). In Panel B, MB equals MC at the new N, which is lower than N in Panel A.

Figure A.2: Equilibrium Enforcement

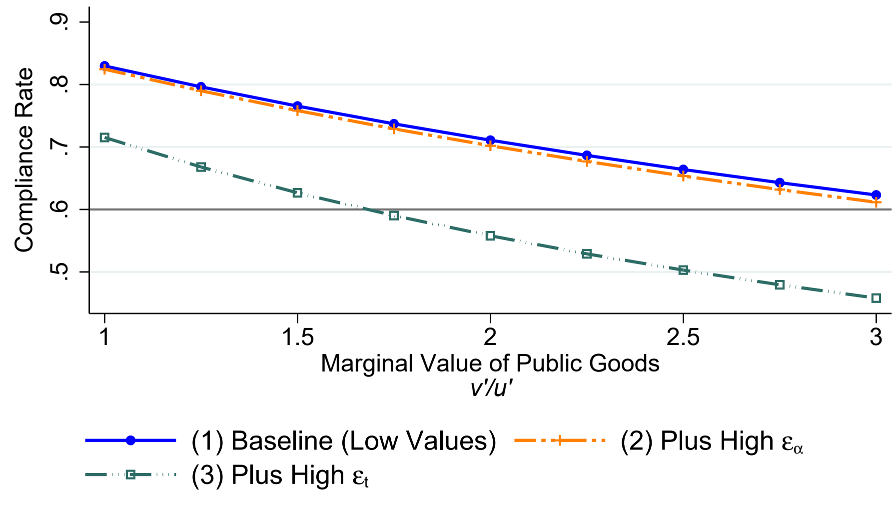

Notes: This figure displays the equilibrium compliance rate discussed in Appendix A.4. The horizontal line represents the observed average compliance rate between 2008 and 2012.

Figure A.3: Role of ω in Relative Desirability of Taxes over Enforcement

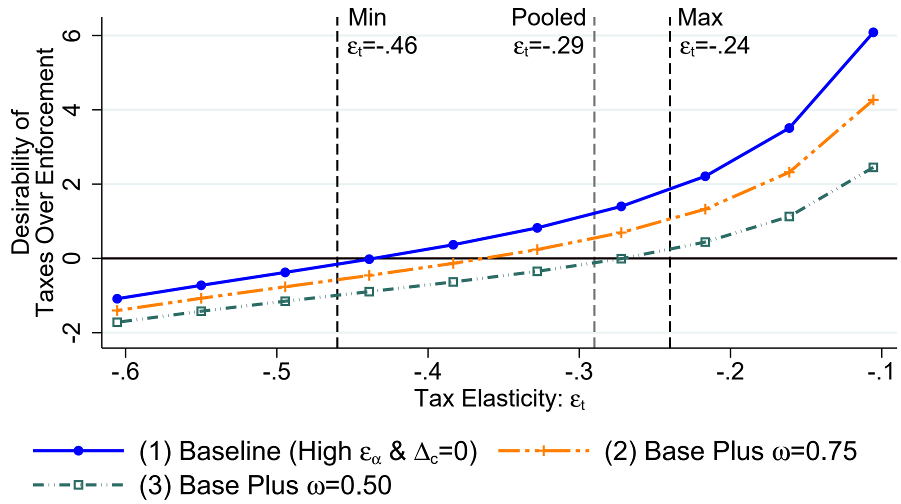

Notes: This figure plots the welfare impact of a balanced-budget policy that increases taxes and reduces enforcement. The figure varies the parameter ω, which under-weights the private cost of enforcement.

Figure A.4: Role of Δ_c in Relative Desirability of Taxes over Enforcement

A: Low Gamma

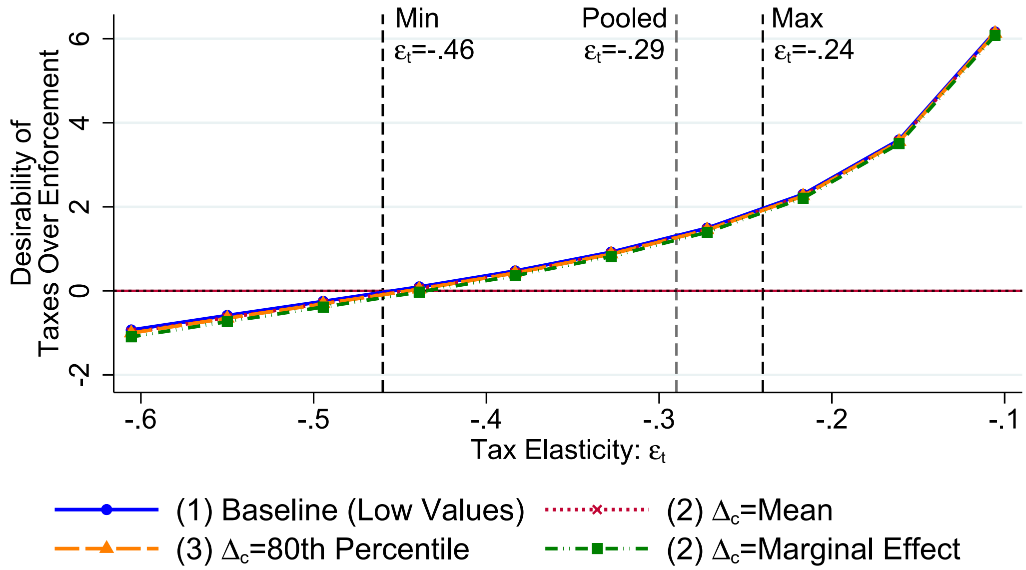

B: High Gamma

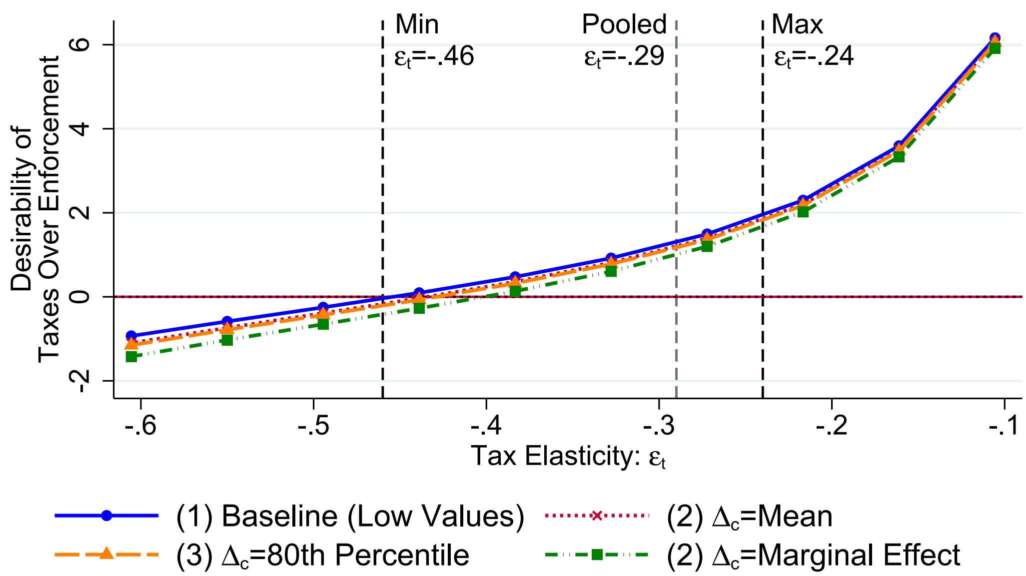

Notes: This figure plots the welfare impact of a balanced-budget policy that increases taxes and reduces enforcement. The figure varies the parameter ΔC, the consumption drop that liquidity-constrained households (80% of all households) experience when property taxes are increased. The baseline scenario assumes $\Delta C = 0$. We then use $\Delta C =$ mean of (property tax payment/income)=3.2% (scenario 2), or $\Delta C =$ 80th percentile of property tax/income=4.2% (scenario 3), or $\Delta C =$ the marginal effect of tax payment on consumption for liquidity-constrained households from our IV estimations in Appendix I, i.e. 9.3% (scenario 4). In each scenario, we multiply the consumption drop with the share of liquidity-constrained households. This figure is discussed in Section 7.1.

B Context Appendix

Table B.1: Consumer Debt in Mexico City

Variable	Mean (1)
Credit take-up	
Informal	.345
Formal	.302
Both	.084
None	.437
Informal borrower shares by type of lender (not exclusive)	
Pawnshop	.135
Friends	.32
Family	.729
Other	.013
Reasons for informality	
Voluntary	.578
Non-eligibility (lack of access)	.288
Initial costs	.125
Other	.009
Formal borrower shares by credit source (not exclusive)	
Credit card	.834
Bank loan	.113
Mortgage	.189
Car/Other	.079
Number of mortgages (liquidity constraints)	
One	1
Two or more	0
Use of formal credit	
Paying a bill	.161
Other	.839
Observations	877

Notes: The table examines consumer debt in Mexico City in 2018, using data from the National Financial Inclusion Survey (*Encuesta Nacional de Inclusión Financiera*, ENIF). The contents of this table are discussed in the Introduction.

Table B.2: Tax Schedule, 2009

Band	Cadastral Value Lower Limit (MXN)	Cadastral Value Upper Limit (MXN)	Lump-Sum Liability (MXN)	Tax Rate on Excess from Lower Limit (percent)	Percent Abatement on Liability
	(1)	(2)	(3)	(4)	(5)
A	0.11	162,740.82	32	0	0
B	162,740.83	325,481.16	37	0	0
C	325,481.17	650,963.56	45	0	0
D	650,963.57	976,444.70	55	0	0
E	976,444.71	1,301,927.10	737.28	0.09542	65
F	1,301,927.11	1,627,408.26	1,047.86	0.11091	45
G	1,627,408.27	1,952,889.39	1,408.85	0.11461	30
H	1,952,889.40	2,278,371.81	1,781.88	0.12522	20
I	2,278,371.82	2,603,852.96	2,189.45	0.13097	15
J	2,603,852.97	2,929,335.38	2,615.73	0.13478	10
K	2,929,335.39	3,254,816.51	3,054.42	0.13892	0
L	3,254,816.52	3,580,297.67	3,506.58	0.1427	0
M	3,580,297.68	3,906,090.04	3,971.04	0.15075	0
N	3,906,090.05	11,718,268.85	4,462.17	0.16278	0
O	11,718,268.86	24,663,843.29	17,178.84	0.16286	0
P	24,663,843.30		38,262.00	0.16902	0

Notes: This table presents an example of the annual tax schedule discussed in Section 2.2, focusing on the year 2009.

Table B.3: Abatements on Gross Tax Liability

Band	2008	2009	2010	2011	2012	2013
	(1)	(2)	(3)	(4)	(5)	(6)
G	30	30	30	30	20	20
H	20	20	20	0	0	0
I	15	15	0	0	0	0
J	10	0	0	0	0	0

Notes: This table displays the abatement rates discussed in Section 2.2.

Table B.4: Early-Bird Discounts and Payment Deadlines

Year	Super Early Bird		Early Bird		Reference Rates		
	Deadline	Discount	Deadline	Discount	Central Bank	Treasury Bonds	Mortgages
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
2008	Jan 31	7%	Feb 28	4%	7.5%	7.42%	12.22%
2009	Jan 31	8%	Feb 28	4%	8.25%	7.59%	12.78%
2010	Jan 31	5%	Feb 28	0%	4.5%	4.49%	12.79%
2011	Jan 10	7%	Jan 31	3%	4.5%	4.14%	12.22%
2012	Jan 17	7%	Jan 31	4%	4.5%	4.27%	12.53%
2013	Jan 31	7%	Feb 28	6%	4.5%	4.15%	12.13%

Notes: This table displays the early-bird discount schedules discussed in Section 2.3 and Appendix H. Discounts are applied to the annual tax liability. All interest rates are annualized.

Table B.5: Home Ownership in Mexico

Panel A: Renting vs Owning

	Homeowners (1)	Renters (2)	Others (3)	P value (4)
Share of the population	62.1	14.2	23.6	
Average Monthly Labor Income (MXN)	3966.1	5394.7	3785.7	0.017
	(263.893)	(473.894)	(363.832)	
Number of Rooms	2.2	1.7	1.8	0.000
	(.048)	(.075)	(.064)	
Number of Household Members	4	3.5	3.7	0.127
	(.092)	(.166)	(.141)	
Age Head of Household	55.5	37	44.8	0.000
	(.748)	(1.466)	(1.348)	

Panel B: Homeowner Characteristics

Home Financing (1)	Current Mortgage (2)	Pays Property Tax (3)	Property Tax Payment Amount (4)	Own a Second Home (5)
14.3	9.9	56.6	300.3	5.4
(1.627)	(1.392)	(2.306)	(20.009)	(1.051)

Panel C: Home Financing

	Gender		Income Quintiles				
Total (1)	Male (2)	Female (3)	Poorest (1)	2nd (2)	3rd (3)	4th (4)	Richest (5)
4.3	4.8	3.9	0.0	2.3	8.0	5.0	6.1
(0.643)	(1.080)	(0.785)	(0.00)	(1.171)	(1.877)	(1.558)	(1.510)

Notes: The table examines home ownership in Mexico, as discussed in Section 2.3. Panels A and B displays summary statistics of Mexican households by ownership status, based on the 2014 ENVI (*Encuesta Nacional de Vivienda*) from the National Institute of Statistics. In Panel A, the home status "Others" includes loaned properties and properties under litigation. The p-values in Panel A evaluate the differences between homeowners and renters. In Panel B, "Home Financing" indicates the share of owners who have received any kind of loan to finance their home purchase. Panel C displays the share of households with a mortgage in the country and its demographic correlates, based on data from the 2017 World Bank Findex database. Standard errors are in parentheses. The difference in the share of observations with a mortgage in Panels B and C is driven by differences in the sample. Panels A and B are for Mexico City, while Panel C is for the whole country.

Figure B.1: Property Taxes in Mexico City

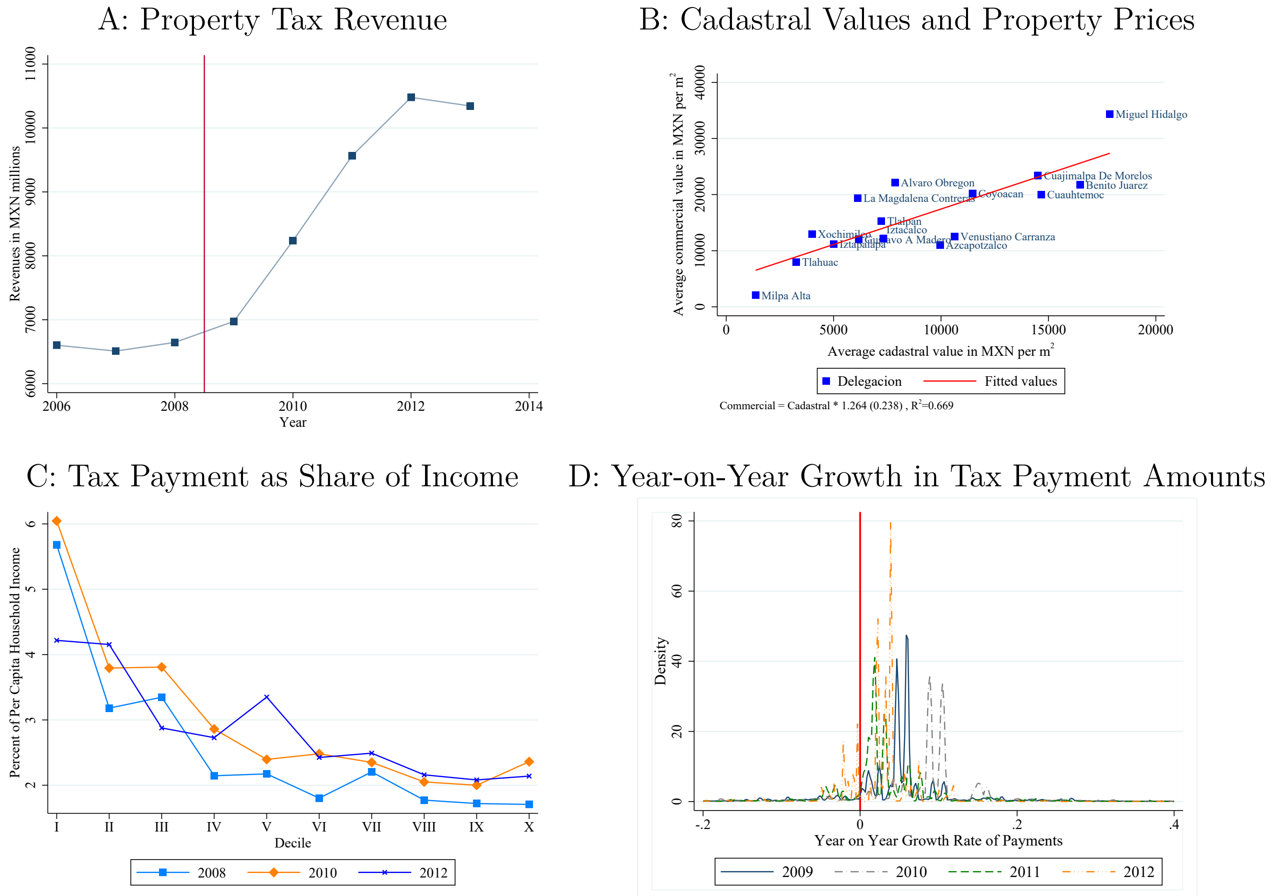

Notes: These figures show key facts about property taxes in Mexico City and are discussed in the Introduction and Sections 2.1 and 2.2. Panel A shows the total property tax revenue by year for Mexico City in nominal terms from government records. Panel B shows the correlation between average cadastral and commercial property values at the *delegación* level. Average commercial prices are obtained from propiedades.com, one of the largest real estate websites in Mexico. Prices were retrieved on the June 4, 2020, and are discounted for inflation using INEGI's inflation calculator. Cadastral values are from the administrative data. Panel C plots the property tax payment reported in the ENIGH household survey (*Encuesta Nacional de Ingresos y Gastos de los Hogares*) as a share of per capita household income, conditional on property tax payment being non-zero. Each line corresponds to a different survey round. Panel D displays the year-on-year growth rate of property tax payments, $\frac{Pay_t - Pay_{t-1}}{Pay_{t-1}}$. The sample is restricted to taxpayers who made a payment in both year t and year $t-1$. The figure shows that there is no anchoring of tax payments at the previous year's liability or tax payment amount, as liabilities are inflation-adjusted each year.

Figure B.2: Relevance and External Validity of the Mexico (City) Context

A: Tax Administration Capacity around the World

A1: Comprehensive Audit Rate

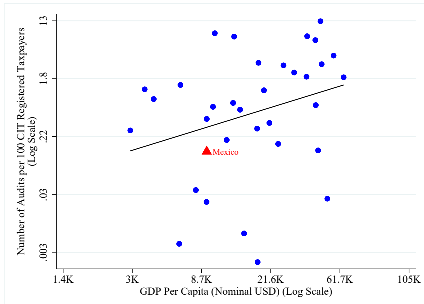

A2: CIT Audit Rate

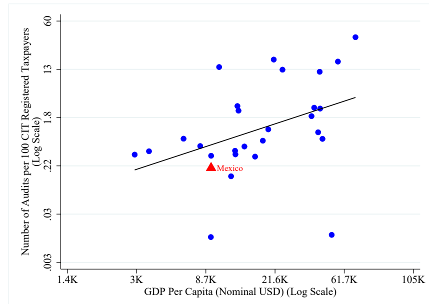

B: Household Liquidity Constraints around the World

B1: Share of Adults with a Bank Account

B2: Share of Adults with a Credit Card

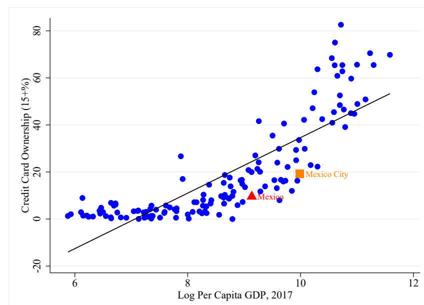

B3: Share of Adults with a Formal Loan

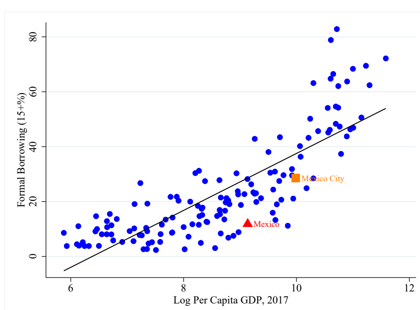

B4: Share of Adults Able to Cover an Emergency Expense from Savings

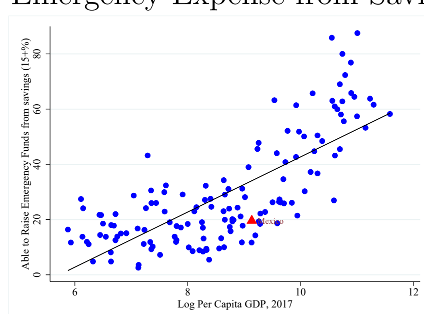

Notes: As discussed in the introduction, this figure shows the correlation between measures of government tax administration capacity and household liquidity constraints with GDP per capita and the levels of these indicators for Mexico, as well as for Mexico City where available. The data for Panels A1 and A2 are from the 2016 Revenue Administration Fiscal Information Tool (RA-FIT). The audit rate consists of the number of audits by each type conducted by the tax authority divided by the number of CIT-registered taxpayers. The data for Panels B1–B4 are from the 2017 World Bank Findex database for all countries and from the 2018 National Financial Inclusion Survey for Mexico City. Panel B4 displays the share of adults who can cover an emergency (an unexpected expense approximately equivalent to 500 USD) from personal savings (formal or informal). This statistic is not available for Mexico City only.

Figure B.3: Property Tax Bill


NOMBRE Y FIRMA DEL CONTRIBUYENTE O REPRESENTANTE LEGAL
CDMX
GOBIERNO DEL DISTRITO FEDERAL
Secretaría de Finanzas
Tesorería

FECHA DE CORTE
13/Diciembre/2015

M²de suelo (terreno)	M²de suelo construcción	Uso - Tipo	IMPUESTO REAL
301.00	221.00	H-02	$945.06
Valor unitario por M²	Valor unitario por M²	Clase	SUBSIDIO OTORGADO POR EL GDF
$1,196.45	$4,366.87	4	$472.06
Valor del suelo	Valor de la construcción	Valor catastral	TOTAL A PAGAR PRIMER BIMESTRE
$360,131.45	$617,650.09	$977,781.54	$473.00

PAGO ANUAL ANTICIPADO (DEL 1 AL 31 DE ENERO)

Vence	Línea de Captura	Importe
31/ENE/16	8007635376000Y57QB2H	$2,608.00

PAGO ANUAL ANTICIPADO (DEL 1 AL 29 DE FEBRERO)

Vence	Línea de Captura	Importe
29/FEB/16	80076353760000E7U6A8	$2,693.00

PAGO 1er. BIMESTRE

Vence	Línea de Captura	Importe
29/FEB/16	80076353760000E5B7AM	$473.00

Propuesta de Declaración de Valor Catastral y pago del Impuesto Predial que se emite con fundamento en los artículos 15, 126, 127, 129, 130 y 131 del Código Fiscal del Distrito Federal.

ESTADO DE CUENTA
076353760000-6
MANZANA
REGIÓN

Notes: This figure displays a typical property tax bill sent to homeowners, as discussed in Section 2.3.

Figure B.4: Distribution of Treated Properties across Neighborhoods in Mexico City

A: Neighborhood Density of Treated Properties

B: Distribution of Treated Properties

C: Count of Treated Properties

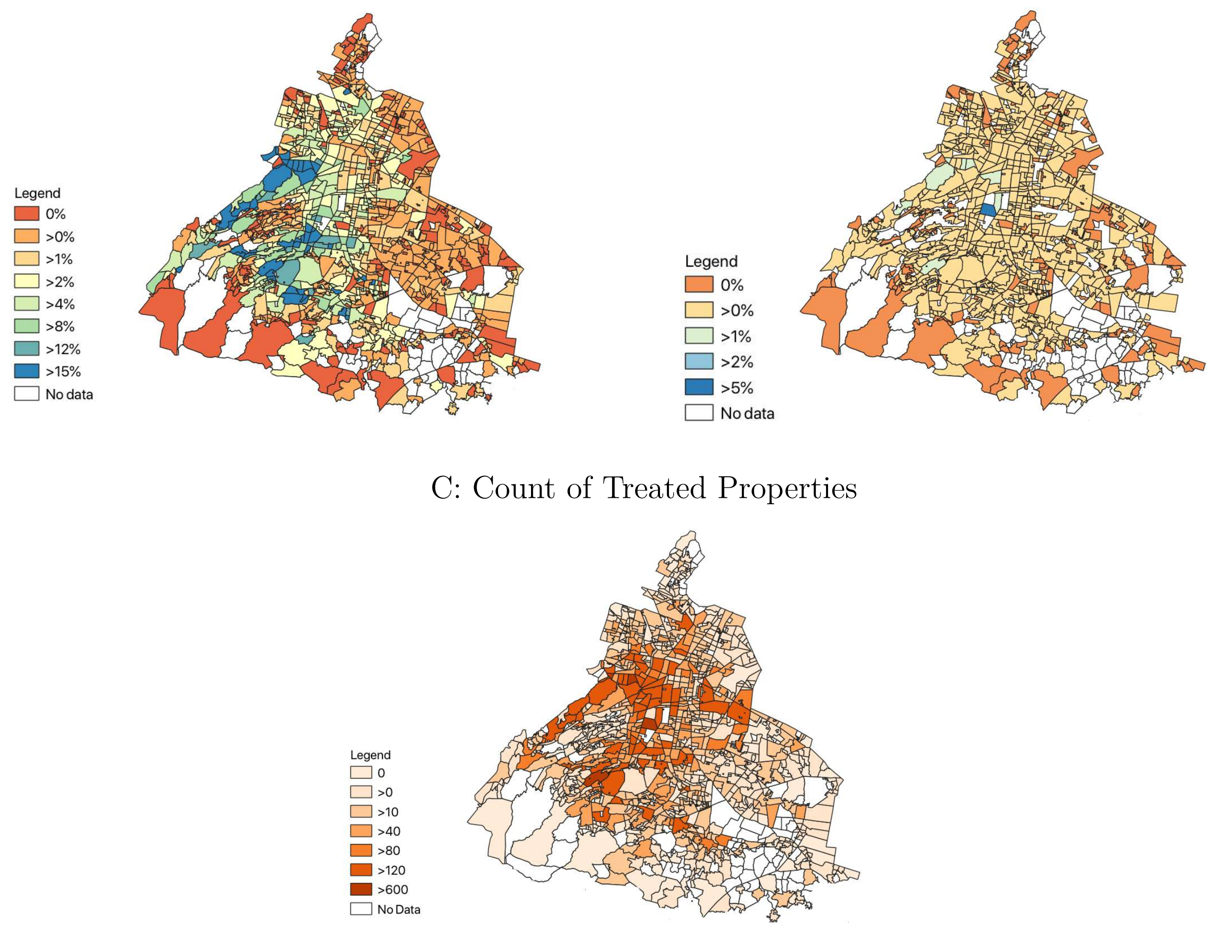

Notes: Panel A shows the density of treated properties within neighborhoods in Mexico City. The density is the number of treated properties in a given neighborhood divided by the total number of properties in the neighborhood. Panel B shows the distribution of treated properties across neighborhoods. For each neighborhood, we show the ratio between the number of treated properties in this neighborhood and the total number of treated properties in the city. Panel C shows the number of treated properties in each neighborhood. These maps are discussed in Section 2.2.

Figure B.5: Tax Payment Timing Around Reminder (Bill) Receipt

A. Compliant Taxpayers

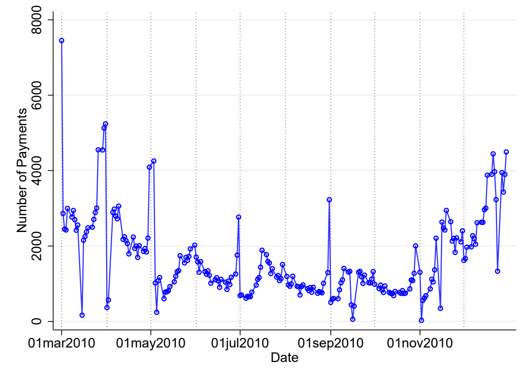

B. Delinquent Taxpayers

B1. Current Balance

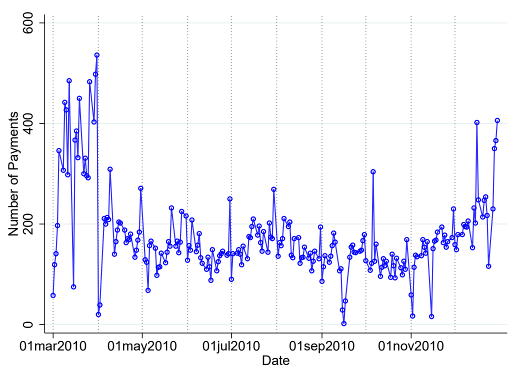

B2. Outstanding Balance

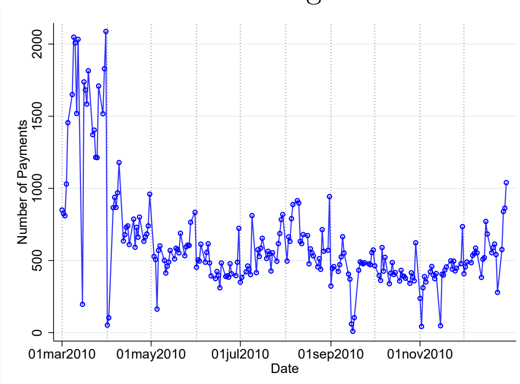

Notes: This figure examines the timing of tax payments throughout the year. For taxpayers who did not pay their annual liability in full by the early-bird deadline, the deadlines to pay the tax liability in installments are February 28, April 30, June 30, August 31, October 31 and December 31 of each year (with minor year-specific divergences). The bimonthly tax bills reminding taxpayers of the upcoming deadline are sent around 5-15 days before the deadline. Featuring data for 2010, this figures show that payments among compliant taxpayers are bunched at the end of each bimonthly period, after bill (reminder) receipt and before the deadline. For non-compliant taxpayers who have outstanding liabilities to pay, this is not the case, suggesting the bimonthly bills do not play a reminder role for these taxpayers.

C Data Appendix

Table C.1: Summary Statistics
Panel A: Property Characteristics

	2008 (1)	2009 (2)	2010 (3)	2011 (4)	2012 (5)
Property Value (MXN)	585,320	617,487	613,493	609,478	605,346
	(1,121,680)	(1,185,320)	(1,180,471)	(1,174,999)	(1,169,283)
Yearly Liability (MXN)	1,457	1,540	1,630	1,704	1,788
	(10,097)	(10,671)	(11,214)	(11,607)	(11,985)
Mean Tax Rate × 100	.1112	.1114	.1198	.1259	.1323
	(.1243)	(.1245)	(.1349)	(.1427)	(.1532)

Panel B: Payment Characteristics

	2008	2009	2010	2011	2012
Payment (current MXN)	680	698	729	709	686
	(4078)	(4126)	(4439)	(4703)	(4563)
Compliance share	.628	.61	.591	.538	.488
	(.931)	(.711)	(.776)	(.62)	(1.149)
Payment type					
Zero payment	.275	.276	.307	.366	.432
	(.446)	(.447)	(.461)	(.482)	(.495)
Partial payment	.172	.174	.174	.136	.097
	(.377)	(.379)	(.379)	(.343)	(.296)
Early Full Payment	.413	.432	.416	.455	.461
	(.492)	(.495)	(.493)	(.498)	(.498)
Non-Early Full Payment	.141	.118	.103	.043	.01
	(.348)	(.323)	(.304)	(.204)	(.1)

Panel C: Penalties and Fees

	2008		2009	
	Mean (1)	Median (2)	Mean (3)	Median (4)
Late Payment Dummy (before due date 2)	.086		.085	
Inflation-Adjusted Liability (dummy)	.065		.053	
Inflation-Adjusted/Original Liability	1.299	1.222	1.164	1.193
Late Payment (after due date 2, within 2 years)	.065		.053	
Penalty Dummy	.005		0	
Surcharge Dummy	.065		.053	
Seizure Dummy	.001		0	
Penalty/Liability	.838	1	.226	.101
Surcharge/Liability	.183	.145	.154	.147
Seizure/Liability	.019	0	0	0
Total/Liability	1.299	1.222	1.164	1.193
Delinquent Taxpayer Dummy	.247		.252	

Notes: Panel A reports summary statistics for 1,420,259 properties with an average land area of 126 (SE=381) square meters and an average construction area of 126 (SE=161) square meters. Panel B reports summary statistics on payments. Panel C reports summary statistics on penalties and fees from the administrative tax data discussed in Section 3. The sample in panel C is slightly different from the sample in panel B as the penalties and fees are unavailable for a small subset of taxpayers. However, the samples are very similar as the panels show that the share of delinquent taxpayers (making zero payment) is around 25% in both samples. Note that Figure 8 in the main paper shows compliance (propensity to pay and payment modality) by property value band.

Figure C.1: Distribution of Property Characteristics by Value Band

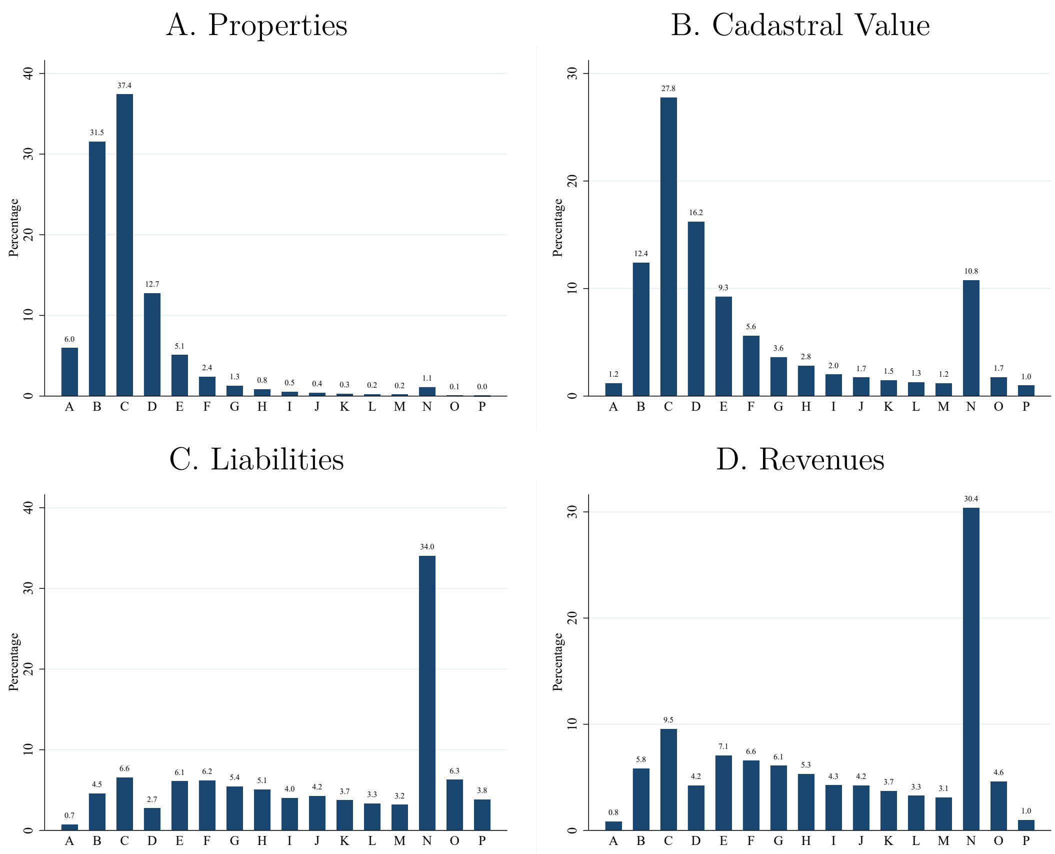

Notes: This figure shows the distribution of property characteristics by cadastral value band, as discussed in Section 3, for the balanced panel of residential properties in Mexico City whose cadastral value did not change between 2009 and 2012.

Figure C.2: Outstanding Property Tax Debt

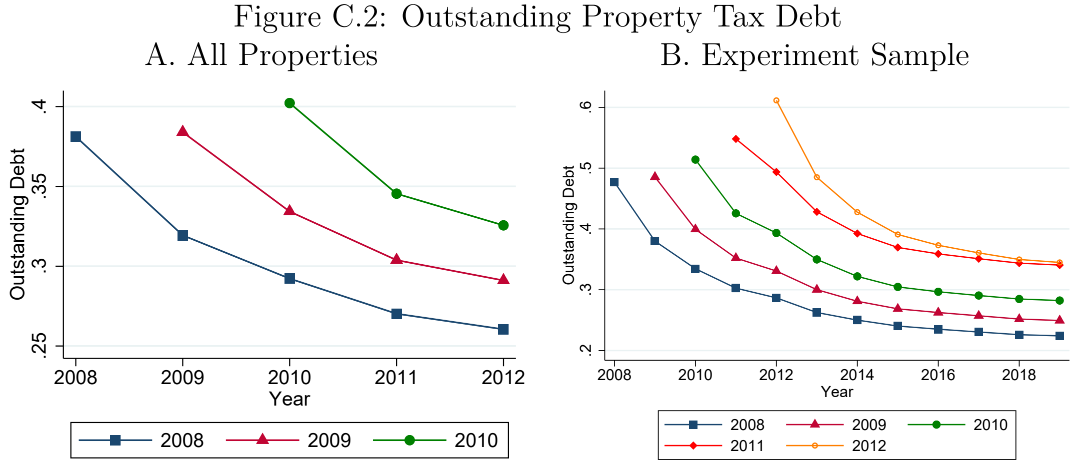

Notes: This figure shows how the share of each year's unpaid tax liability evolves over time. This is referenced in Section 1. Panel A includes all taxpayers with outstanding tax debt. Panel B restricts the sample to taxpayers targeted in the enforcement intervention.

D Regression Discontinuity Appendix

Table D.1: Identification Check for Regression Discontinuity Estimation

	Band I (2010)	Band H (2011)	Band G (2012)
	(1)	(2)	(3)
P-value ($H_0 : f^+_{\text{Value}} = f^-_{\text{Value}}$)	.444	.828	.752
q	138	190	250
N (left)	1647	2237	3076
N (right)	1059	2109	3501
Total N	2706	4346	6577
Effective N (left)	74	97	122
Effective N (right)	64	93	128

Notes: This table reports results from the RD validity test proposed by Bugni and Canay (2020), as discussed in Section 4.1. This test examines the continuity of the running variable at the cut-off, an implication of the assumption of no manipulation. In particular, the fraction of units under treatment and control should be similar on both sides. The test statistic exploits the fact that under the null, the number of treated units out of the q observations closest to the cut-off is approximately distributed as a binomial with sample size q and probability $\frac{1}{2}$. The paper proposes a data-dependent rule for q, the number of "effective" observations near the cut-off.

Table D.2: Robustness of Regression Discontinuity Estimation – Using Local Linear Regressions with Optimal Bandwidth

	Mean Tax Rate (basis points)	Payment Amount (MXN thousands)	Payment in Full (percentage points)	Compliance Share × 100
	(1)	(2)	(3)	(4)
	I. Estimates for the 2009-2010 treatment			
T	8.923 ***	.61 *	-8.426 **	-2.471
	(.105)	(.323)	(3.606)	(2.641)
Properties	17864	17864	17864	17864
Mean at Baseline (treated band)	50.112	5.836	36.626	47.881
Implied Elasticity		.587	-1.292	-.29
		(.311)	(.553)	(.31)
	II. Estimates for the 2012-2011 treatment			
T	12.109 ***	.667 ***	-5.8 **	-3.019
	(.033)	(.16)	(2.573)	(2.032)
Properties	28094	28094	28094	28094
Mean at Baseline (treated band)	47.461	4.734	35.072	47.478
Implied Elasticity		.552	-.648	-.249
		(.133)	(.288)	(.168)
P-value ($H_0 : \epsilon_{2011} = \epsilon_{2010}$)		.918	.302	.908
	III. Estimates for the 2011-2012 treatment			
T	17.958 ***	.644 ***	-9.615 ***	-3.833 *
	(.018)	(.134)	(1.567)	(2.051)
Properties	48838	48838	48838	48838
Mean at Baseline (treated band)	41.06	3.287	37.969	44.885
Implied Elasticity		.448	-.579	-.195
		(.093)	(.094)	(.104)
P-value ($H_0 : \epsilon_{2012} = \epsilon_{2011}$)		.52	.819	.785
P-value ($H_0 : \epsilon_{2012} = \epsilon_{2010}$)		.668	.204	.772

Notes: This table is similar to Table 2, Panel A, but uses local linear regressions with optimal bandwidth as in Calonico et al. 2014. The estimates are statistically indistinguishable from the results of the preferred specification.

Table D.3: Robustness of Regression Discontinuity Estimation – Differences-in-Discontinuities Estimates

	Mean Tax Rate (basis points)	Payment Amount (MXN thousands)	Payment in Full (percentage points)	Compliance Share × 100
	(1)	(2)	(3)	(4)
	I. Estimates for the 2009-2010 treatment			
β_1	8.23 ***	1.019 ***	-5.377 ***	1.521
	(.05)	(.183)	(1.695)	(1.399)
Properties	17864	17864	17864	17864
Years of Data	4	4	4	4
Adjusted R-Squared	.318	.005	.002	.001
Mean at Baseline (treated band)	50.112	5.836	36.626	47.881
Implied Elasticity		1.063	-.894	.193
		(.191)	(.282)	(.178)
P-value ($H_0 : \beta_0 = \beta_1$)	.000	.000	.000	.915
	II. Estimates for the 2010-2011 treatment			
β_1	9.71 ***	.445 ***	-2.568 **	-3.584 ***
	(.066)	(.12)	(1.248)	(1.081)
Properties	28094	28094	28094	28094
Years of Data	4	4	4	4
Adjusted R-Squared	.31	.005	.003	.002
Mean at Baseline (treated band)	47.461	4.734	35.072	47.478
Implied Elasticity		.459	-.358	-.369
		(.124)	(.174)	(.111)
P-value ($H_0 : \beta_0 = \beta_1$)	.000	.000	.009	.013
	III. Estimates for the 2011-2012 treatment			
β_1	20.017 ***	.503 ***	-11.621 ***	-7.285 ***
	(.038)	(.071)	(1)	(.932)
Properties	48838	48838	48838	48838
Years of Data	4	4	4	4
Adjusted R-Squared	.934	.006	.003	.002
Mean at Baseline (treated band)	41.06	3.287	37.969	44.885
Implied Elasticity		.314	-.628	-.333
		(.045)	(.054)	(.043)
P-value ($H_0 : \beta_0 = \beta_1$)	.000	.000	.000	.000

Notes: This table shows, as discussed in Section 4.1, the effect of the tax rate changes driven by the abatement removal, in excess of the effect of the smaller year-on-year tax rate changes, given by β_1 in the estimating equation $\Delta Y_{i,t} = \alpha_0 + \beta_0 T_i + f_0(\hat{V}_i) + g_0(\hat{V}_i)T_i + [\alpha_1 + \beta_1 T_i + f_1(\hat{V}_i) + g_1(\hat{V}_i)T_i]D_t + \epsilon_{i,t}$, where D_t is an indicator for the time period when the abatement is removed. The results of this estimation are very similar to our main estimates in Table 2, Panel A.

Figure D.1: Identification Tests for Regression Discontinuity Estimation

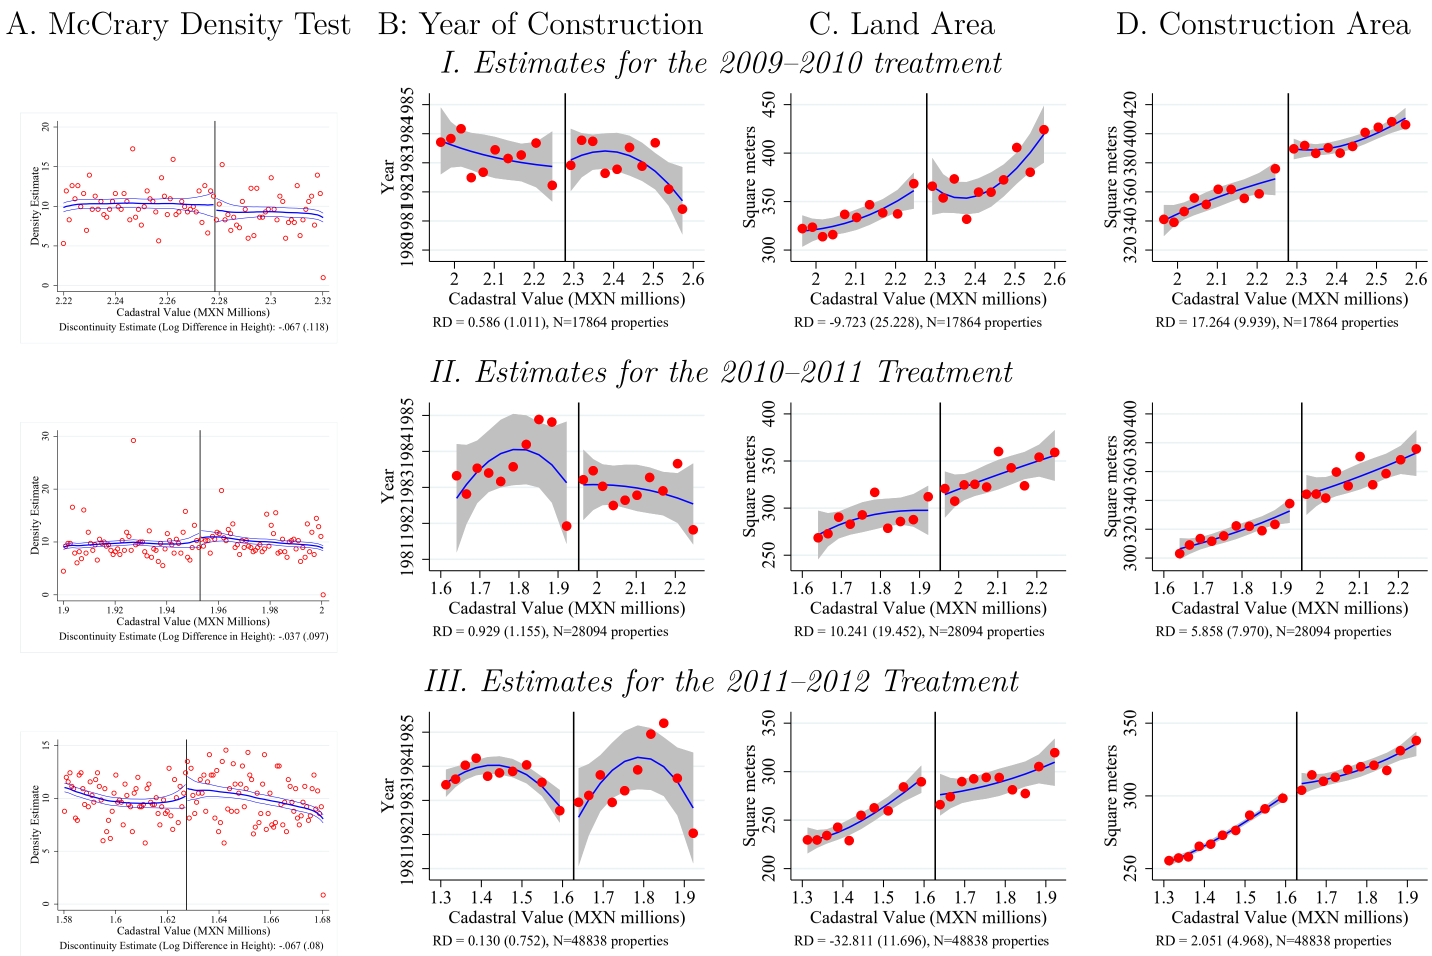

Notes: The panels in Column (A) display the density of properties by cadastral value around the lower threshold of bands I, H and G, which were treated with large tax rate increases. We test for a discontinuity in the density at the threshold, as proposed by McCrary 2008. In Columns (B)–(D), we test for discontinuities in property characteristics around the treatment thresholds and find no significant differences. The graphs are constructed as in Figures 3 and 4. These results are discussed in Section 4.1.

Figure D.2: Robustness of Regression Discontinuity Estimates by Bandwidth and Degree of Polynomial

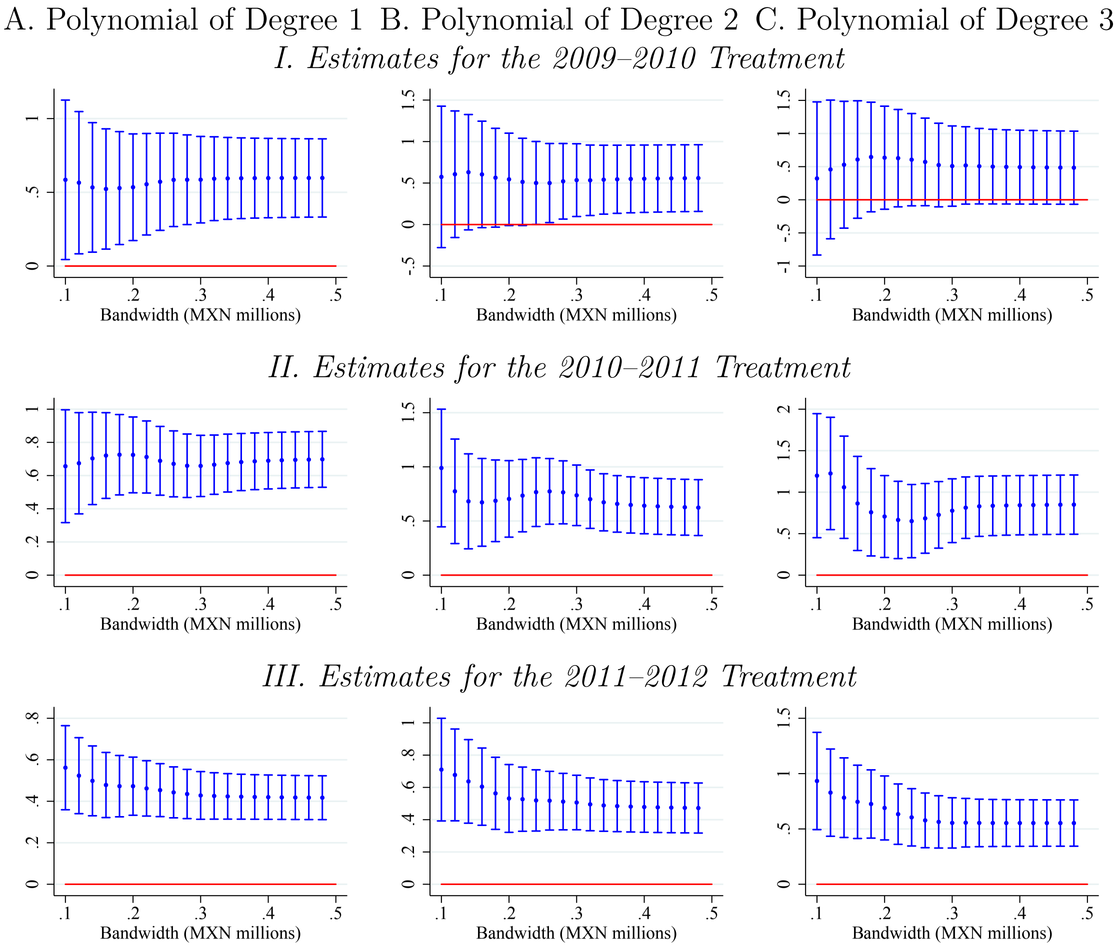

Notes: This figure documents the robustness of our RD estimations, as discussed in Section 4.1 and displayed in Figures 3. We focus on the payment amount as outcome. Each panel plots the point estimates and 95 percent confidence intervals of the treatment effect for different bandwidth values between 0.1 and 0.5 million MXN in 0.02 million increments (horizontal axis). Each row reports results for the different treatments, and each column presents the estimates for a different polynomial in cadastral value.

Table D.4: The Effect of Tax Rates on Tax Payment – Regression Discontinuity Estimates
Robustness to Dropping Taxpayers with Exemptions

	Mean tax rate (basis points)	Payment amount (MXN thousands)	Payment in full (percentage points)	Compliance share × 100
	(1)	(2)	(3)	(4)
	I. Estimates for the 2010 treatment			
T	9.115 ***	.592 **	-4.916 **	-2.836
	(.066)	(.259)	(2.33)	(2.184)
Properties	15190	15190	15190	15190
Adjusted R-squared	.98	.008	.003	.001
Mean at baseline (treated band)	50.109	5.635	39.227	46.239
Implied elasticity		.577	-.689	-.337
		(.252)	(.327)	(.26)
Implied Semi-elasticity		.012	-.014	-.007
		(.005)	(.007)	(.005)
	II. Estimates for the 2011 treatment			
T	12.141 ***	.679 ***	-6.18 ***	-3.268 **
	(.035)	(.133)	(1.717)	(1.298)
Properties	23882	23882	23882	23882
Adjusted R-squared	.993	.017	.003	.001
Mean at baseline (treated band)	47.473	4.546	37.465	45.552
Implied elasticity		.584	-.645	-.281
		(.114)	(.179)	(.111)
Implied Semi-elasticity		.012	-.014	-.006
		(.002)	(.004)	(.002)
P-value ($H_0 : \epsilon_{2011} = \epsilon_{2010}$)		.981	.906	.841
	III. Estimates for the 2012 treatment			
T	17.998 ***	.571 ***	-10.469 ***	-6.001 ***
	(.028)	(.092)	(1.535)	(1.295)
Properties	41040	41040	41040	41040
Adjusted R-squared	.996	.027	.008	.003
Mean at baseline (treated band)	41.073	3.367	40.627	46.04
Implied elasticity		.387	-.588	-.297
		(.062)	(.086)	(.064)
Implied Semi-elasticity		.009	-.014	-.007
		(.002)	(.002)	(.002)
P-value ($H_0 : \epsilon_{2012} = \epsilon_{2011}$)		.131	.775	.895
P-value ($H_0 : \epsilon_{2012} = \epsilon_{2010}$)		.465	.765	.882

Notes: This table is similar to Table 2, Panel A, but excludes taxpayers who ever received a subsidy or exemption. It demonstrates the robustness of the RD results to dropping these taxpayers.

E DiD Appendix

Figure E.1: The Effect of Tax Rates on Appeals

A. 2009–2010 Treatment

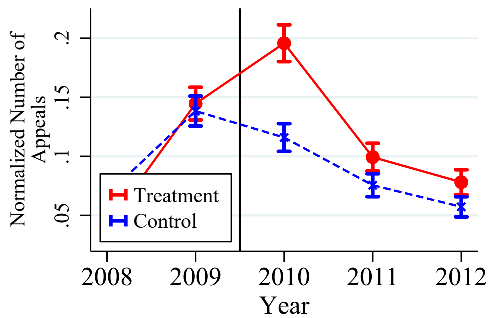

DD = 0.968 (0.097), N=5345 properties per year

B. 2010–2011 Treatment

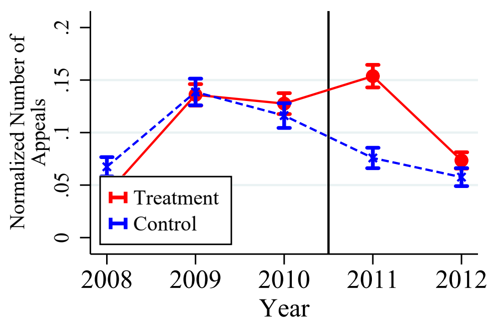

DD = 0.823 (0.069), N=7175 properties per year

C. 2011–2012 Treatment

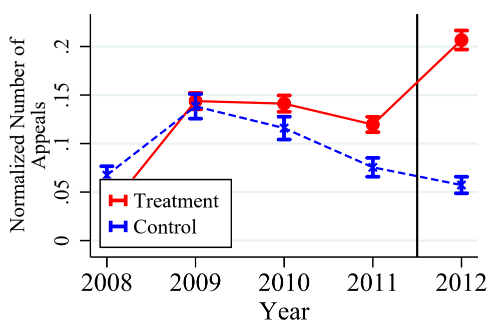

DD = 1.288 (0.057), N=9393 properties per year

Notes: This figure is identical to Figure 5, but the outcome variable is a dummy capturing whether the property owner appealed against the property tax liability. We restrict the sample to taxpayers making a payment every year, as we observe the appeals for these taxpayers only. Hardly any of the appeals result in a change in the tax liability (footnote 28). When calculating the DiD estimates displayed under the graphs, we take into account the first post-reform year only, as the graphical evidence suggests that the increase in appeals is primarily limited to the first post-reform year. Among taxpayers in the treatment group (in the restricted sample of taxpayers making a payment each year), the share filing an appeal is 14.46 percent in 2009, 12.76 percent in 2010 and 11.96 percent in 2011. These results are discussed in Section 4.2.

Table E.1: The Effect of Tax Rates on Tax Payment – Differences-in-Differences Estimates
Robustness to Using One Post-Reform Period Only

	Mean tax rate	Payment amount	Payment in full	Compliance share
	(1)	(2)	(3)	(4)
		I. Estimates for the 2010 treatment		
DD	.171 ***	.110 ***	-.136 ***	-.034 ***
	(.000)	(.018)	(.019)	(.012)
Adjusted R-squared	.996	.014	.006	.002
Properties (treatment)	5747	5747	5747	5747
Properties (control)	6510	6510	6510	6510
Implied elasticity		.643	-.794	-.200
		(.106)	(.113)	(.07)
		II. Estimates for the 2011 treatment		
DD	.236 ***	.163 ***	-.178 ***	-.063 ***
	(.000)	(.022)	(.017)	(.011)
Adjusted R-squared	.993	.008	.006	.009
Properties (treatment)	9661	9661	9661	9661
Properties (control)	6511	6511	6511	6511
Implied elasticity		.690	-.753	-.267
		(.092)	(.071)	(.047)
P-value ($H_0 : \epsilon_{2011} = \epsilon_{2010}$)		.777	.812	.279
		III. Estimates for the 2012 treatment		
DD	.401 ***	.196 ***	-.300 ***	-.182 ***
	(.000)	(.029)	(.014)	(.010)
Adjusted R-squared	.994	.003	.016	.038
Properties (treatment)	15227	15227	15227	15227
Properties (control)	6508	6508	6508	6508
Implied elasticity		.489	-.747	-.455
		(.073)	(.035)	(.025)
P-value ($H_0 : \epsilon_{2012} = \epsilon_{2011}$)		.016	.021	.000
P-value ($H_0 : \epsilon_{2012} = \epsilon_{2010}$)		.023	.048	.000

Notes: This table is similar to Table 2, Panel B, but uses one post-reform period only for each reform episode. This confirms the robustness of our main DiD estimates, in which we use between one and three post-reform years, depending on the reform. These results are discussed in Section 4.2.

Table E.2: The Effect of Tax Rates on Tax Payment – Differences-in-Differences Estimates Robustness to Different Control Bands

	Mean Tax Rate (basis points)	Payment amount (MXN thousands)	Payment in full	Compliance share
	(1)	(2)	(3)	(4)
	I. Estimates for the 2009-2010 treatment			
K and L	.166	.116	-.103	-.041
	(.000)	(.017)	(.017)	(.011)
E and F	.255	.169	-.029	-.001
	(.000)	(.010)	(.013)	(.009)
M and N	.275	.163	.001	.024
	(.000)	(.010)	(.013)	(.009)
K,L,E,F,M and N	.268	.165	-.018	.01
	(.000)	(.010)	(.013)	(.009)
	II. Estimates for the 2010-2011 treatment			
K and L	.232	.169	-.149	-.07
	(.000)	(.02)	(.015)	(.010)
E and F	.331	.175	-.073	-.038
	(.000)	(.008)	(.010)	(.007)
M and N	.355	.166	-.021	.018
	(.000)	(.008)	(.010)	(.007)
K,L,E,F,M and N	.347	.17	-.049	-.008
	(.000)	(.008)	(.010)	(.007)
	III. Estimates for the 2011-2012 treatment			
K and L	.401	.196	-.3	-.182
	(.000)	(.029)	(.014)	(.010)
E and F	.517	.16	-.216	-.145
	(.000)	(.009)	(.009)	(.007)
M and N	.545	.138	-.161	-.087
	(.000)	(.009)	(.009)	(.007)
K,L,E,F,M and N	.536	.151	-.189	-.114
	(.000)	(.009)	(.009)	(.006)

Notes: This table is similar to Table 2, Panel B, but displays the point estimates from the DiD estimation using alternative value bands as control groups. In each panel, the first line reproduces our preferred estimates from Table 2, using properties in bands K and L as the control group. The remaining lines display estimates using alternative control groups. Using band J as a control group is not an option, as properties in this band were treated in 2008–2009. The table shows that our preferred specification yields results similar to those of the alternative specifications but slightly less optimistic about the tax rate increases. In our preferred specification, tax payments increase a bit less and compliance falls a bit more than in alternative specifications. Yet we still conclude that tax rate increases are welfare improving. Choosing a different specification would only strengthen this conclusion. These results are discussed in Section 4.2.

Figure E.2: Parallel Trends with Different Control Groups – 2009–2010 Treatment

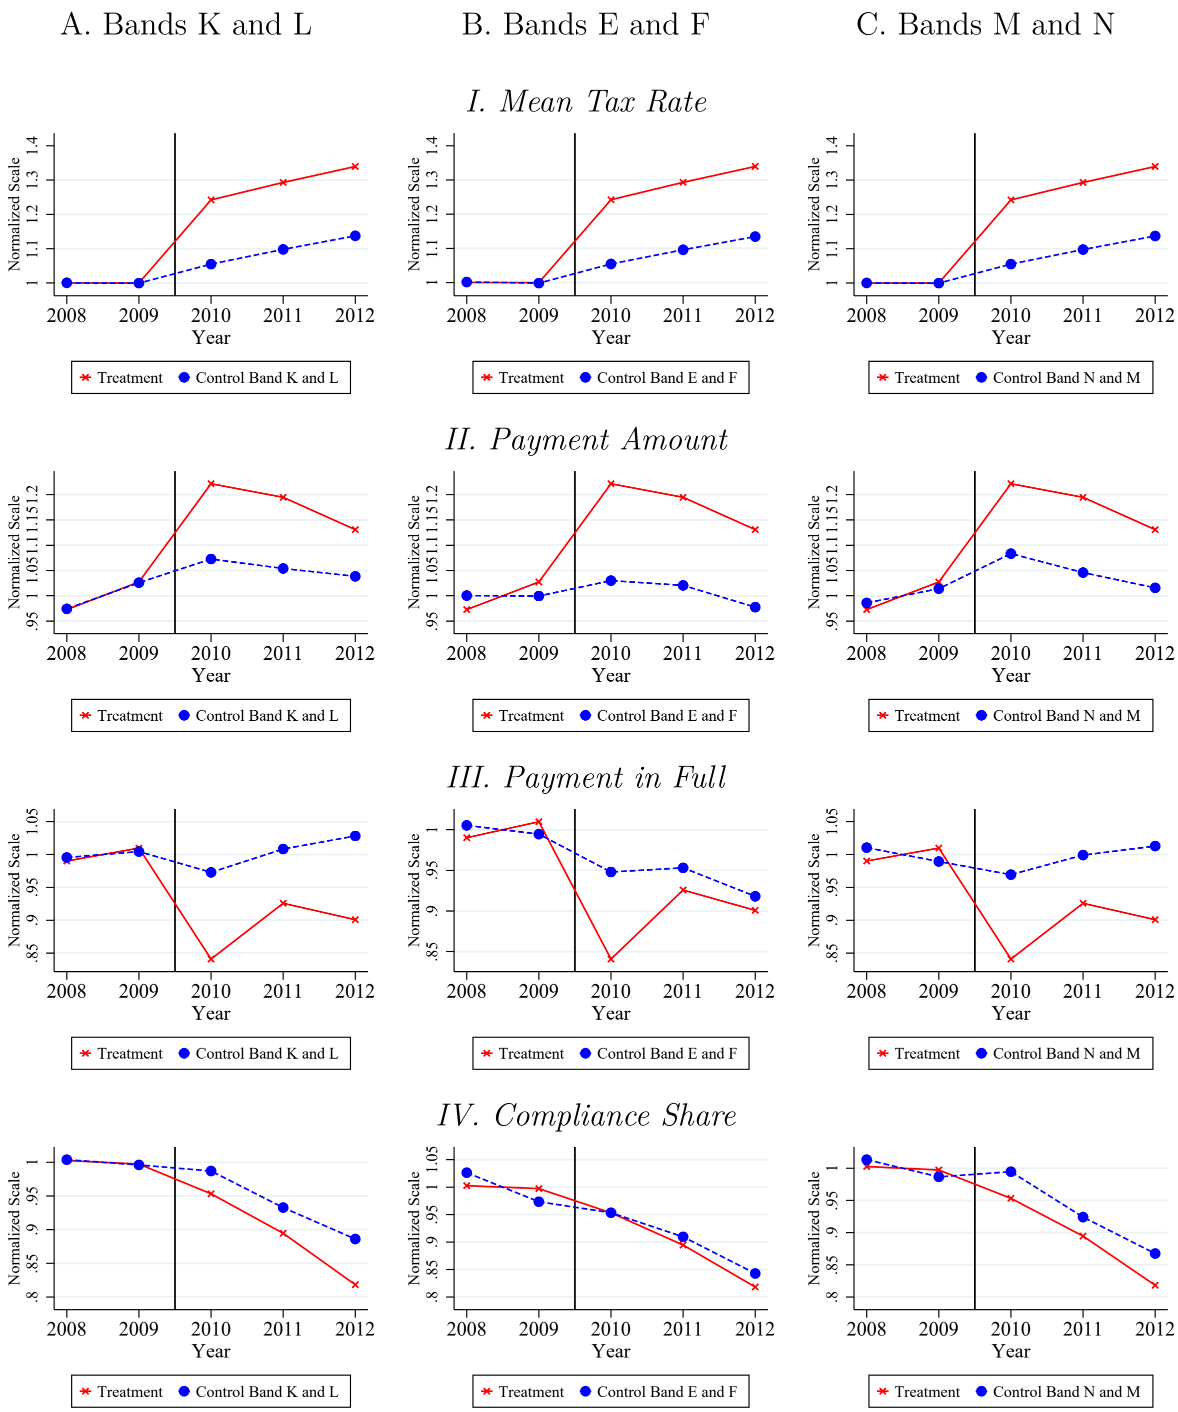

Notes: This figure is similar to Figures 5 and 6 but considers two alternative control groups: bands E and F Column (B)), and bands M and N (Column (C)). Using band J as a control group is not an option, as properties in this band were treated in 2008–2009. Column (A) reproduces the preferred specification in which bands K and L are the control group. The treated band in this figure contains properties in band I, as the figure focuses on the 2009–2010 reform, in which properties in band I were treated with a quasi-exogenous tax rate increase. Figures E.3 and E.4 show similar graphs, focusing on the 2010–2011 and the 2011–2012 reforms, respectively. The corresponding DiD estimates for all reform episodes are displayed in Table E.2. These results are discussed in Section 4.2.

Figure E.3: Parallel Trends with Different Control Groups – 2010–2011 Treatment

A. Bands K and L B. Bands E and F C. Bands M and N

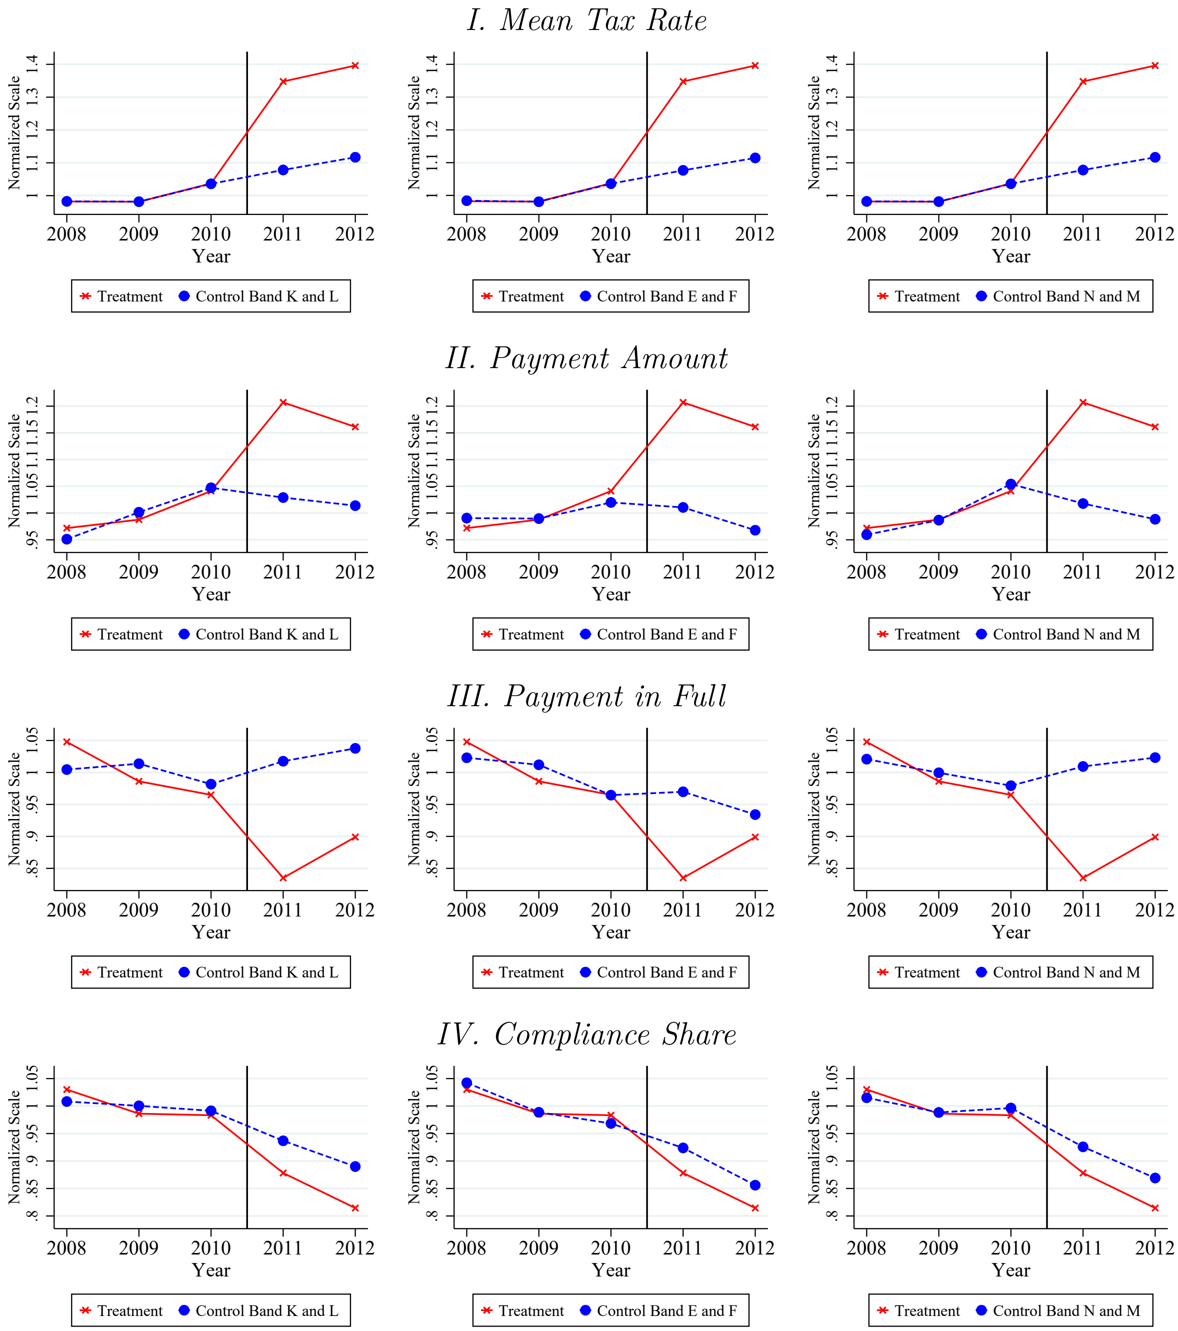

Notes: This figure is identical to Figure E.2 but focuses on the 2010–2011 reform, in which properties in band H were treated with a quasi-exogenous tax rate increase. The treatment group in this figure thus contains properties in band H. The corresponding DiD estimates are displayed in Table E.2.

Figure E.4: Parallel Trends with Different Control Groups – 2011–2012 Treatment

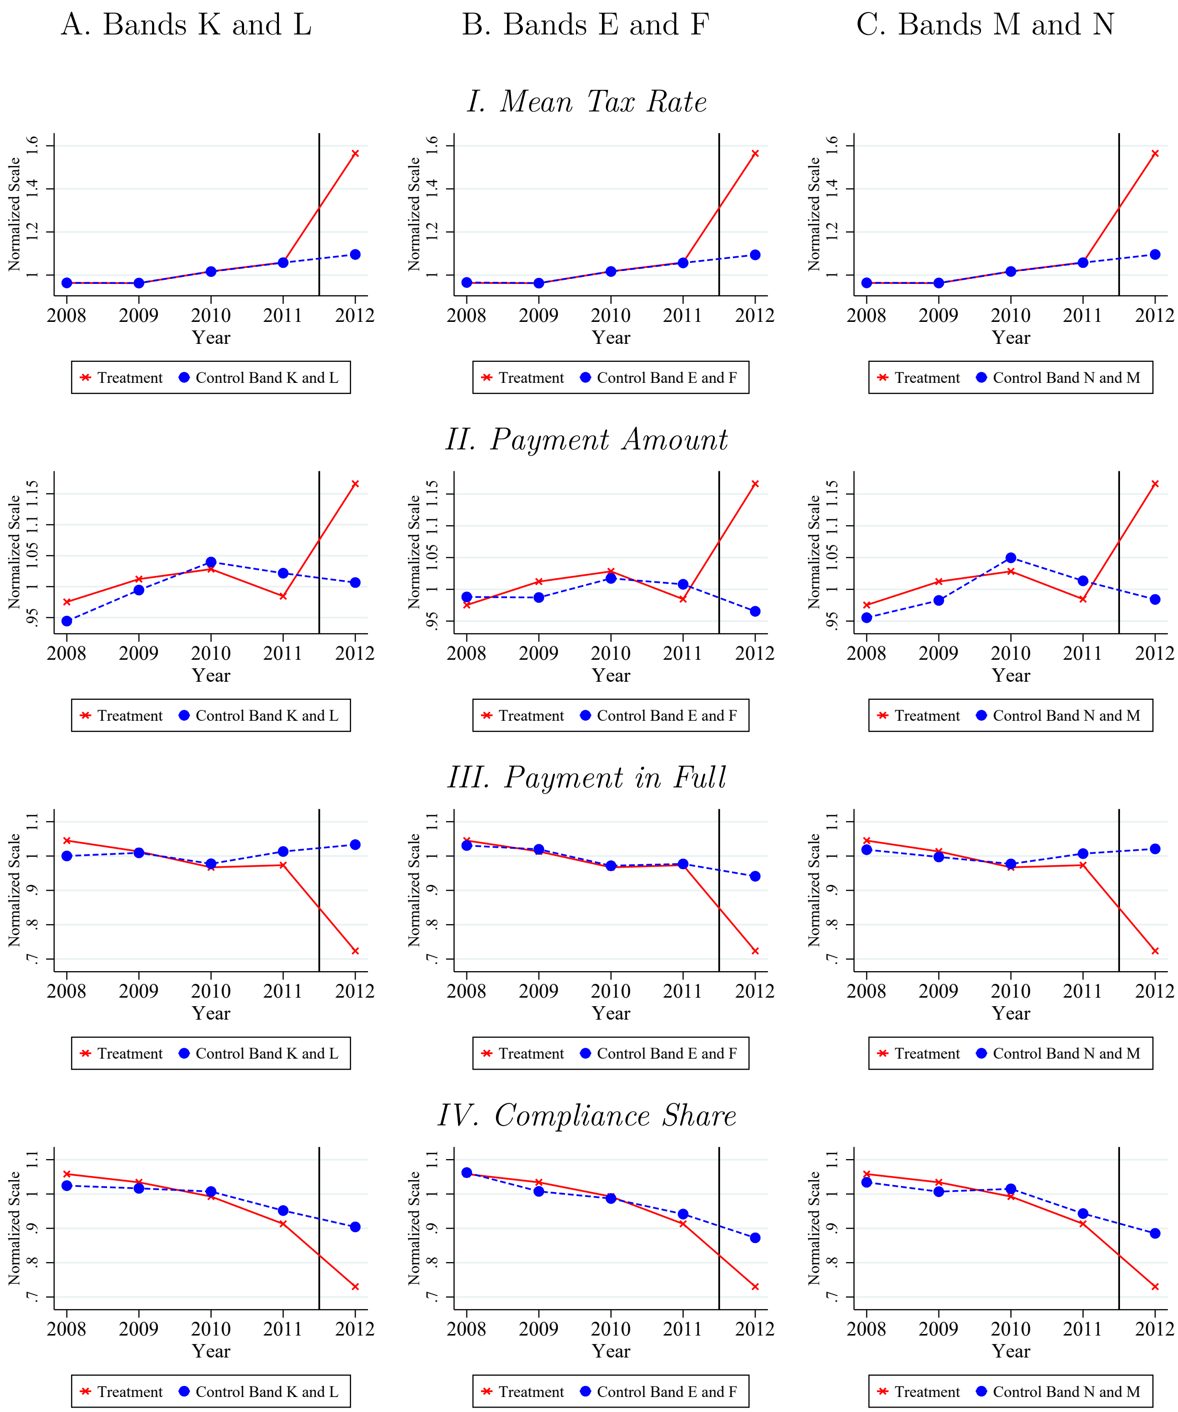

Notes: This figure is identical to Figure E.2 but focuses on the 2011–2012 reform, in which properties in band G were treated with a quasi-exogenous tax rate increase. The treatment group in this figure thus contains properties in band G. The corresponding DiD estimates are displayed in Table E.2.

Table E.3: The Effect of Tax Rates on Tax Payment – Differences-in-Differences Estimates Robustness to Dropping Taxpayers with Exemptions

	Mean tax rate	Payment amount	Payment in full	Compliance share
	(1)	(2)	(3)	(4)
		I. Estimates for the 2010 treatment		
DD	.167 ***	.084 ***	-.093 ***	-.048 ***
	(.000)	(.019)	(.018)	(.012)
Adjusted R-squared	.998	.011	.005	.004
Properties (treatment)	4854	4854	4854	4854
Properties (control)	5530	5530	5530	5530
Implied elasticity		.502	-.558	-.289
		(.112)	(.107)	(.074)
Implied Semi-elasticity		.01	-.011	-.006
		(.002)	(.002)	(.001)
		II. Estimates for the 2011 treatment		
DD	.232 ***	.107 ***	-.142 ***	-.077 ***
	(.000)	(.021)	(.015)	(.011)
Adjusted R-squared	.995	.012	.004	.008
Properties (treatment)	8194	8194	8194	8194
Properties (control)	5531	5531	5531	5531
Implied elasticity		.463	-.611	-.332
		(.092)	(.066)	(.047)
Implied Semi-elasticity		.01	-.013	-.007
		(.002)	(.001)	(.001)
P-value ($H_0 : \epsilon_{2011} = \epsilon_{2010}$)		.489	.331	.393
		III. Estimates for the 2012 treatment		
DD	.401 ***	.14 ***	-.282 ***	-.16 ***
	(.000)	(.031)	(.015)	(.011)
Adjusted R-squared	.993	.009	.013	.019
Properties (treatment)	12928	12928	12928	12928
Properties (control)	5528	5528	5528	5528
Implied elasticity		.35	-.705	-.399
		(.077)	(.037)	(.027)
Implied Semi-elasticity		.009	-.017	-.01
		(.002)	(.001)	(.001)
P-value ($H_0 : \epsilon_{2012} = \epsilon_{2011}$)		.182	.000	.006
P-value ($H_0 : \epsilon_{2012} = \epsilon_{2010}$)		.071	.000	.002

Notes: This table is similar to Table 2, Panel B, but excludes taxpayers who ever received a subsidy or exemption. It demonstrates the robustness of the DiD results to dropping these taxpayers. These results are discussed in Section 4.2.

Table E.4: The Effect of Tax Rates on Tax Payment – Differences-in-Differences Estimates
Robustness to Different Clustering Levels

		Standard Error Cluster Level:		
	Point Estimate	Taxpayer	Postal Code	Bootstrap Delegacion
	(1)	(2)	(3)	(4)
	I. Estimates for the 2009-2010 treatment			
Mean Tax Rate (basis points)	.166	0	0	0
Payment amount (MXN thousands)	.116	.017	.015	.01
Payment in full (percentage points)	-.103	.017	.017	.015
Compliance share × 100	-.041	.011	.011	.008
	II. Estimates for the 2010-2011 treatment			
Mean Tax Rate (basis points)	.232	0	0	.001
Payment amount (MXN thousands)	.169	.02	.024	.022
Payment in full (percentage points)	-.149	.015	.016	.01
Compliance share × 100	-.07	.01	.011	.008
	III. Estimates for the 2011-2012 treatment			
Mean Tax Rate (basis points)	.401	0	.001	.001
Payment amount (MXN thousands)	.196	.029	.029	.039
Payment in full (percentage points)	-.3	.014	.015	.014
Compliance share × 100	-.182	.01	.01	.011

Notes: This table is similar to Table 2, Panel B. Column (1) displays the point estimates from the DiD estimation. Columns (2)–(4) display the standard errors at different clustering levels. Column (4) performs a bootstrap where the 16 *delegaciones* are the resampling clusters.

Table E.5: Robustness of Tax Rate Elasticity Estimations – Panel Regressions

	Mean tax rate (1)	Payment amount (2)	Payment in full (3)	Compliance share (4)
		A. Value Band FE		
PE Tax Rate		.694 ***	-.203 ***	-.132 ***
		(.086)	(.046)	(.045)
Tax Rate Elasticity		.694	-.551	-.292
		(.086)	(.125)	(.098)
Mean Tax Rate	49.53			
Mean		11301.63	.37	.45
Cadastral Value Band FE		Yes	Yes	Yes
Year FE		Yes	Yes	Yes
Band G, H, I * Tax Rate		No	No	No
Observations		80	80	80
		B. Heteregeneity by treated bands		
PE Tax Rate		.702 **	-.294 **	-.115
		(.275)	(.146)	(.142)
PE Treated Bands * Tax Rate		-.006	.072	-.014
		(.205)	(.109)	(.106)
Tax Rate Elasticity Control Bands		.702	-.799	-.254
		(.275)	(.396)	(.313)
Tax Rate Elasticity Treated Bands		.696	-.604	-.284
		(.038)	(.055)	(.044)
Mean Tax Rate	49.53			
Mean		11301.63	.37	.45
Cadastral Value Band FE		Yes	Yes	Yes
Year FE		Yes	Yes	Yes
Band G, H, I * Tax Rate		Yes	Yes	Yes
Observations		80	80	80

Notes: This table demonstrates the robustness of the estimates for the elasticity of tax compliance to the tax rate, displayed in Tables 2. Panel A here displays estimates for $Y_{it} = \beta_1 R_{it} + \gamma_i + \delta_t + \epsilon_{it}$, where Y_{it} is the average outcome for band i in year t, R_{it} is the log average tax rate for band i in year t , δ_t and γ_i denote year and value band fixed effects, and ϵ_i is the error term. Payment is in logs; the other outcomes are not. The elasticities are calculated as in Table 2. Standard errors are in parentheses. Panel B is similar to Panel A but includes an interaction between the tax rate and an indicator for bands G, H and I (the treated bands in the RD and DiD estimations). The elasticity estimates are very similar to our main RD and DiD estimates. These results are discussed in Section 4.2.

F Real Response Appendix

A potential concern for tax rate increases is that while they may raise revenue, they may also be accompanied by a reduction in real estate investment. One possibility is that owners of existing properties might be less likely to invest in maintaining or upgrading their properties. However, because the cadastral value of a given property would not be affected by maintenance or upgrades, raising property tax rates does not disincentivize these forms of investment. Nonetheless, it may be the case that higher property tax rates can decrease investment in new housing units.

One way in which developers can respond to tax rate hikes is by using the cadastral value formula to design houses so that they fall below the threshold of a given value band. To test for this, Panels A and B of Figure F.1 plot the number of properties built after the 2010 and 2011 tax increases.[66] We group properties into small cadastral value bins around the lower threshold of each treated value band. If the tax rate increase dissuaded developers from building properties in this value band, we should see a bunching of new properties just below the lower threshold of the treated band. The graphs show that this is not the case: the number of new properties is weakly decreasing with property value and is smooth around the threshold of the treated band. That is, we find no evidence of bunching. To further investigate this type of response, in Panels C and D of Figure F.1, we plot the percentage change in new properties by bin, relative to the average number of new units in the previous two years. As in the previous graphs, there is no sign of bunching and no discontinuity at the threshold of the treated band.

Building on our analyses in Section 4.2, we also conduct a DiD analysis on the number of new property developments. To do so, we first rank properties by cadastral value and divide each value band into 5 sub-bins of equal size. We then construct a count of the number of new properties in each sub-bin and year. Finally, we estimate a regression similar to Equation 7 where the outcome is the log number of new properties at the sub-bin–year level. Panel E in Figure F.1 shows the results of this estimation where we stack the 2010 and 2011 reforms and where we use properties in bands K and L as controls. This figure shows that we do not find a decrease in the number of new properties in bands that experience increases in property tax rates.

Although other studies have found evidence of real responses to property tax changes (e.g., Singh 2020), our results can be rationalized when we consider the context. The quasi-exogenous tax rate increases that we study apply to a very small range of property values. To determine whether their future properties would fall into one of the treated bands, property developers need precise knowledge of the tax code and clarity on the exact features of the property to be constructed. The applicability of the tax rate changes may be too narrow to warrant such an analysis. In addition, developers may anticipate future tax rate changes, reducing their responsiveness to recent reforms.

Overall, we do not find evidence that increases in property tax rates disincentivized the construction of new housing units. Based on these results, we focus the policy analysis in Section 7 on the roles that compliance, enforcement, and liquidity constraints play in the administration of the property tax.

[66] We exclude the 2012 reform because our data end in 2013, which limits the number of new units that we observe.

Figure F.1: Absence of a Real Estate Response to Tax Rate Increases

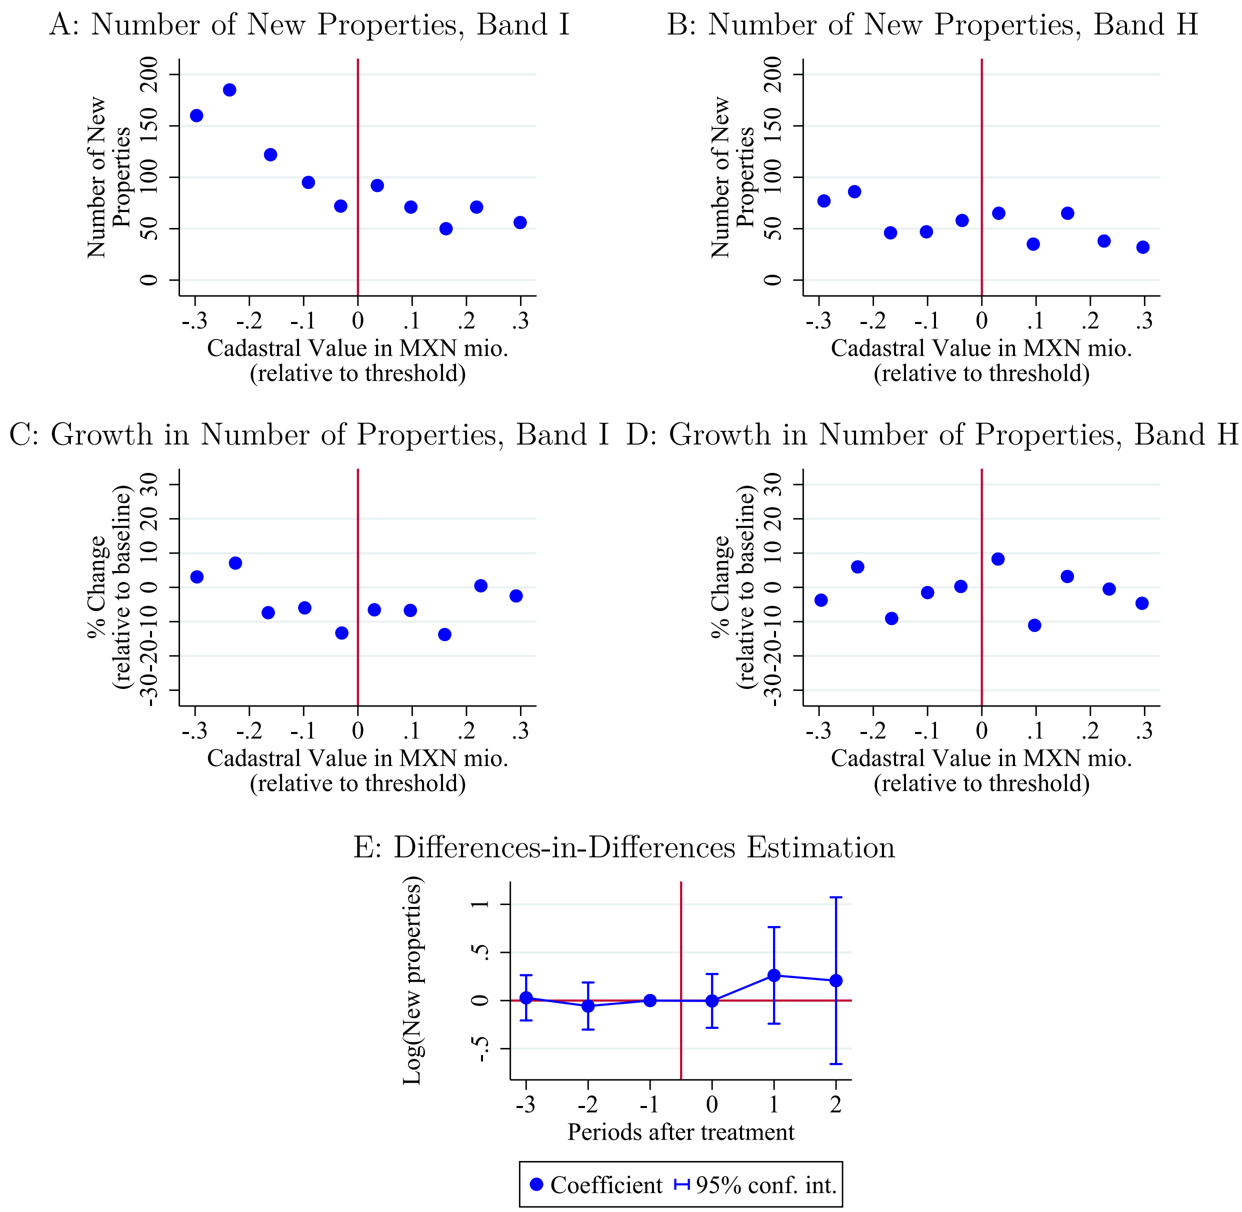

Notes: This figure examines the effect of tax rate increases on real estate investment. Panel A plots the number of new properties constructed around the lower threshold of band I during the post-reform years 2011–2013. We plot the number of new properties in equally sized cadastral value bins in bands H and I, within a 0.3 million MXN cadastral value range around the threshold. Panel B is similar but plots the number of new properties in bands G and H, constructed in the years 2012 and 2013. Panels C and D are similar to Panels A and B, respectively, but plot the growth rate in new properties. The numerator of the growth rate is the number of new properties plotted in Panel A (B). The denominator is the yearly average number of properties in the last two (three) pre-reform years (we use an average to minimize noise). Panel E plots the results of the difference-in-difference estimation $Log(N_{bt}) = \alpha_b + \mu_t + \gamma \cdot Treat_b \cdot Post_t + \epsilon_{bt}$, where N_{bt} is the number of new properties constructed in property value bin b in year t and α_b and μ_t are bin and time fixed effects. Each value band is divided into equally sized bins, and standard errors are clustered at the bin level. Value bands I, H, and G are treated; value bands J, K and L serve as controls. The other bands are omitted. The reforms are stacked, so that t is the time relative to reform.

G Experiment Appendix

Table G.1: Cadastral Value Distribution of Experiment Sample

	All Taxpayers	Experiment Sample
Mean	589,530.8	533,087.3
SE	(636.2)	(2310.3)
Min	993.7	17,178.4
Max	11,711,063.3	11,670,532.6
20th Percentile	229,784.6	256,034.0
50th Percentile	391,487.2	419,170.5
80th Percentile	730,281.1	678,949.6

Notes: As discussed in Section 5.1, this table compares the cadastral value between the delinquent taxpayers targeted in the tax compliance intervention and the full population of taxpayers, showing very similar distributions.

Figure G.1: Experiment Design

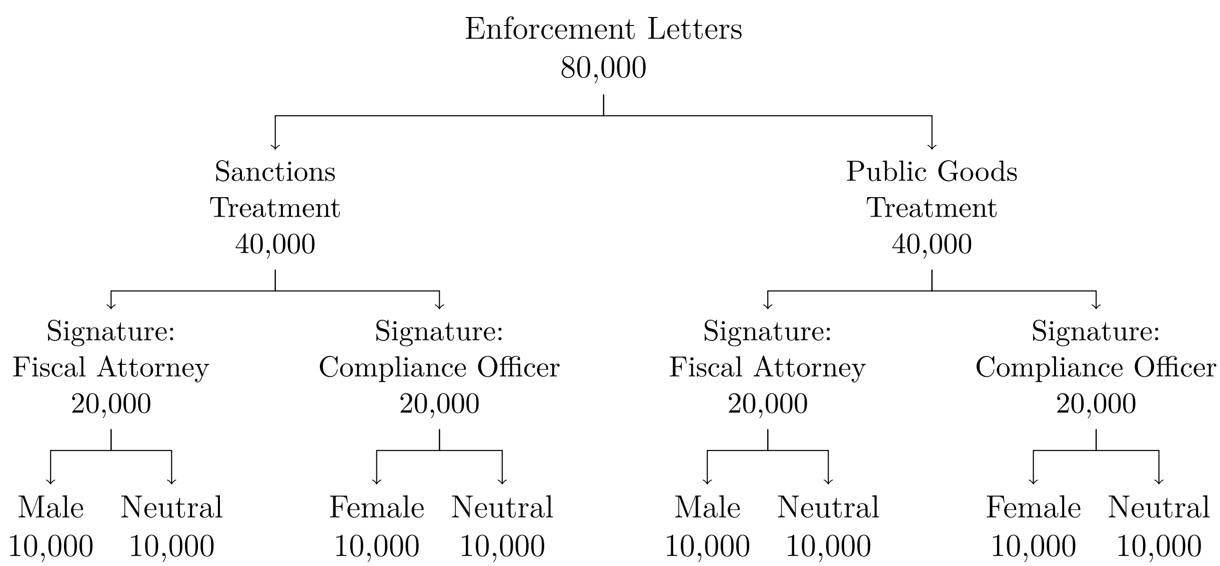

Notes: This diagram represents the different treatment arms of the enforcement intervention discussed in Section 5, in which the Ministry of Finance sent letters to encourage payment of outstanding property tax debt. Letter recipients were selected from a pool of taxpayers who had become delinquent between bimester 4 of 2009 and bimester 3 of 2014. The letters were sent between July 28 and August 11, 2014. A control group of 10,000 delinquent taxpayers received no letters.

Table G.2: The Effect of Enforcement Letters on Tax Payment: Robustness

1: Letter Content		2: Sender Position		3: Male Sender		4: Female Sender	
Sanctions	Public Good	Fiscal Attorney	Compliance Officer	Male	Neutral	Female	Neutral
Panel A: Cross Sectional Regressions with Controls							
Panel A1: Any Payment							
9.463***	4.815***	7.585***	6.637***	7.059***	8.07***	7.039***	6.308***
(.312)	(.292)	(.304)	(.3)	(.361)	(.364)	(.357)	(.353)
Panel A2: Payment Amount (MX Pesos)							
39.018***	9.306***	28.95***	18.961***	27.041***	30.98***	21.843***	17.169***
(1.98)	(1.678)	(1.872)	(1.801)	(2.36)	(2.251)	(2.209)	(2.151)
Panel B: Difference in Difference Estimates							
Panel B1: Any Payment							
9.591***	4.975***	7.814***	6.751***	7.173***	8.455***	7.004***	6.499***
(.256)	(.238)	(.249)	(.245)	(.296)	(.302)	(.293)	(.291)
Panel B2: Payment Amount (MX Pesos)							
55.564***	16.049***	43.656***	27.956***	40.283***	47.029***	29.431***	26.482***
(2.292)	(1.957)	(2.192)	(2.075)	(2.642)	(2.648)	(2.454)	(2.444)

Notes: This table is similar to Table 3. Each of Columns (1)–(4) and each row displays results from one regression. Panel A displays estimates from a cross-sectional regression that includes the cadastral value and age of the property as controls in the estimation. Panel B displays estimates for β_5 and β_6 from the difference-in-difference specification $Y_{igt} = \alpha_i + \lambda_t + \beta_1 T1_{i20} + \beta_2 T2_{i20} + \beta_3 T1_{i40} + \beta_4 T2_{i40} + \beta_5 T1_{i60} + \beta_6 T2_{i60} + \epsilon_{igt}$, where Y_{igt} is the outcome for property i in treatment group g on day t, α_i and λ_t are property and day fixed effects, $T1_{i20}$ and $T2_{i20}$ are treatment indicators that switch on for the first 20 days after all letters were sent, $T1_{i40}$ and $T2_{i40}$ are treatment indicators that switch on between days 20 and 40, and $T1_{i60}$ and $T2_{i60}$ are treatment indicators that switch on between days 40 and 60.

Figure G.2: The Effect of Enforcement Letters on Tax Payment

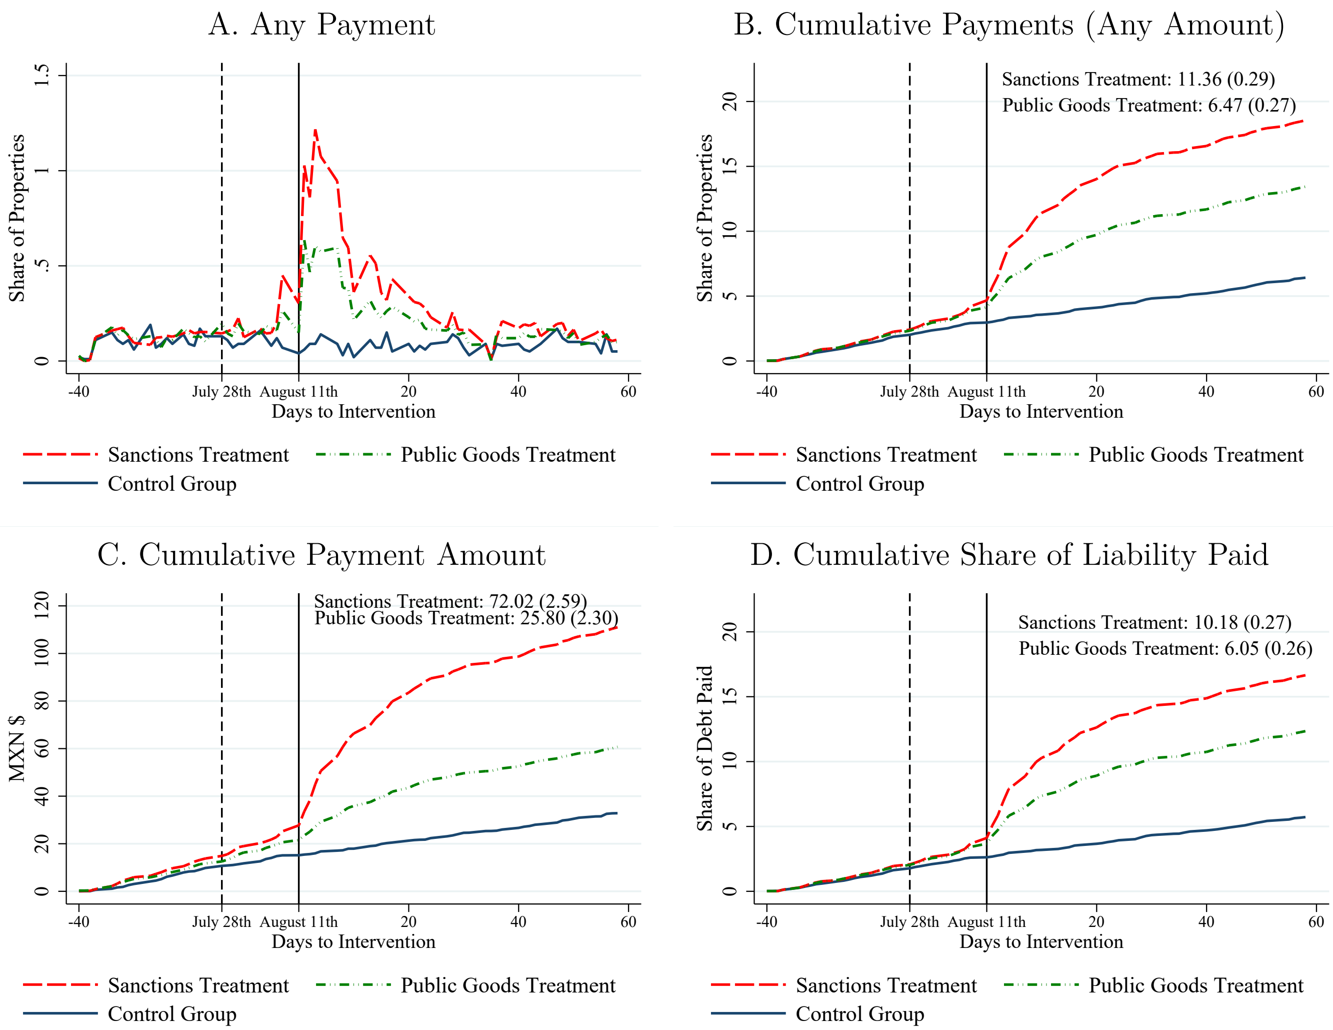

Notes: This figure is identical to Figure 7 except that it displays the raw data before detrending and the estimates are from a regression run on the non-detrended data. This results in slightly larger treatment effect estimates than those of our preferred specification.

Figure G.3: Translation of Enforcement Letters

Taxpayer name
Taxpayer address

Sanctions Treatment

Avoid major inconvenience and regularize your tax status

Invitation Letter for Payment of Property Tax

According to the registers of the Federal District's Treasury, you have outstanding **property tax** debt for the tax period(s) *****. We would therefore be grateful if you could update your tax status within 15 working days of receipt of this letter.

Delay in property tax payment can be sanctioned with fines and interest costs, and with interventions that the fiscal authority conducts to ensure effective tax collection, as per the Tax Code, **which can lead to the seizure of property.**

Avoid major inconvenience and regularize your tax status.

Public Goods Treatment

With our tax payment, we all contribute to improving our city

Invitation Letter for Payment of the Property Tax

As you know, a large part of the social programs and investments in infrastructure and security that the Government of Mexico City implements are financed by property tax revenues. Your contribution is therefore very important, and we would be pleased if you could update your property tax account as soon as possible and cover the outstanding tax debt for the above mentioned building for the tax period(s) ***** within 15 working days upon receipt of this letter. We ask you to update your account to avoid incurring surcharges.

With the revenues obtained from property taxes in your city, we finance the following public goods, among others:

- Food pensions for the elderly;
- School uniforms and school supplies for children;
- The operation of health centers and hospitals of the Government of Mexico City;
- Street lights and sidewalks in your neighborhood.

Boxed: Information about payment and further details on the back

Signature: Name, Title
(For gender-neutral signatures, only the initials of the first name are provided.)

H Payment Timing Appendix

As we show in Section 1, the interaction between property taxes and liquidity constraints implies that governments can improve tax administration by providing liquidity. The government of Mexico City provides liquidity by offering discounts to taxpayers who pay in full before a given deadline.[67] In this section, we study how the timing of payments is affected by changes in discounts—through both deadlines and rates—and use this variation to infer households' value for liquidity.

Households choose to pay the property tax on a given day by trading off a lower overall tax payment against the cost of giving up interest-bearing liquid assets and the hassle cost of paying taxes on a given day. In the absence of a value for liquidity, taxpayers would prefer to pay ahead of the deadline to avoid uncertain hassle costs and time constraints that might prevent them from obtaining the discount. If households face liquidity constraints, they may risk missing out on the discount for the benefit of holding on to liquid assets until immediately before the deadline.

The government offered up to three types of discounts between 2009 and 2013. Figure H.1 plots the discount rate and the histogram of property tax revenue for every year. The size of the discounts and the deadlines varied substantially over the years. These discount deadlines can be interpreted as time notches in the sense of Slemrod (2013). Figure H.1 shows a clear pattern where taxpayers respond very strongly to these incentives by bunching payments immediately before the deadline. This strong response to discounts is consistent with a high value for liquidity.[68]

To quantify the value for liquidity, we use this variation to estimate a dynamic model of payment timing where households trade off discounts, interest income, and the hassle costs of paying the tax.[69] For a given date t in year y, households obtain utility $v_0(t, y)$ when they pay their taxes. $v_0(t, y) = \theta_{t,y} + \theta_1 \text{Tax}(t, y)$, where $\theta_{t,y}$ captures the benefit of paying taxes net of the hassle cost of paying on day t and where $\text{Tax}(t, y)$ equals one minus the discount applicable on day t and year y.[70] Alternatively, households can choose to delay paying taxes at time t. The value of this choice is given by $v_1(t, y) = \theta_2 \text{Interest}_y + \beta EV_1(t + 1, y)$, which captures the interest income from delaying payment and the discounted value function in the next period. Finally, every day households face idiosyncratic hassle costs of paying taxes that follow a logistic distribution.

We study the timing of payments by modeling $P_0(t, y)$: the probability of paying on any given day conditional on not having paid yet. Following Hotz and Miller (1993), $P_0(t, y)$ captures the value of delaying payment since $EV_1(t+1, y) = v_0(t+1, y) - \ln P_0(t+1, y) + \gamma$, where γ is Euler's constant.

[67]Discounts provide liquidity by effectively lending money to those who do not take up the discount. Assuming the discount rate is d, households can pay $1 - d$ today or borrow from the government at the rate of $\frac{d}{1-d}$. Governments provide liquidity through several mechanisms, including unemployment insurance schemes or accelerated depreciation deductions for firms.

[68]It is worth noting that leading behavioral models are not consistent with these data patterns. First, because deadlines change across years, the bunching patterns that we observe are not consistent with the salience of specific dates—e.g., first-of-the-month effects. Second, because the bunching patterns are very pronounced, these patterns are likely not driven by rational inattention. Finally, one may worry that these patterns are a result of hyperbolic discounting. Fang and Silverman (2004) develop a model of present-biased preferences to study the effects of time limits in public policy. A prediction of this model is that present-biased households would likely miss the deadline, which is inconsistent with the sharp bunching patterns in the data.

[69]Hassle costs include time spent withdrawing cash and traveling to a payment location. Note that automatic payments were not allowed and online payments are only now being introduced.

[70]The problem ends when households pay the tax, so that action 0 constitutes a terminal option.

The relative log-likelihood of paying on any given day is then:

$$\begin{aligned}\ln\left(\frac{P_0(t,y)}{1-P_0(t,y)}\right) &= v_0(t,y)-v_1(t,y)=v_0(t,y)-\beta v_0(t+1,y)-\theta_2\text{Interest}_y-\beta\ln P_0(t+1,y)-\beta\gamma \\ &= -\theta_1\{\text{Tax}(t,y)-\beta\text{Tax}(t+1,y)\}-\theta_2\text{Interest}_y \\ &\quad -\beta\ln P_0(t+1,y)+\theta_{t,y}-\beta\theta_{t+1,y}.\end{aligned} \tag{H.1}$$

The first line follows from the logistic distribution and the expression for $EV_1(t+1,y)$ above. The second line substitutes the definitions of choice value $v_0(t,y)$. To implement this model, we assume that the hassle costs $\theta_{t,y}$ have three components: a day-of-the-year effect, θ_t; a day-of-the-week effect, $\theta_{t,d}$; and a residual component, $\varepsilon_{t,y}$.[71] The identifying assumption of this equation is that conditional on day-of-the year and day-of-the-week fixed effects, the daily changes in residual time costs, $\Delta_t\varepsilon_{t,y}$, are unrelated to yearly variation in interest rates or in the size and timing of discounts. This assumption is plausible since Interest_y is set by the broader market and since the policy variation in $\text{Tax}(t,y)$ features significant changes in the number of deadlines, due dates, and magnitude of the discounts that are unrelated to daily hassle costs.

Given the rich variation in discounts, we estimate the parameters of the dynamic discrete choice model $(\theta_1,\theta_2,\beta)$ via non-linear least squares, where the fixed effects recover day-of-the-year and day-of-the-week hassle costs. Figure H.2 plots $P_0(t,y)$ along with the model fit and shows that this relatively simple model does a remarkably good job of matching the data patterns. Panel A of Table H.1 reports the main estimates from this model. These estimates reveal the importance of liquidity concerns in a couple of ways. First, we can interpret the discount rate as implying that households would pay an interest rate of approximately $8\%(\frac{1}{\beta}-1)$ to delay paying property taxes by a single day.[72] This very high rate of discounting highlights the value of liquidity. Second, the effect of changes in taxes is only about 10 $(\approx\frac{\theta_1}{\theta_2})$ times larger than changes in interest income. In a world without liquidity constraints, households would place a much smaller value on short-term interest relative to the value placed on a permanent discount on their taxes.

We now use the estimated model to quantify the liquidity value provided by the discounts. Panel B of Table H.1 shows the gain in consumer surplus relative to a world without discounts.[73] Column (2) shows that taxpayers value discounts at between 3.26 and 10.35 percent of the value of their property taxes. Comparing these values to the deadlines in Column (3), we find higher values in years with later deadlines. This makes sense: discounts raise welfare more when they allow households to hold on to liquid assets longer. Later deadlines also lower the risk of facing a high hassle cost in the early days of the year. Column (4) lists the realized fiscal cost of the discounts as a percentage of the property tax. As would be expected, the value of the discounts in Column (2) is positively related to the fiscal costs. Finally, Column (5) displays the ratio of the value of the discounts to the fiscal cost (Columns (2) over (4)). This ratio would fall below unity if households incur higher

[71] Day-of-the-year effects capture holidays that are tied to specific days of the year (e.g., January 6^{th} is the day of the Three Wise Men). Day-of-the-week fixed effects capture the fact that different days of the week have different hassle costs (e.g., taxpayers can remit taxes during weekends by paying at convenience stores).

[72] While this is a very high rate of discounting, it is also consistent with the very high interest rates charged by payday lenders in the US.

[73] Following Train (2009), consumer surplus is given by $\ln\left[\sum_{j=0,1}\exp\{v_j(t=0,y|\{\text{Tax}(t,y)\}_t,\text{Interest}_y)\}\right]\times\frac{1}{\theta_2}$. We evaluate welfare on the first day of the year ($t=0$) given a daily profile of discounts $\{\text{Tax}(t,y)\}_t$. We divide the log-sum by the coefficient on interest income to interpret this quantity as a monetary measure of consumer surplus. Finally, we evaluate this expression setting Interest_y to the average value in our sample.

Table H.1: Dynamic Model of Payment Timing

Panel A. Model Estimates

	(1) Tax Coefficient θ_1	(2) Interest Coefficient θ_2	(3) Discount Factor β
Estimate	0.936***	0.096	0.924***
	(0.180)	(0.064)	(0.041)

Panel B. Welfare Estimates

(1) Year	(2) Consumer Surplus From Discounts	(3) Discount Deadlines	(4) Fiscal Cost	(5) Relative Value of Discounts
2009	10.35	Jan 31, Feb 28	7.80	1.33
2010	3.92	Jan 31	4.95	0.79
2011	3.26	Jan 10, 17, 31	6.53	0.50
2012	5.23	Jan 17, 31	6.84	0.77
2013	7.76	Jan 31, Feb 28	6.95	1.12

Notes: Panel A reports the parameter estimates from the dynamic discrete choice model presented in this Appendix. The parameters were estimated by solving Equation H.1 via non-linear least squares. The model parameters reflect the role of liquidity constraints through the low discount rate and the relatively high utility value of interest income. The underlying data on payment probabilities are weighted to reflect tax collections per day. Panel B uses the model estimates to compute the consumer surplus from discounts (as a percentage of tax payment). Discounts have a larger effect on consumer surplus when taxpayers face longer deadlines and when the discounts are more generous. In these cases, the value of the discounts is greater than the fiscal cost (also as a percentage of tax payment), showing that the government can lower the welfare cost of property taxes by providing liquidity to taxpayers.

hassle costs to obtain the discount. Conversely, this ratio would be greater than one if the value of liquidity exceeds both the additional hassle costs and the fiscal cost to the government. Column (5) shows that in years where households have a longer time to obtain the discounts, the value of the discount to households exceeds the fiscal cost of the discount. That is, when deadlines provide meaningful increases in liquidity, the welfare gains experienced by households exceed the fiscal costs to the government.

The results in this section provide further evidence that payments for property taxes interact with liquidity constraints. This result is evident both in reduced-form patterns that showcase the sensitivity of the timing of payments to discounts and in the estimates from the dynamic discrete choice model. Because we find a value for liquidity even among the households that are willing to pay in full, relaxing liquidity constraints for a broader set of households can lower the welfare cost of paying property taxes and improve the design of the property tax system.

Figure H.1: Taxpayer Response to Super-Early-Bird and Early-Bird Discounts

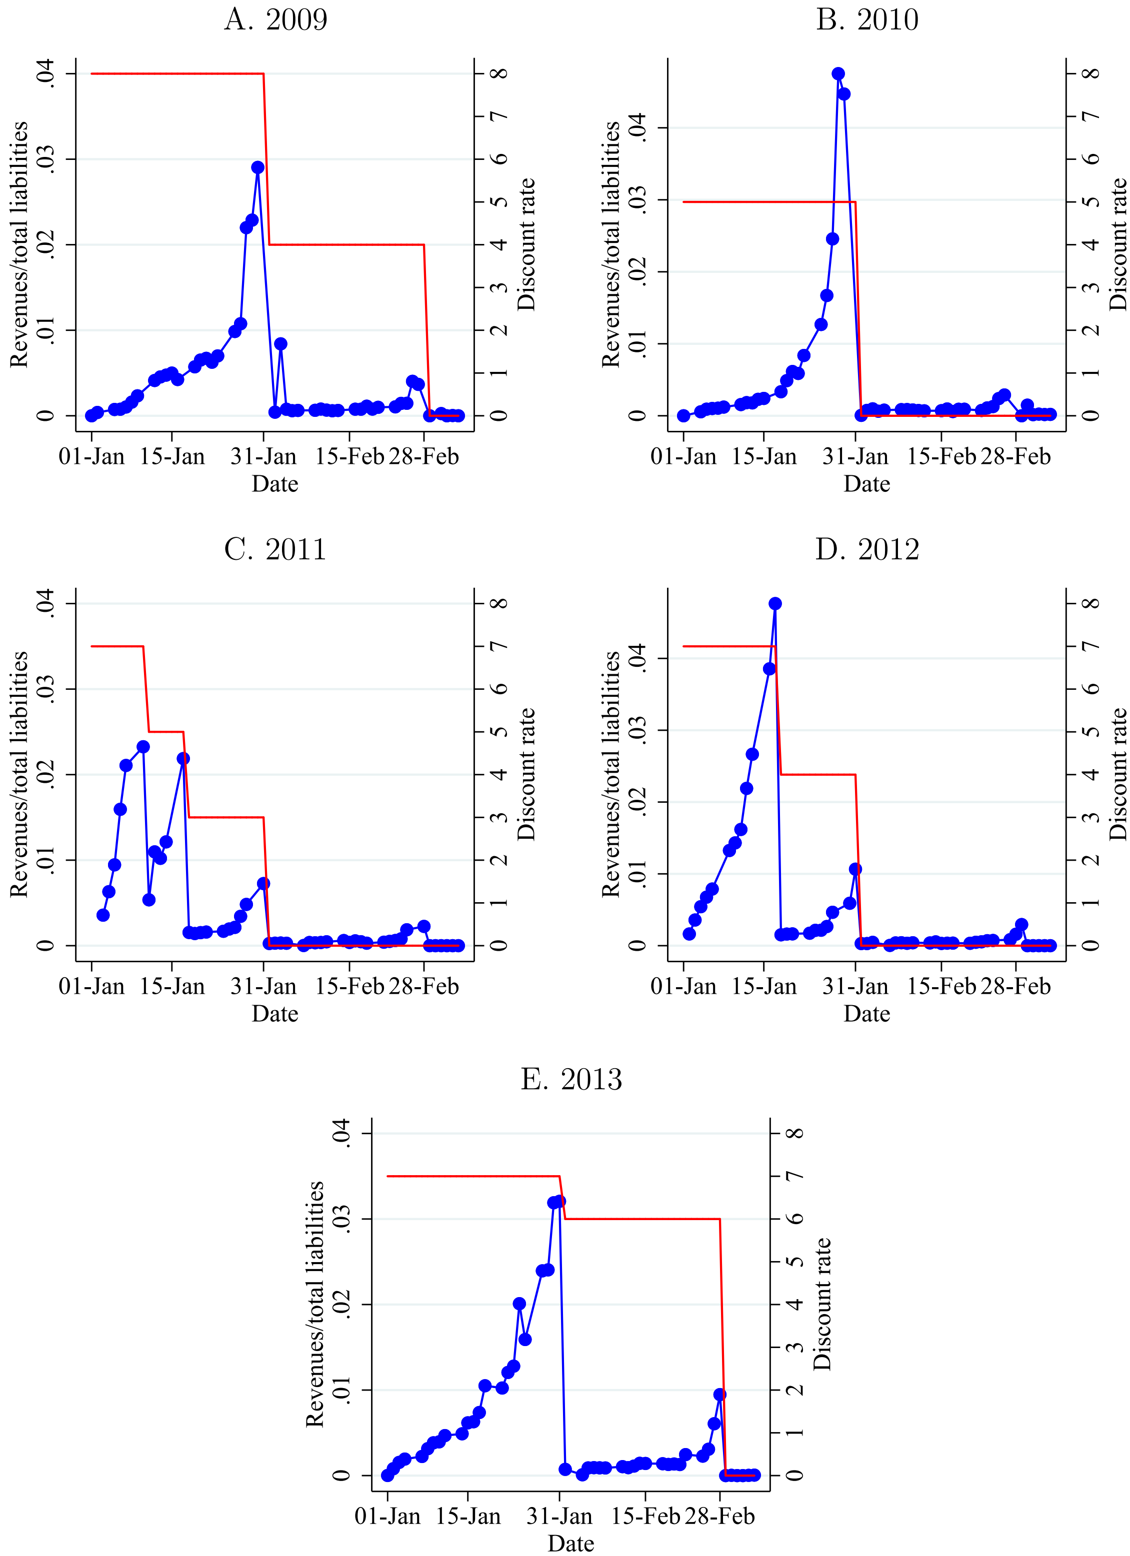

Notes: This figure plots the early-bird discount rates and the timing of property tax payments, as discussed in this appendix. The red line in each panel represents the annual super-early-bird and early-bird discount rates. The blue dots represent the revenues of the government each day, expressed as a fraction of the total yearly liabilities in Mexico City.

Figure H.2: Estimates from Dynamic Model of Payment Timing

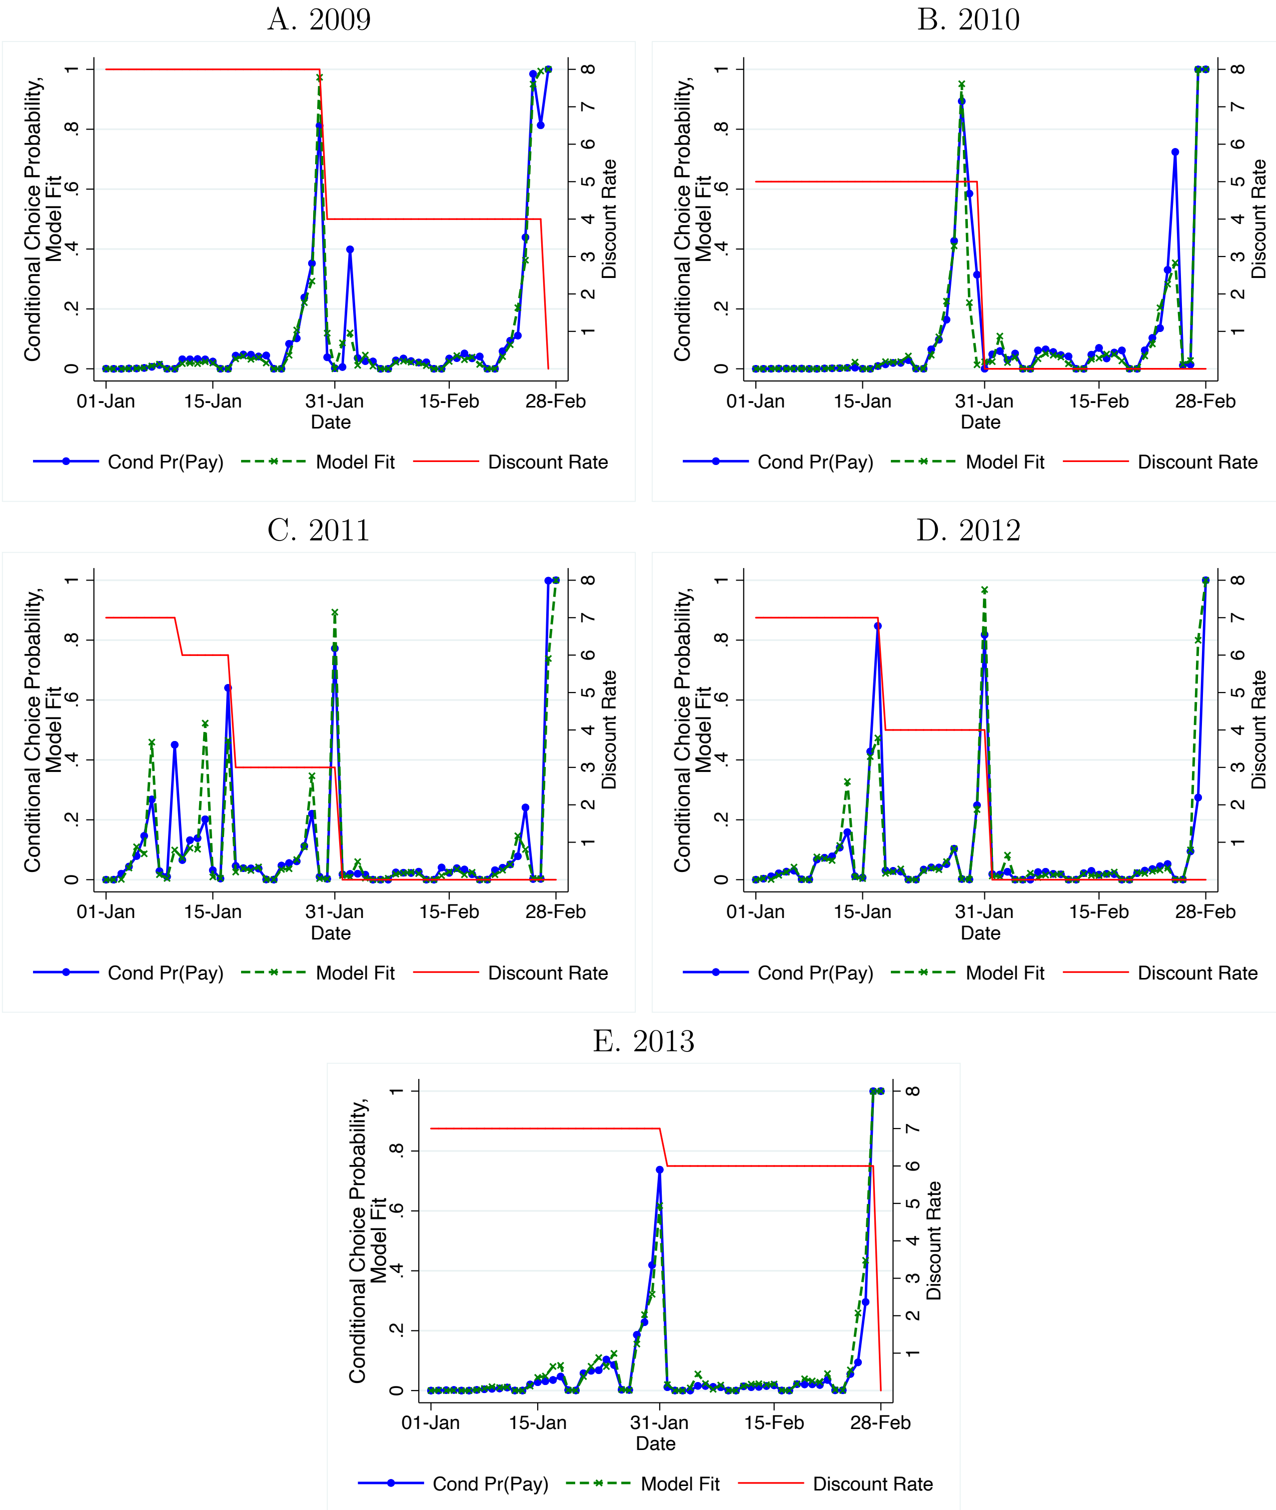

Notes: This figure displays estimates from the dynamic discrete choice model of payment timing presented in this appendix. The red lines display $P_0(t, y)$: the probability of paying taxes on any given day conditional on not having paid. These probabilities are weighted to reflect tax collections per day. The blue lines display the corresponding model fit from the dynamic discrete choice model. The black lines display tax obligations net of discounts.

I Instrumental Variable Appendix

To study the impact of taxes on consumption, we estimate the following specification:

$$\ln C_{it} = \alpha + \beta_1 \ln P_{it} + \delta_1 \ln I_{it} + \kappa_1 A_{it} + \gamma_1 X_{it} + \varepsilon_{it}, \tag{I.1}$$

where P_{it} is the property tax payment of household i in year t, C_{it} is household per capita consumption, I_{it} is household per capita income, A_{it} is a dummy indicating whether the household has access to a credit card, and the control vector X_{it} contains year dummies and *delegación* (i.e., city district) fixed effects. In this equation, β_1 measures the consumption elasticity with respect to property tax payments. If paying property taxes leads households to decrease their consumption, we would expect to find that $\beta_1 < 0$.

Previous work has estimated regressions similar to Equation I.1 to study the impacts of permanent and transitory income shocks and unemployment insurance payments (e.g., Gruber, 1997; Kroft and Notowidigdo, 2016) on consumption. An important concern with Equation I.1 is that property tax payments are likely correlated with other factors that influence consumption. For instance, households facing unmeasured income shocks may decide to skip a property tax payment and may also decrease everyday consumption. These kind of shocks would lead to upwardly biased estimates of β_1.[74]

To provide unbiased estimates of β_1, we use the tax increases that we analyze in Section 4 to isolate variation in property tax payments that is unrelated to other drivers of consumption. Consider the following first-stage equation:

$$\ln P_{it} = \pi_1 + \pi_2 Z_{it} + \delta_2 \ln I_{it} + \kappa_2 A_{it} + \gamma_2 X_{it} + \epsilon_{it}, \tag{I.2}$$

where the instrument Z_{it} is a dummy indicating whether household i is in the treated cadastral value band in year t. As our RD and DiD estimations show, tax rate changes resulting from abatement removals have significant effects on tax payments, suggesting that Z_{it} would be a relevant and statistically strong instrument. The exclusion restriction is that these tax rate changes impact consumption through property tax payments only. Our results in Section 4 show that tax changes lead to sharp variation between treated and untreated households and that these households have parallel trends on a number of outcomes prior to the tax increases. These results suggest that the tax rate variation driven by the abatement removals in Z_{it} is likely unrelated to other determinants of household consumption.[75]

Our instrumental variables strategy requires three key variables: consumption, C_{it}; tax payments, P_{it}; and a dummy indicating whether the household is in the value band treated with an abatement removal, Z_{it}. Unfortunately, while the ENIGH data measure consumption and property tax payments, they do not record cadastral values. We overcome this issue by using a split-sample instrumental variables strategy (e.g., Angrist and Krueger, 1992; Card and McCall, 1996) that combines information from our administrative tax data and our household survey data. Specifically, we follow a three-step procedure to construct the probability that household i is in the treated cadastral value band in year t. First, we use the administrative tax data to construct coarse land area $\times$ construction area bins and calculate within each bin and year the share of households that are treated.

[74] Indeed, OLS estimations of Equation I.1 yield positive estimates of β_1.

[75] Unlike in the US, property tax revenue in Mexico is not used to fund neighborhood schools or other amenities. Any benefits from additional tax revenue would be spread among the more than 20 million residents of Mexico City.

Second, we use the ENIGH data on the number of rooms in the home to estimate each home's probability to be in a given land area×construction area bin, using a multinomial logit. Finally, we utilize the shares calculated in the first step and the probabilities in the second step to estimate the probability that the home is in a treated cadastral value band in a specific year. To account for the uncertainty in our proxy for Z_{it}, we bootstrap this procedure to calculate standard errors.

Table I.1 presents our IV results, estimated using the sample of surveyed households in Mexico City that report paying property taxes in 2010 and 2012. Column (1) reports the first-stage regression estimates. We find a strongly significant parameter estimate $\hat{\pi}_2 = 0.509$ ($SE = 0.063$), which implies an F-statistic over 65. Column (2) includes an interaction of the instrument with the log of per-capita income, where we normalize log income relative to the cross-sectional mean. In this specification, the mean effect of tax increases on tax payment is 0.47; we also estimate larger effects of high-income households than for low income households. The model parameter π_2 has an intuitive interpretation: the effect of belonging to the treatment group on the tax rate in percentage terms. It is, therefore, reassuring that our estimate of π_2 has a similar magnitude to the tax hike of 49 percent in 2012, experienced by properties in value band G.

Columns (3) and (4) report reduced-form regressions of log per-capita consumption on the predicted likelihood that the household faced a tax increase. Both columns report small consumption effects of facing a tax increase of -1.7 and -2.9 percent, on average. Column (4) shows that the consumption of lower income households is more responsive to property tax increases: we estimate that facing a tax increase lowers consumption by 7.1 percent for households with log per-capita income one standard deviation below the mean.

Columns (5) and (6) report second-stage estimates. Column (5) reports a negative but statistically insignificant effects of tax payments on consumption on average of -3.3 percent. Because we expect tax payments to have larger impacts for lower-income households, we augment Equation I.1 by including an interaction with per capita household income.[76] Column (6) shows the estimates produced when we interact tax payment with household income and helps us understand the magnitude and heterogeneity in the effects of property taxes on consumption. We estimate that doubling property taxes leads to a decline in consumption of 9.3 percent for households whose income is one standard deviation below the cross-sectional mean.[77] To gauge the magnitude of these effects relative to other policies, note that Gruber (1997); Kroft and Notowidigdo (2016) estimate that losing a job without unemployment insurance (UI) would lead to a 23 percent drop in consumption in the US and that increasing the UI replacement rate by 10 percentage points would reduce this drop by 2.7 percent. Moreover, Chetty (2004) notes that income shocks can lead to larger changes in consumption when households face consumption commitments, such as with housing.

These estimates show that policy makers are right to be concerned about the potential of property taxes to exacerbate liquidity constraints. In Section 7, we quantify the extent to which liquidity constraints increase the welfare cost of taxes, using a range of estimates that includes our IV estimates.

[76]We also include an interaction between the instrument and per capita household income in the first stage in Column (2) and the reduced form in Column (4).

[77]While doubling property taxes might seem extreme, recall that the sanctions treatment in our field experiment in Section 5 led to a doubling of tax payments relative to payments made by the control group. To put these estimates in the perspective of our tax changes, recall that property taxes increased by 20 percent in 2010, by 27 percent in 2011, and by 47 percent in 2012.

Table I.1: The Effect of Property Taxes on Per Capita Consumption – IV Results

	First-Stage		Reduced-Form		Second-Stage	
	(1)	(2)	(3)	(4)	(5)	(6)
Z : (A)	0.509***	0.471***	-0.017	-0.029		
	(0.063)	(0.059)	(0.024)	(0.024)		
	[0.390, 0.647]	[0.362, 0.601]	[-0.066, 0.027]	[-0.079, 0.016]		
Z × log(pc income): (B)		0.153***		0.049**		
		(0.057)		(0.023)		
		[0.036, 0.265]		[0.004, 0.096]		
log(Pay): (C)					-0.033	-0.046
					(0.049)	(0.050)
					[-0.132, 0.057]	[-0.144, 0.044]
log(Pay) × log(pc income): (D)						0.054**
						(0.023)
						[0.007, 0.099]
log(pc income)	0.480***	0.405***	0.763***	0.739***	0.779***	0.471***
	(0.030)	(0.042)	(0.017)	(0.022)	(0.032)	(0.131)
	[0.421, 0.538]	[0.323, 0.495]	[0.727, 0.797]	[0.694, 0.783]	[0.718, 0.841]	[0.207, 0.723]
Lack of credit	-0.421***	-0.405***	-0.270***	-0.265***	-0.284***	-0.266***
	(0.059)	(0.059)	(0.030)	(0.030)	(0.037)	(0.037)
	[-0.542, -0.307]	[-0.528, -0.289]	[-0.329, -0.212]	[-0.323, -0.207]	[-0.355, -0.212]	[-0.335, -0.192]
F-statistic (excluded instruments)	65.71	33.18				
P-value	0.000	0.000				
Post-estimation:						
Marginal effect of Z for high income: (A)+σ×(B)		0.604***		0.013		
		(0.082)		(0.030)		
		[0.471, 0.746]		[-0.035, 0.062]		
Marginal effect of Z for low income: (A)-σ×(B)		0.338***		-0.071**		
		(0.072)		(0.033)		
		[0.224, 0.458]		[-0.128, -0.018]		
Marginal effect of log(Pay) for high income: (C)+σ×(D)						0.001
						(0.049)
						[-0.080, 0.082]
Marginal effect of log(Pay) for low income: (C)-σ×(D)						-0.093
						(0.059)
						[-0.195, -0.003]

Notes: This table reports the results from the main IV strategy. The instrument Z is constructed as the predicted probability that a household's property is part of the treated cadastral value band in each year. The outcome of the first-stage equation is the log property tax payment, while the outcome variable for the reduced-form and second-stage regressions in the log of per capita household consumption. All regressions include *delegación* fixed effects and year dummies. In post estimation, we use the standard deviation of log income σ to calculate the marginal effects of the variables Z and log(Pay) at one standard deviation above and below mean income. Bootstrapped standard errors based on 1,000 replications are in parentheses. Percentile confidence intervals at the 95% level are in squared brackets. N=2,609. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

J Revenue-Maximizing Tax Rates

In this appendix, we assume that the government's single aim is to maximize tax revenue, and we use the different tax changes analyzed in Section 4 to estimate the revenue-maximizing tax rate. Our empirical results from Section 4 show that larger tax increases imply smaller revenue elasticities and that tax increases have a significant effect on the fraction of delinquent taxpayers. These results raise the possibility that further tax increases may have small or even null effects on revenue. To evaluate this possibility, we use our empirical estimates to characterize the degree to which current tax rates are close to the revenue-maximizing tax rate.

Building on the corporate tax literature (Clausing, 2007; Devereux, 2007; Kawano and Slemrod, 2015; Suárez Serrato and Zidar, 2018), we estimate a quadratic relation between taxes and revenue:

$$\ln Rev_t = \beta_1 \tau_t + \beta_2 (\tau_t)^2.$$

Revenue is a concave parabola of taxes when $\beta_1 > 0$ and $\beta_2 < 0$. Intuitively, $\beta_1 > 0$ implies that introducing a small tax raises revenue, and $\beta_2 < 0$ implies that the marginal impact on revenue $(\beta_1 + 2\beta_2\tau_t)$ is smaller for higher tax rates. At the revenue-maximizing rate, the marginal impact of a tax increase is zero, which implies that revenue is maximized by $\tau^* = \frac{-\beta_1}{2\beta_2}$. Key empirical questions are then whether $\beta_2 < 0$ and whether large values of β_2 imply small values of τ^*.

To connect this framework to our results, write the effect of a tax change on revenue as:

$$\underbrace{\frac{\Delta \ln Rev_t}{\Delta \tau_t}}_{\text{Semi-Elasticity: } \eta_t} = \beta_1 + 2\beta_2\tau_t.$$

This expression implies that we can estimate β_1 and β_2 from multiple estimates of the revenue semi-elasticity at different values of τ_t.[78] Let $\hat{\boldsymbol{\eta}} = [\hat{\eta}_{2010}, \hat{\eta}_{2011}, \hat{\eta}_{2012}]'$ be the vector of semi-elasticities from the three tax changes, and define the matrix $\boldsymbol{W} = [\mathbf{1}_t, 2\boldsymbol{\tau_t}]$. Using a simple application of classical minimum distance (CMD), we estimate β_1 and β_2 as a linear combination of the semi-elasticities: $[\hat{\beta}_1, \hat{\beta}_2]' = (\boldsymbol{W}'\boldsymbol{W})^{-1}(\boldsymbol{W}'\hat{\boldsymbol{\eta}})$.[79] We then use these estimates to test whether $\beta_2 < 0$ and to study the implied revenue-maximizing rates τ^*.

The elasticity and semi-elasticity estimates resulting from our estimations discussed Section 4 are displayed in Table J.1. Applying the method discussed above, we obtain estimates of $\beta_2 = 1.64(SE = 2.34)$ when using the regression discontinuity estimates of $\hat{\boldsymbol{\eta}}$ and $\beta_2 = 3.29(SE = 0.74)$ when using the difference-in-difference estimates (Table J.2).[80] The result that both estimates of β_2 are positive implies that current property tax rates are significantly below the revenue-maximizing rate. This result is driven by the fact that our semi-elasticity estimates are not decreasing in τ_t. As Tables J.1 shows, we estimate larger semi-elasticities for larger values of τ_t.

[78] One potential concern is that the three tax changes estimate effects from households in different parts of the home value distribution, leading to different elasticities. This is not the case. The thresholds for the three reforms were approximately 2.275, 1.95, and 1.625 million MXN. The three elasticities are based on comparable properties that are 325 thousand MXN—about 16 thousand USD—apart.

[79] See Chamberlain (1984) for a guide to CMD and Suárez Serrato and Zidar (2016) for a recent application.

[80] While we can reject the null hypothesis that $\beta_2 < 0$ with a p-value < 0.001 when we use the DiD estimates, we cannot reject this hypothesis when we use the RD estimates. See Table J.2 for details. Importantly, this result is not driven by a lack of statistical precision. Estimates of β_1 and β_2 yield precisely estimated revenue semi-elasticities at the average tax rate of $0.010(SE = 0.002, t - stat = 5.29)$ for the regression discontinuity case and $0.012(SE = 0.001, t - stat = 19.78)$ for the difference-in-difference case.

While the point estimates for β_2 are positive, we also consider how uncertainty in these estimates affects our policy analysis. To explore the role of uncertainty, we simulate 10,000 values of β_1 and β_2 based on their joint distribution and characterize the resulting distribution of τ^*. This exercise shows that 80 percent of the time, the revenue-maximizing rate is greater than 159 basis points. We also find that 90 percent of the simulated values yield estimates of τ^* above 73 basis points and that only 5 percent of the estimates are below 61 basis points.[81] Given that the highest tax rate in the three reforms was 50 basis points, these results show that the government can raise the property tax rate by 20–50 percent with a very limited risk of going beyond the revenue-maximizing rate.

The policy takeaway from this analysis is that rigorous empirical evidence from recent tax increases shows that current tax rates are significantly below the revenue-maximizing tax rate.[82]

While the government may be able to collect additional tax revenue by taxing property at rates below τ^*, the welfare costs from increasing tax rates or tightening enforcement may exceed the value that taxpayers obtain from using the additional revenue to provide public goods. For this reason, it is possible that the optimal tax rate may fall significantly below the revenue-maximizing rate. In addition, the revenue-maximizing analysis does not provide any guidance as to whether the government should rely on tax rate increases or enforcement actions to collect revenue. In Section 7 of the paper, we implement the welfare-maximizing model from Section 1 that incorporates these important insights.

Table J.1: The Effect of Tax Rates on Tax Payment – Semi-Elasticity Estimates

	A. Regression Discontinuity			B. Difference-in-Differences		
	Payment Amount	Payment in Full	Compliance Share	Payment Amount	Payment in Full	Compliance Share
	(1)	(2)	(3)	(4)	(5)	(6)
I. Estimates for the 2009-2010 Treatment						
Implied Elasticity	.55	-.822	-.368	.697	-.617	-.244
	(.226)	(.318)	(.234)	(.101)	(.104)	(.067)
Implied Semi-Elasticity	.011	-.016	-.007	.014	-.012	-.005
	(.005)	(.006)	(.005)	(.002)	(.002)	(.001)
II. Estimates for the 2010-2011 Treatment						
Implied Elasticity	.475	-.717	-.407	.728	-.642	-.300
	(.108)	(.170)	(.107)	(.085)	(.064)	(.043)
Implied Semi-Elasticity	.01	-.015	-.009	.015	-.014	-.006
	(.002)	(.004)	(.002)	(.022)	(.001)	(.001)
III. Estimates for the 2011-2012 Treatment						
Implied Elasticity	.314	-.658	-.316	.489	-.747	-.455
	(.059)	(.083)	(.060)	(.073)	(.035)	(.025)
Implied Semi-Elasticity	.008	-.016	-.008	.012	-.018	-.011
	(.001)	(.002)	(.001)	(.002)	(.001)	(-001)

Notes: This table reports the semi-elasticity estimates implied by our regression discontinuity and difference-in-difference estimations, for which the estimates are reported in Table 2.

[81] These simulations are based on our regression discontinuity results. The difference-in-difference estimates imply larger revenue-maximizing rates in all cases.

[82] Our findings contrast with the results of Haughwout et al. (2004), who find that in three of four major US cities, property tax rates are close to the peak of the Laffer curve.

Table J.2: Classical Minimum Distance Estimates of Revenue-Maximizing Tax Rates

	(1) RD Estimates	(2) DiD Estimates
β_1	-0.005	-0.019***
	(0.020)	(0.007)
β_2	1.640	3.293***
	(2.347)	(0.738)
Mean Semi-Elasticity	0.010***	0.012***
	(0.002)	(0.001)
20th Percentile of τ^*	158.968	250.000
10th Percentile of τ^*	73.431	250.000
5th Percentile of τ^*	61.521	250.000

Standard errors in parentheses.

* $p < .1$, ** $p < .05$, *** $p < .01$.

Notes: This table reports results for the classical minimum distance estimates of revenue-maximizing tax rates. Column (1) uses estimates from the RD estimation to estimate β_1 and β_2, while Column (2) uses estimates from the DiD estimation. In both cases, we find that $\beta_2 > 0$, indicating that current tax rates are substantially below the revenue-maximizing rates. For both cases, we report the mean semi-elasticity by evaluating the revenue curve at the average tax rate. Finally, the last panel reports estimates of revenue-maximizing tax rates τ^* from 10,000 simulated values of β_1 and β_2. We compute τ^* in each simulation and report the percentiles of this distribution. We top-code τ^* at 250 basis points when $\beta_2 > 0$.

K Spillover Effects

To measure the welfare effects of different tax administration instruments, including the tax rate and enforcement, it may be important to account for their spillover effects on tax compliance. Recent literature suggests that policy changes can have spillover effects on tax compliance through social motives, including reputational impacts (Besley et al., 2023) and tax morale or fairness perceptions (Best et al., 2020). To address this possibility, we begin by providing an econometric framework that helps us identify and estimate spillover effects in the data. We then empirically test whether the tax rate changes studied in Section 4 and the enforcement letters described in Section 5 had substantial spillover effects on tax compliance. In our setting, we do not find economically important spillover effects of taxes or enforcement letters. For this reason, our baseline estimates of the welfare effects of these policies focus on direct impacts.

K.1 Econometric Framework

Consider a model where taxpayers may be treated themselves and where their neighbors may also be treated. Define the treatment dummy for household i in period t as $T_{i,t}$. Similarly, let NT_{it} denote a dummy for neighbor treatment. The pair (T_i, NT_i) defines four potential outcomes: (1,1) treated with neighbors treated, (0,1) untreated with neighbors treated, (1,0) treated with neighbors untreated, and (0,0) untreated with neighbors untreated. Let $Y_{i,t}^{(T_i,NT_i)}$ denote these potential outcomes for a treated household i in period t, where $t = 1$ corresponds to the period before the treatment and $t = 2$ denote the post period. The observed outcome is

$$Y_{i,t} = T_{i,t}(Y_{i,t}^{(1,0)} + NT_{i,t}(Y_{i,t}^{(1,1)} - Y_{i,t}^{(1,0)})) + (1 - T_{i,t})(Y_{i,t}^{(0,0)} + NT_{i,t}(Y_{i,t}^{(0,1)} - Y_{i,t}^{(0,0)})).$$

The treatment can be a tax hike or an enforcement letter. In both cases we are interested in identifying the Indirect Treatment Effect on the Non-Treated (ITEN), defined as follows:

$$ITEN = E[Y_{i,2}^{(0,1)} - Y_{i,2}^{(0,0)}],$$

and the Indirect Treatment Effect on the Treated (ITET), defined as:

$$ITET = E[Y_{i,2}^{(1,1)} - Y_{i,2}^{(1,0)}].$$

While these parameters need not be the same, both could be the equal to each other if the indirect effect of neighbor treatment is orthogonal to the taxpayer's treatment status (i.e., $Y_{i,2}^{(T_{i,2},1)} - Y_{i,2}^{(T_{i,2},0)} \perp T_{i,2}$). In other words, the ITEN and ITET parameters are the same under the assumption that the indirect effect of neighbor compliance does not vary with the taxpayer's treatment status. We make use of this assumption for estimation purposes.

K.2 Differences-in-Differences Estimation

We begin by explaining the differences-in-differences strategy we use to identify the ITEN parameter. When utilizing differences-in-differences, one can estimate the ITEN parameters for individuals who were not treated and whose neighbors were treated. Having this in mind, consider now the before-after difference for the group of untreated taxpayers with treated neighbors:

$$\begin{aligned}\Delta_1|\{T_{i,1}=0,T_{i,2}=0,NT_{i,1}=0,NT_{i,2}=1\} &= E[Y_{i,2}-Y_{i,1}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=1,NT_{i,1}=0]\\ &= E[Y_{i,2}^{(0,1)}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=1,NT_{i,1}=0]\\ &- E[Y_{i,1}^{(0,0)}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=1,NT_{i,1}=0]\\ &= E[Y_{i,2}^{(0,1)}-Y_{i,2}^{(0,0)}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=1,NT_{i,1}=0]\\ &+ E[Y_{i,2}^{(0,0)}-Y_{i,1}^{(0,0)}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=1,NT_{i,1}=0].\end{aligned} \tag{K.1}$$

Consider now the before-after difference for the group of untreated taxpayers with untreated neighbors:

$$\begin{aligned}\Delta_0|\{T_{i,1}=0,T_{i,2}=0,NT_{i,1}=0,NT_{i,2}=0\} &= E[Y_{i,2}-Y_{i,1}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=0,NT_{i,1}=0]\\ &= E[Y_{i,2}^{(0,0)}-Y_{i,1}^{(0,0)}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=0,NT_{i,1}=0].\end{aligned} \tag{K.2}$$

Subtracting Equation (K.2) from Equation (K.1), we find that the differences-in-differences estimator recovers the ITEN for the group of taxpayers who were not treated but had treated neighbors:

$$\Delta_1|\{T_{i,1}=0,T_{i,2}=0,NT_{i,1}=0,NT_{i,2}=1\}-\Delta_0|\{T_{i,1}=0,T_{i,2}=0,NT_{i,1}=0,NT_{i,2}=0\}=ITEN,$$

provided that the following parallel trends assumption holds:

$$\begin{aligned}&E[Y_{i,2}^{(0,0)}-Y_{i,1}^{(0,0)}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=1,NT_{i,1}=0]-\\ &E[Y_{i,2}^{(0,0)}-Y_{i,1}^{(0,0)}|T_{i,1}=0,T_{i,2}=0,NT_{i,2}=0,NT_{i,1}=0]=0.\end{aligned}$$

A similar reasoning can be applied to derivation the following identifying equation for the ITET:

$$\Delta_1|\{T_{i,1}=0,T_{i,2}=1,NT_{i,1}=0,NT_{i,2}=1\}-\Delta_0|\{T_{i,1}=0,T_{i,2}=1,NT_{i,1}=0,NT_{i,2}=0\}=ITET.$$

Under the assumption that the indirect treatment effect is orthogonal to the taxpayer's treatment status, both differences-in-difference strategies recover the same indirect treatment effect (i.e., ITE=ITEN=ITET). Therefore, for estimation purposes, we use these estimation strategies interchangeably depending on data availability.

K.3 Spillover Effects of Tax Rate Hikes

We estimate the spillover effect of the tax rate hikes described in Section 4 for the group of taxpayers who were not directly treated with these hikes. That is, we estimate the ITEN by comparing the average tax compliance for the group of untreated taxpayers in zipcodes where a high density of neighbors is treated with the average compliance for the group of taxpayers in zipcodes where a zero density of neighbors is treated, before and after treatment. Specifically, we estimate the parameter β in the following regression model via OLS on a sample containing only indirectly-treated and non-treated taxpayers:

$$Y_{i,t}=\alpha+NT_{i,t}\beta+\gamma_i+\delta_t+\varepsilon_{i,t}, \tag{K.3}$$

where γ_i are property fixed effects, which control for time-invariant property characteristics that correlate with tax compliance, and δ_t are time dummies, which control for shocks to tax compliance that are common to all taxpayers. Standard errors are robust to heteroskedasticity of unknown form and are clustered at the zipcode level, as we define neighbor exposure at that geographical level.

We select the sample for estimating Equation K.3 for each tax hike, following a two-step procedure each year. First, we calculate the share of treated properties in all zip codes, and tag non-treated properties in zip codes with a sufficiently high share of treated properties. We show below that our results are robust to different threshold definitions. These tagged properties are the indirectly treated ones.[83] Second, we match each property in the indirectly treated group with a property of similar cadastral value, land area, construction area, and construction year in the sample of untreated properties in zip codes without treated properties, using a propensity score matching algorithm. Our strategy is therefore similar to the matching and differences-in-differences approach (MDID) described in Blundell and Costas Dias (2009).

The identifying assumption is that, in the absence of neighbor treatment, compliance would change differentially among properties that were indirectly treated compared to properties that were not indirectly treated. While this assumption is not testable, Figure K.1 provides graphical evidence that the gap in compliance between the indirectly-treated and the non-treated group was statistically indistinguishable from zero before treatment. The figures also indicate that there is no detectable evidence for a spillover effect after treatment and that we can rule out even very small spillovers.

Table K.1 summarizes the difference-in-difference estimates. Each column corresponds to a different treatment group definition, obtained by varying the threshold in the distribution of direct treatment density across zip codes. Regardless of what cutoff we use to tag the indirectly treated properties and which reform we use, the spillover estimates in all specifications are small and standard errors indicate that they are precisely estimated. Since, in our setting tax increases do not have economically important spillover effects on tax compliance, we do not include spillover effects in our policy analysis in Section 7.

Figure K.1: Difference in Tax Compliance Between Indirectly Treated and Untreated Taxpayers

Panel A. 2010 | Panel B. 2011 | Panel C. 2012

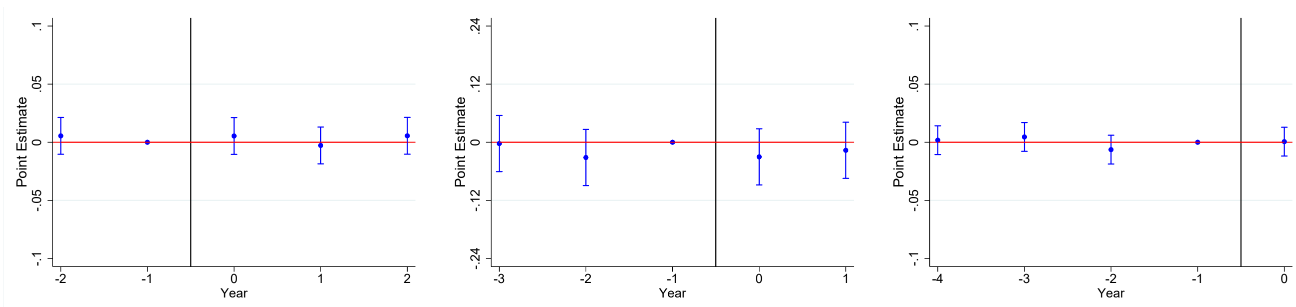

Notes: These figures present evidence of parallel trends in tax compliance between indirectly treated and matched untreated taxpayers prior to the 2010, 2011, and 2012 tax hikes, as discussed in Section K.3. For each year, the figure displays the point estimates and 95% confidence intervals of year×treatment group interactions in a regression of tax compliance on taxpayer fixed effects and time dummies, where the interaction for the year prior to indirect treatment is excluded. The estimating sample each year includes only indirectly treated properties and their matched pairs and excludes all properties treated directly with the tax hike. Standard errors are robust to heteroskedasticity of unknown form and are clustered at the zipcode level.

[83] For each reform, we also drop properties in zip codes that are indirectly treated in a later reform.

Table K.1: Indirect Effect of Tax Hikes on Tax Compliance

	Mean	75th Percentile	80th Percentile
	(1)	(2)	(3)
Panel I: Estimates for the 2010 Treatment			
MDID ATT	0.005	-0.008	0.009
	(0.009)	(0.015)	(0.007)
R^2	0.648	0.650	0.597
N	82,240	746,470	389,420
Panel II: Estimates for the 2011 Treatment			
MDID ATT	-0.017	0.011	-0.010
	(0.025)	(0.008)	(0.014)
R^2	0.412	0.575	0.417
N	171,030	203,480	326,150
Panel III: Estimates for the 2012 Treatment			
MDID ATT	0.005	-0.007	0.007
	(0.007)	(0.015)	(0.020)
R^2	0.594	0.568	0.555
N	799,230	287,090	198,440

Notes: This table shows the indirect effect of tax rates increases on an indicator for compliance, as mentioned in Section K.3. Indirectly treated properties are untreated properties in zip codes with a sufficiently high treatment density. We also drop directly treated properties and properties in zip codes that are indirectly treated in later years. The column titles indicate the cutoff in treatment density we use to tag the indirectly treated properties. We then use nearest-neighbor propensity score matching to find control properties in zipcodes with zero direct treatment density for the indirectly treated properties in high density zipcodes. The propensity score Pr[Treat=1]=$F(X_i'\beta)$, where X_i' is a vector of covariates including cadastral value, land area, construction area, construction year and *delegación*. The controls are then matched to the treated and the reweighted sample is used to compute the ITEN using DID. High density is defined as per column labels. * $p < .1$, ** $p < .05$, *** $p < .01$.

K.4 Spillover Effects of Enforcement Letters

We now estimate the spillover effect of the enforcement letters studied in Section 5 for the group of taxpayers who were directly treated.[84] That is, we estimate the ITET of enforcement letters by comparing the effects of the two types of enforcement letters (i.e., sanctions and public goods) on the payment behavior of the group of treated taxpayers residing in high-density zipcodes with the corresponding effects on the payment behavior of the group of treated taxpayers residing in low-density zipcodes 40 days after the letters were sent. High and low density refers to the density of properties that received an enforcement letter, and we divide the sample by the mean of the distribution across zip codes. We estimate the parameters β_3 and β_4 of the following regression model via OLS:

$$Y_i = \alpha + \beta_1 T1_i + \beta_2 T2_i + \beta_3 (T1_i \times NT_i) + \beta_4 (T2_i \times NT_i) + \beta_5 NT_i + \varepsilon_i, \quad \text{(K.4)}$$

where NT_i is a dummy indicating whether property i is in a zipcode where a high density of neighbors is treated with either of the two letters, $T1_i$ is an indicator for the enforcement letter, and $T2_i$ is an indicator for the public goods letter. Standard errors are robust to heteroskedasticity of unknown form and are clustered at the zipcode level, since indirect treatment occurs at this geographical level.

To account for the pre-treatment differential trends in payment rates across groups in our esti-

[84] Data limitations prevent us from estimating the spillover effect of the enforcement letters on the non-treated. We follow previous literature on spillover effects and RCTs in development economics. Athey and Imbens (2017) provide a review of other RCTs involving interactions.

mation, we estimate the following regression on the pre-intervention data:

$$Y_{igt} = \mu_g \cdot t + \alpha_i + \lambda_t + \varepsilon_{igt},$$

where t indicates days and g treatment groups. We then subtract the estimated trend $\hat{\mu}_g \cdot t$ from each treatment group. We conduct this corrective procedure separately for properties in high- and low-density zipcodes.

Our identifying assumption is that the density of treated neighbors at the zipcode level is uncorrelated with other unobservable determinants of taxpayer compliance. We build confidence on this assumption by progressively adding controls to the regression, including cadastral value, land and construction area, and construction year, since coefficient stability is generally taken in the literature as evidence for the absence of unobserved correlates with the variables of interest (see Oster (2019) for an ample discussion). Furthermore, we add interactions of these controls with the treatment group dummies in estimation.

Figure K.2 presents graphical evidence of a differential impact of both types of enforcement letters on taxpayer compliance in high-density relative to low-density zipcodes. To formalize this graphical finding, we present our OLS regression estimates of Equation (K.4) in Table K.2. The estimates in column (1) suggest that the indirect treatment has a differential and statistically significant impact of 1.24 percentage points on the compliance of properties located in high-density neighborhoods that were treated with the sanctions letter, whereas the public goods treatment also displays a differential albeit statistically insignificant impact of 0.431 percentage points on the same groups properties. Column (2) shows that we find similar spillover effects when including controls in the regression. Columns (3) and (4) estimate spillover effects on payment amounts. Both columns show economically small and statistically insignificant estimates of spillover effects. Since, in our setting enforcement letters do not have economically important spillover effects on tax compliance, we do not include spillover effects in our policy analysis in Section 7.

Figure K.2: Spillover Effect of the Enforcement Letters on Tax Payment

Panel A. High-Density Zipcodes | Panel B. Low-Density Zipcodes

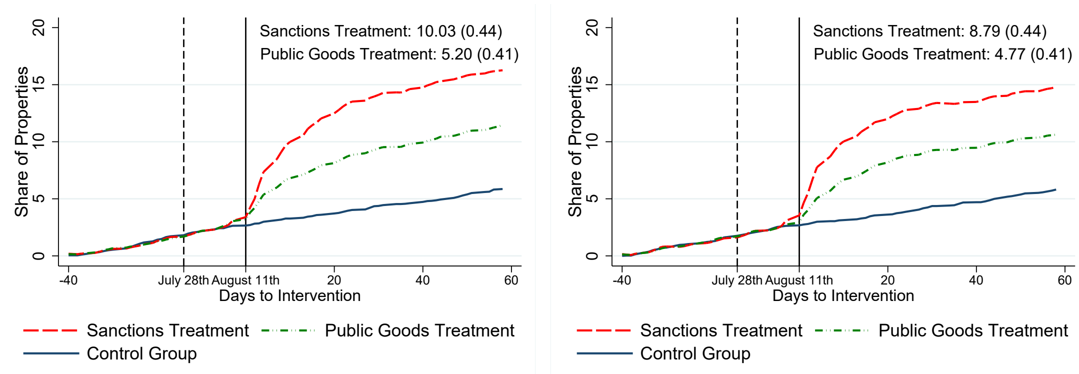

Notes: These figures present evidence of the spillover effects of enforcement letters for both the sanctions and the public goods treatments, as discussed in Section K.4. This figure shows the raw data and the treatment estimates which we obtain from detrended data via a two-step procedure. First, we run the following regression on the pre-intervention data: $Y_{igt} = \mu_g \cdot t + \alpha_i + \lambda_t + \epsilon_{igt}$, where t indicates days and g treatment groups. Second, we subtract the trend $\mu_g \cdot t$ from each treatment group. This is reasonable because the pre-intervention trend is indeed almost perfectly linear and the control group trend continues linearly after the intervention. The outcome variable is a payment dummy, which we construct taking into account only payments made by the taxpayers participating in the experiment described in Section 5 between July and November 2014 against outstanding debt for the period from bimester 4 of 2009 to bimester 3 of 2014. Standard errors are robust to heteroskedasticity of unknown form and are clustered at the zipcode level.

Table K.2: Spillover Effect of the Enforcement Letters on Tax Payment

	A. Any Payment		B. Payment Amount	
	(1)	(2)	(3)	(4)
Sanctions Treatment	8.788	8.553	47.525	41.220
	(0.441)	(0.441)	(3.462)	(3.462)
Public Goods Treatment	4.767	4.556	14.160	9.564
	(0.413)	(0.415)	(2.904)	(2.405)
Sanctions $\cdot HighDensity$	1.240	1.610	-7.439	0.317
	(0.622)	(0.622)	(4.027)	(4.027)
Public Goods $\cdot HighDensity$	0.431	0.763	-2.083	4.431
	(0.582)	(0.607)	(3.304)	(3.460)
R^2	0.009	0.010	0.006	0.028
N	72,593	72,593	72,593	72,593
Density in High Density Zipcodes	0.118			
Density in Low Density Zipcodes	0.050			

Notes: This table shows the indirect effect of enforcement letters on property tax payment, as described in Section K.4. Column (1) presents OLS estimates of the coefficients β_3 and β_4 in Equation (K.4), while Column (2) presents estimates of the same parameters after including controls in the regression. The vector of controls (which are demeaned) includes cadastral value, land and construction area, and construction year, as well as interactions of these variables with the treatment group dummies $T1_i$ and $T2_i$. The outcome variable is a payment dummy, which we detrend before estimation following a two-step procedure. First, we run the following regression on the pre-intervention data: $Y_{igt} = \mu_g \cdot t + \alpha_i + \lambda_t + \epsilon_{igt}$, where t indicates days and g treatment groups. Second, we subtract the trend $\mu_g \cdot t$ from each treatment group. This is reasonable because the pre-intervention trend is indeed almost perfectly linear and the control group trend continues linearly after the intervention. Our sample includes only payments made by the taxpayers participating in the experiment described in Section 5 between July and November 2014 against outstanding debt for the period from bimester 4 of 2009 to bimester 3 of 2014. Standard errors are robust to heteroskedasticity of unknown form and are clustered at the zipcode level. * $p < .1$, ** $p < .05$, *** $p < .01$.

L Quantifying Liquidity Constraints

This appendix estimates the share of liquidity-constrained consumers in Mexico City using household income and consumption data. To draw inferences about the prevalence of liquidity constraints, we follow the literature and (1) estimate the representative household's income process, (2) compare the implied consumption responses to unexpected permanent and transitory income shocks under different theoretical assumptions about the degree to which households are liquidity constrained, and (3) quantify the prevalence of liquidity constraints in the data by comparing household income and data moments with their expected values according to economic theory.

As discussed in the main body of the paper, we cannot directly estimate each individual household's income process, since doing so would require longitudinal data. Our data on household consumption and income is from the ENIGH survey, which is a repeated cross-section. For this reason, we estimate the income process from cohort-level consumption and income moments, as proposed in the seminal article of Blundell and Preston (1998).

Section L.1 presents the stochastic process of income. Sections L.2 and L.3 describe the household's consumption rule under the permanent income hypothesis and for liquidity-constrained households, respectively. Finally, in Section L.4, we show that we can identify the share of liquidity-constrained households from an OLS regression of the first difference of the consumption variance on the first difference of the variance of household income, together with the temporary income shocks share of the total household income variance.

L.1 The Stochastic Process for Income

We follow Blundell and Preston (1998) and start by considering a permanent-transitory decomposition for income. The income of individual i in cohort c and year t is written as:

$$y_{ict} = y_{ict}^P + u_{ict}, \tag{L.1}$$

where y_{ict} represents the permanent component of income and u_{ict} the transitory shock in year t.

We assume that the permanent component follows a random walk, or

$$y_{ict}^P = y_{ic,t-1}^P + v_{ict}, \tag{L.2}$$

where v_{ict} denotes a permanent shock, which is assumed to be orthogonal to u_{ict}. Furthermore, we assume that the variances of the shocks are the same in any period for all individuals in any cohort but that these variances are not constant over time. The cross-sectional covariances of the shocks with previous periods' incomes are assumed to be zero. In this discussion, we assume that shocks are independently distributed across individuals.

The process for income can be written as

$$y_{ict} = y_{ic,t-1} + u_{ict} - u_{ic,t-1} + v_{ict}. \tag{L.3}$$

Defining $\text{var}_{ct}(y)$ to be the cross-sectional variance of the transitory shock for cohort c in year t and $\text{var}_{ct}(v)$ be the corresponding variance of permanent shocks, the growth in the cross-section variance of income for cohort c can be seen from Equation (L.3) to take the form,

$$\Delta\text{var}_{ct}(y) = \Delta\text{var}_{ct}(u) + \text{var}_{ct}(v). \tag{L.4}$$

L.2 Consumption Growth under the Permanent Income Hypothesis

Assuming quadratic preferences, with the discount rate equal to the interest rate, the martingale property of consumption follows from Hall (1978):

$$\Delta c_{ict} = v_{ict} + \frac{1}{\rho_t}\frac{r}{1+r}u_{ict}, \tag{L.5}$$

where $\rho_t = 1 - (1+r)^{-(T-t+1)}$ is a annuitization factor.

Taking variances from both sides of Equation (L.5), we obtain

$$\Delta \text{var}_{ct}(c) = \frac{1}{\rho_t^2}\frac{r^2}{1+r^2}\text{var}_{ct}(u) + \text{var}_{ct}(v).$$

For T large and r small, we obtain the following expression:

$$\Delta \text{var}_{ct}(c) \approx \text{var}_{ct}(v) = \Delta \text{var}_{ct}(y) - \Delta \text{var}_{ct}(u). \tag{L.6}$$

L.3 Consumption Growth for Liquidity Constrained Consumers

We define liquidity constraints as the inability to access credit or saving mechanisms through formal institutions or informal means. The main implication of this definition is that the consumption of fully liquidity-constrained consumers follows their income one-to-one, or $c_{ict} = y_{ict}$. Taking variances and first differences on both sides of this equation, we have

$$\Delta \text{var}_{ct}(c) = \Delta \text{var}_{ct}(y). \tag{L.7}$$

L.4 Estimating the Share of Non-Constrained Consumers

Let h be a dummy for liquidity-constrained consumers. Then, we can decompose $\Delta \text{var}_{ct}(c)$ as follows:

$$\Delta \text{var}_{ct}(c) = p(h=0)\Delta \text{var}_{ct}(c|h=0) + p(h=1)\Delta \text{var}_{ct}(c|h=1). \tag{L.8}$$

Consider now the linear projection of $\Delta \text{var}_{ct}(c)$ on $\Delta \text{var}_{ct}(y)$,

$$\Delta \text{var}_{ct}(c) = \Delta \text{var}_{ct}(y)\beta + \varepsilon_{ct}.$$

The parameter β is identified as

$$\beta = \frac{\text{cov}(\Delta \text{var}_{ct}(c), \Delta \text{var}_{ct}(y))}{\text{var}(\Delta \text{var}_{ct}(y))}. \tag{L.9}$$

From equations (L.4), (L.6), (L.7), (L.8), and (L.9), it follows,

$$\beta = 1 - p(h=0)\frac{\text{var}(\Delta \text{var}_{ct}(u))}{\text{var}(\Delta \text{var}_{ct}(u)) + \text{var}(\text{var}_{ct}(v))}. \tag{L.10}$$

Thus, the extent to which β deviates from 1 depends on (1) the probability that consumers are non-constrained and (2) the share of the income variance change explained by transitory shocks.

To identify the latter, consider now the first-order auto-covariance of y_{ict} for cohort c,

$$\text{cov}(\text{var}_{ct}(y), \text{var}_{c,t-1}(y)) = \text{var}(\text{var}_{c,t-1}(y)) - \text{var}(\text{var}_{c,t-1}(u)). \tag{L.11}$$

From equation (L.11), we have

$$\begin{aligned}\text{var}(\Delta\text{var}_{ct}(u)) &= \text{var}(\text{var}_{ct}(u)) + \text{var}(\text{var}_{c,t-1}(u)) \\ &= \text{var}(\text{var}_{c,t}(y)) - \text{cov}(\text{var}_{c,t+1}(y), \text{var}_{c,t}(y)) + \text{var}(\text{var}_{c,t-1}(y)) - \text{cov}(\text{var}_{ct}(y), \text{var}_{c,t-1}(y)).\end{aligned} \tag{L.12}$$

Furthermore, taking variances on both sides of equation (L.4), we have

$$\text{var}(\Delta\text{var}_{ct}(u)) + \text{var}(\text{var}_{ct}(v)) = \text{var}(\Delta\text{var}_{ct}(y)). \tag{L.13}$$

Substituting equations (L.12) and (L.13) into equation (L.10), we arrive at the following equation for the share of non-constrained consumers:

$$p(h=0) = (1-\beta) \times \frac{1}{\frac{\text{var}(\text{var}_{c,t}(y)) - \text{cov}(\text{var}_{c,t+1}(y), \text{var}_{c,t}(y)) + \text{var}(\text{var}_{c,t-1}(y)) - \text{cov}(\text{var}_{ct}(y), \text{var}_{c,t-1}(y))}{\text{var}(\Delta\text{var}_{ct}(y))}}. \tag{L.14}$$

Hence, from Equation (L.14), it follows that two parameter estimates are sufficient statistics for the share of non-constrained consumers: (1) the coefficient from a regression of the first difference of the consumption variance on the first difference of the income variance, and (2) the first difference of the temporary income shocks variance, expressed as a share of the first difference of the total household income variance. We estimate both parameters below.

L.5 Estimating $p(h = 1)$

We estimate the temporary income shock share of the change in the variance of income at 0.579, using income moments from the ENIGH. This estimate, together with our $\hat{\beta}$ estimate of 0.588 in Column (2) of Table 4, implies a share of liquidity-constrained consumers of 0.288. This share encompasses consumers who face liquidity constraints resulting from a lack of access to formal credit markets and informal insurance mechanisms. Overall, these results confirm the importance of liquidity constraints in shaping consumer and taxpayer behavior in Mexico City.

M Social Norms

A central concern with tax policies involving reductions in enforcement in developing countries is that they may negatively impact social norms. For example, reductions in enforcement could weaken tax morale and the reputational benefits taxpayers derive from compliance, hence reducing the incentive for compliance.

We build on work by Besley et al. (2023) and examine how adding a social component to taxpayer preferences changes our optimal policy prescriptions. We show that allowing social norms to play a role in our theoretical framework does not yield clear-cut policy implications. In particular, whether accounting for social norms strengthens or moderates the welfare impact of marginal policy changes compared to our baseline model depends on whether social norms generate a utility reward for compliant taxpayers that outweighs the utility cost for delinquent taxpayers. Section M.1 presents a modified version of our theoretical model that accounts for social norms. We present the optimal tax rate and enforcement level under this modified model version in sections M.2 and M.3, respectively.

M.1 Theoretical Model

Suppose first that households know how the share of compliant taxpayers, denoted by λ, varies with the tax rate and enforcement. Namely, they know the function $\lambda(t, \alpha)$. Now, following Besley et al. (2023), write the social component of preferences as $S(\mathbb{1}(\text{Delinquent}), \lambda(t, \alpha))$. Denote the relative social payoff of compliance as

$$\Delta(\lambda(t, \alpha)) = S(0, \lambda(t, \alpha)) - S(1, \lambda(t, \alpha)).$$

We assume that $\frac{\partial \lambda(t,\alpha)}{\partial t} < 0$, $\frac{\partial \lambda(t,\alpha)}{\partial \alpha} > 0$, $\frac{dS(0,\lambda)}{d\lambda} > 0$, and $\frac{dS(1,\lambda)}{d\lambda} < 0$.

Households solve the problem:

$$\begin{aligned}
\max_{s,\text{Delinquent}} \quad & u(c) + \beta u(c') + v(g) + S(0, \lambda(t, \alpha)) - (M_i(\alpha) + \Delta(\lambda(t, \alpha))) \times \mathbb{1}[\text{Delinquent}] \\
c \quad &= \quad y - tH(1 - \mathbb{1}[\text{Delinquent}]) - s \\
c' \quad &= \quad s(1 + r) + H(1 - z \times \mathbb{1}[\text{Delinquent}]) \quad \text{and} \quad s \geq 0,
\end{aligned}$$

where savings receive an interest rate r and the level of public goods g is set by the government.

Consider now the effects of taxes on mean utilities:

$$\frac{\partial V^{\text{Delinquent}}}{\partial t} = \frac{\partial S(1, \lambda(t, \alpha))}{\partial t} > 0 \quad \text{and} \quad \frac{\partial V^{\text{Pay}}}{\partial t} = -u'(c_s^{\text{Pay}})H + \frac{\partial S(0, \lambda(t, \alpha))}{\partial t} < 0.$$

Note how, when we incorporate the social payoff of compliance into the analysis, delinquent taxpayers benefit from increases in the tax rate because they lower the social punishment from delinquency. The reason is that they lower the compliance rate, reducing social stigma. From a converse argument, it follows now that tax rate increases are more harmful than before for compliant taxpayers.

Consider now the effects of increasing enforcement by raising α:

$$\frac{\partial V^{\text{Delinquent}}}{\partial \alpha} = -m'(\alpha) + \frac{\partial S(1, \lambda(t, \alpha))}{\partial \alpha} < 0 \quad \text{and} \quad \frac{\partial V^{\text{Pay}}}{\partial \alpha} = \frac{\partial S(0, \lambda(t, \alpha))}{\partial \alpha} > 0.$$

Note that enforcement now has a more detrimental effect on delinquent taxpayers, as they suffer from an increased stigma as compliance increases in response to enforcement. On the other hand, enforcement favors compliant taxpayers, as it increases the social reward from payment.

M.2 Optimal Property Tax

Examining the welfare impact of a tax increase yields the following expression for the marginal value of public funds:

$$MVPF_t = \underbrace{\frac{v'(g)}{u'(c)}}_{\text{Value of Public Goods}} - \underbrace{\frac{1-\gamma(\pi_{0,c}^{\text{Pay}}\Delta c_0^{\text{Pay}} + \pi_{s,c}^{\text{Pay}}\Delta c_s^{\text{Pay}}) - \kappa_t}{1+(1-\tilde{z})\varepsilon_t^{\text{Pay}}}}_{\text{Welfare Cost Per Dollar of Revenue}}, \tag{M.1}$$

where:

$$\kappa_t = \frac{1}{u'(c)HN^{\text{Pay}}}\left[N^{\text{Pay}}\frac{\partial S(0,\lambda(t,\alpha))}{\partial t} + (1-N^{\text{Pay}})\frac{\partial S(1,\lambda(t,\alpha))}{\partial t}\right].$$

Equation (M.1) is similar to Equation (1), with an additional term on the right-hand side, which accounts for the additional welfare effect from social norms changing when the government increases the tax rate to levy an additional dollar of tax revenue. Furthermore, we obtain the following modified expression for the optimal tax rate by setting $MVPF_t = 0$ and rearranging terms:

$$t = \frac{(1+(1-\tilde{z})\varepsilon_t^{\text{Pay}})\frac{v'(g)}{u'(c)} - 1 + \kappa_t}{\gamma(\pi_{s,c}^{\text{Pay}}\eta_{t,s}^c + \pi_{0,c}^{\text{Pay}}\eta_{t,0}^c)}.$$

Note that the sign of κ_t can be positive or negative, depending on whether a marginal increase in the tax rate leads to an improvement in the social reward for compliant taxpayers that outweighs the social punishment worsening for delinquent taxpayers. Thus, the optimal property tax rate could be higher or lower than in the baseline case, depending on the relative magnitude of these effects.

M.3 Optimal Enforcement

Examining the welfare impact of an increase in enforcement yields the following expression for the marginal value of public funds:

$$MVPF_\alpha = \frac{v'(g)}{\omega m'(\alpha)} - \frac{1-N^{\text{Pay}} - \kappa_\alpha}{(1-\tilde{z})N^{\text{Pay}}\frac{\varepsilon_\alpha^{\text{Pay}}Ht}{\alpha} - 1}, \tag{M.2}$$

where:

$$\kappa_\alpha = \frac{1}{\omega m'(\alpha)}\left[N^{\text{Pay}}\frac{\partial S(0,\lambda(t,\alpha))}{\partial \alpha} + (1-N^{\text{Pay}})\frac{\partial S(1,\lambda(t,\alpha))}{\partial \alpha}\right].$$

Relative to Equation (3), accounting for social norms in Equation (M.2) can reduce or increase the welfare cost per dollar of revenue, as enhanced levels of enforcement increase the social reward of compliant taxpayers but worsen the punishment for delinquent taxpayers. Thus, the optimal level of enforcement can be higher or lower than in the baseline case.

Equations M.1 and M.2 show that the effects of policies on social norms can impact the overall welfare effects of policies. In our baseline analysis in Section 7, we abstract from these effects for two reasons. First, as the discussion of these equations shows, it is possible that the effects of policies on social norms can increase or decrease the welfare effects of tax increases or enhanced enforcement. Importantly, it is not always the case that incorporating social norms into the analysis makes enforcement more desirable from a welfare perspective. Second, while the social component of preferences may be important in general, the fact that we do not estimate economically important spillover effects in Appendixes K.3 and K.4 casts doubt on the importance of this force in our setting and for the range of policy changes that we study.