A Bad Bunch: Asset Value Under-Reporting in the Mumbai Real Estate Market

Santosh Anagol Vimal Balasubramaniam
Tarun Ramadorai Antoine Uettwiller*

December 2022

Abstract

Real-estate values are often under-reported to evade taxes and hide wealth built from tax-evaded income. We develop a new method to estimate under-reporting, and employ it on large and granular administrative data from the Mumbai real estate market. The approach compares bunching of reported values around government-assessed guidance values with a third-party measure of true underlying transactions prices. We estimate that 13 percent of value in Mumbai real estate is under-reported between 2013 and 2022. Secondary market transactions witness a post-demonetization decline in under-reporting. Properties with mortgages from public-sector banks and from banks with high non-performing assets exhibit greater under-reporting.

*Anagol: Wharton School, University of Pennsylvania. Email: anagol@wharton.upenn.edu. Balasubramaniam: Queen Mary University of London. Email: v.balasubramaniam@qmul.ac.uk. Ramadorai: Imperial College London and CEPR. Email: t.ramadorai@imperial.ac.uk Uettwiller: Imperial College London. Email: a.uettwiller17@imperial.ac.uk. We thank Cristian Badarinza, Raj Iyer, Ben Lockwood, Shing-Yi Wang, and seminar participants at Imperial College London, UC San Diego, the SITE Financial Regulation Conference, Wharton, Johns-Hopkins SAIS, and the WEFIDEV Workshop in Finance and Development for helpful comments. We also thank Karthik Suresh, Karan Gulati, and Alexandru Zanca for able research assistance, and Raja Seetharaman at Propstack for data access.

1 Introduction

Low-income countries collect less tax revenue than high-income countries as a fraction of GDP, and a perennial question is the extent to which this gap is attributable to differences in the under-reporting of economic activity (Gordon and Li 2009, Kleven, Kreiner and Saez 2016). The problem of under-reporting is particularly vexing when attempting to use property values as a tax base either for property or transaction taxes, as properties are heterogeneous along multiple dimensions. This renders objective valuation difficult, providing strong incentives to property owners to under-report valuations, with substantial tax revenue losses in developed and developing countries alike (Harberger 1965, OECD 2007). This issue is also important in light of calls to increase property taxes in discussions about progressive taxation as an antidote to rising wealth and income inequality (Stiglitz 2015, Piketty 2015).

We propose a new method to detect the under-reporting of property transaction values, and apply the method on a large and granular administrative dataset in Mumbai, India. Studying this question in this setting is interesting for several reasons. First, the decision of how much value to report is a high-stakes decision for households. Home values are typically multiples of household income, and reported home values serve as the basis for transaction taxes, capital gains taxes, and annual property taxes.[1] Second, property tax revenues are important for governments. Property stamp taxes constitute approximately 20% of state government revenues in India, and property taxes constitute over 5% of all government tax revenues in OECD countries, on average.[2] Third, for decades, the Indian authorities have suspected that real estate buyers under-report valuations as a way of laundering unreported income, such as cash earnings and bribes (so called "black money"), and economically massive policy interventions, such as India's 2016 demonetization, have been motivated by the desire to reduce such black money. Fourth, in India, the government's method to prevent under-reporting is to create formulaic assessments of property value based on the physical location of properties, and to set the tax base as the higher of this government-assessed value and the sales price reported by the buyer. This is a commonly used method around the world, that to our knowledge has not yet been carefully analyzed by economists; based on our review of transaction tax policies in the 82 largest cities in the world, 35 of these cities employ this specific system.[3]

[1] In Mumbai, there is a 5% stamp tax, a 1% registration tax, and (small) property taxes are levied in certain sub-regions.

[2] See `https://data.oecd.org/tax/tax-on-property.htm`, accessed March 2022.

[3] Appendix A.1 presents the cities that employ this guidance value system along with their transac-

The statutory requirement in the Mumbai setting is that property purchasers report the true transaction value to the government, to serve as the tax base for the stamp tax and other relevant taxes. If the buyer's reported value falls below the government assessed value (also known as the "guidance" or "circle" value), then the tax base is set at the government assessed value. An important mode of evasion under this system is that buyers and sellers can collude to minimize transaction tax liabilities by under-reporting the true transaction price to the government.[4]

To derive an accurate measure of property value under-reporting in this setting, we develop a new approach that entails comparing the distribution of reported property values around government-assessed guidance values with the distribution of a measure of true underlying transactions prices, which we source from a third-party. In addition to the specific results we uncover, this novel method can be more generally applied to quantify asset value under-reporting in other settings.

To apply the method, we source data on the universe of all property registrations in Mumbai between 2013 and 2022, comprising all property registrations. After applying various filters, our final dataset, which offers the closest geographical match between our property registration data and third-party measured market value data, comprises 260,614 property registration documents (we also work with a dataset of 156,645 transactions with a match to the same project as the transaction, in which we confirm all our findings). The data reveal prominent bunching of self-reported property transaction values at the government assessed values: 10.2% of reported transaction values bunch within 1% of the guidance value, and an additional 22.3% report more than 1% below the guidance value. Interestingly, 67.5% of transactions are reported at 1% or more above the guidance value, suggesting that penalties and/or moral concerns provide substantial disincentive to simply report the minimum government value for a large part of the transacting population.

Standard practice in the bunching literature, as developed by Saez (2010) and surveyed in Kleven (2016), is to use an optimizing model to translate such bunching behavior at points where marginal tax rates change to infer the elasticity of under-reporting and real behavior changes in response to tax rates. Our context does feature a "kink," in the sense that the marginal tax rate below the guidance value is zero,

tion tax rate. A detailed spreadsheet of valuation systems for the top 82 cities of the world can be found here.

[4] The Indian tax administration (and anecdotal reports) discuss that the difference between the reported value and the true transaction price is often transferred from buyer to seller in currency notes, to avoid detection of tax evasion through the formal financial system. See Indian Department of Revenue's "White Paper on Black Money," 2012, `https://dor.gov.in/sites/default/files/FinalBlackMoney.pdf`.

but then discretely increases to 5% (the transaction tax rate) for the marginal rupee reported above the guidance value. Viewed through this lens, the bunching of transactions prices at guidance values appears consistent with under-reporting, and as we later report, we estimate "evasion elasticities" using this approach. The problem, however, is that such bunching could also be consistent with truthful reporting for at least two reasons. First, government-assessed property values could be extremely accurate and timely estimates of true underlying transactions prices. Second, buyers and sellers might perfectly anchor transactions at guidance values.[5] If true underlying market values were observable, we could easily distinguish bunching due to under-reporting versus bunching due to these alternative explanations, but the existence of under-reporting itself renders market values difficult to observe.

To address this challenge we develop a method that uses a noisy measure of true underlying market prices to estimate under-reporting behavior. The key insight is that bunching due to the coincidence of market prices with guidance values should produce small differences between more aggregated measures of reported values and more aggregated (even if noisily measured) market values. In contrast, bunching due to under-reporting produces an aggregate reported value distribution that lies below aggregate measured market values by the extent of under-reporting. The aggregation is fundamental here, as it smooths out the noise inherent in imperfectly measured market transactions prices. The method also exploits the idea that a simple economic model of under-reporting suggests that under-reporting should be largest for "buncher" transactions, and decline as we look at transactions where the buyer endogenously chooses increasingly higher values to report than the guidance value.

We implement the method by matching our administrative data on reported values to a third-party provided price dataset of new buildings developed and sold during our sample period. As we later describe, these "mystery shopping" data are based on collecting pricing sheets and other marketing materials directly from developers, developer sales offices, and mail flyers sent by developers. While these data are purchased and used by banks, developers, and investors in the real estate space, we nevertheless expect these third-party price data to be a noisy measure of true transaction prices. When we measure average under-reporting using our approach applied to these data, we estimate under-reporting rates of approximately 20% for "buncher" transactions, and roughly the same rate for transactions with reported values less than

[5] While we are unaware of direct evidence that market prices anchor on government assessed values, both Genesove and Mayer (2001) and Andersen et al. (2021) show that property sellers anchor on original purchase prices.

guidance values. These estimated under-reporting rates decline linearly for transactions where the buyer chose to report more over the guidance value; for buyers reporting greater than 50% above the guidance values we estimate zero under-reporting. This pattern of declining estimated under-reporting as buyers report higher and higher prices above the government-assessed guidance value is precisely what we would expect if we had true measures of transaction prices. This also lends credence to our specification of our observed proxy prices as a noisily measured version of true transaction prices. Aggregating across all transactions, we estimate a 13% under-reporting rate. Assuming these under-reporting rates are representative for Maharashtra state, this translates into a large annual revenue loss for the state government: in 2021, a loss of US$ 475mn lower tax revenues from property transactions (or 1% of total revenue from this source).

Governments have strong incentives to set accurate guidance values. That said, it is unlikely that they are perfect measures of true underlying value, given the fairly broad geographical regions over which they are set (this is especially true in our context, given the density of physical properties in Mumbai). Moreover, these values are infrequently updated, which leads to both staleness and inaccuracy. We use this insight to uncover additional behavioral responses using our approach, studying the revision of government-assessed values across multiple neighborhoods, which occur in three of the seven years in our sample.

We find large spikes in the volume and value of registered transactions in the days and months directly before guidance values are scheduled to change, suggesting substantial gaming behavior. We also find that market participants rush transactions right before guidance values change, and uncover evidence that they back-date transactions to exploit lower guidance values. These spikes translate into approximately 6%-12% higher estimated under-reporting rates in months immediately before guidance value changes. We argue these time-series patterns are unlikely to be driven by anchoring, since if bunching comes from transactors setting market prices by anchoring on the guidance value, we would expect similar amounts of bunching over time—it is unclear why such behavior would suddenly become more frequent in the month before a guidance value change. We also argue that the time-series evidence is not consistent with the hypothesis that bunching is primarily determined by guidance values perfectly tracking true market prices. The third-party data suggest that market prices have drifted upwards over time, so we would expect guidance values to be least accurate in the months right before scheduled guidance value changes, and most accurate in the month after these changes, meaning that this explanation predicts greater bunching *after* rather than *before* guidance value changes, contrary to the data.

To understand the underlying economics, we study how measured under-reporting using our approach varies across important events and types of market participants. We first study how India's "demonetization" changes under-reporting, which in theory should have made the acquisition of cash more difficult (in the short run), and raised the potential cost of accepting cash in the long run. While we do not see a strong aggregate pattern of reduced transaction numbers or values in the weeks surrounding demonetization, we do find that under-reporting rates for re-sale secondary market transactions (but not developer-sold primary market transactions) dropped by roughly one-half following demonetization.[6] This finding is consistent with the theory in Kleven, Kreiner and Saez (2016), who argue that it is more difficult for large organizations to maintain the collusive agreements that underpin tax evasion, as secondary market transactions only require a single buyer and seller to maintain such an agreement.

We also assess heterogeneity in under-reporting behavior across different types of sellers, by transaction value, and by whether the transaction is associated with a mortgage. In addition to the patterns surrounding demonetization, we find more generally that under-reporting rates in secondary market resale transactions are higher than those in primary-market developer sale transactions. Moreover, we find that under-reporting rates are higher for lower-value transactions, where detection probabilities may be perceived to be lower, and the economic penalties for under-reporting are lower in absolute value.

Finally, we merge reported transactions values with administrative data on mortgage values. We find that the reported values of transactions with low mortgage loan-to-value (LTV) ratios exhibit the greatest extent of bunching at government-assessed values, while transactions with high LTV ratios exhibit the least bunching. This suggests an association between financial constraints, as expressed in the desire for mortgage credit, and tax evasion. We also find correlations between the bunching behavior of borrowers and the ownership structure of lending banks (i.e. government versus private sector banks), as well as with lending banks' non-performing loan rates, which is consistent with a link between financial intermediation, credit screening technology, and tax evasion through under-reporting. More specifically, we estimate the greatest under-reporting for properties with mortgages from cooperative and public-sector

[6] We also study the effects of the introduction of the Goods and Services Taxes (GST) and the Real Estate Regulation Act (RERA), and find little evidence of changes in reporting behavior associated with these events.

banks, and least for those with private and foreign banks; mortgages from banks with high overall non-performing loans are also associated with high bunching behavior.

The remainder of this paper is organized as follows. The remainder of this section reviews related literature. Section 2 describes the institutional background for property valuation for tax purposes in our setting and elsewhere. Section 3 sets up a simple model to guide our empirical work. Section 4 describes our data. Section 5 documents our baseline results. Section 6 documents heterogeneity in our measures of under-reporting, and Section 7 concludes.

1.1 Related Literature

Our paper contributes to the large literature on tax evasion. Our approach is most similar to previous work that uses third-party derived estimates of economic activity to study evasion behavior. Fisman and Wei (2004) estimate tariff evasion on imports into China from Hong Kong by comparing reported imports in China to the more accurately measured exports from Hong Kong to China. Slemrod (2007) discusses randomized audits, conducted by tax authorities, as a method of estimating aggregate U.S. income tax evasion, and Kleven et al. (2011) reviews the impact of audits on evasion. Pissarides and Weber (1989) pioneered examining consumption expenditure as an indicator of true income, finding that the self-employed have higher rates of consumption relative to their reported incomes, and Braguinsky, Mityakov and Liscovich (2014) more recently applies this method to administrative data on car ownership in Russia combined with reported earnings. Artavanis, Morse and Tsoutsoura (2016) use bank determined credit capacity as an independent signal of true income, finding relatively greater credit limits conditional on reported income for the self-employed.

We find under-reporting that is large in absolute terms, but our estimated under-reporting rates of 13% are low relative to prior small-sample estimates in India (40% under-reporting) and income tax evasion estimates for the self-employed (40%-80% across developed and developing countries). Our estimate of the property value under-reporting rate is potentially useful for studies on how neighborhood change, transportation infrastructure, and zoning reforms affect real estate prices, as these studies often use government assessed property values in lieu of frequently unavailable market price data (Anagol, Ferreira and Rexer 2021, Tsivanidis 2019, Gechter and Tsivanidis 2020, Harari, Wong et al. 2018).

In terms of possible remedies, we note that Pomeranz and Vila-Belda (2019) survey research with tax authorities, much of which focuses on policy interventions aimed at increasing tax revenues. To our knowledge this work has not studied real estate under-reporting, especially in contexts where agents can choose to report at or above government-assessed values, though Casaburi and Troiano (2016) study the political economy consequences of an Italian national reform that aimed to force property owners to register their land so as to enter the tax base (an extreme form of asset value under-reporting is hiding property ownership from the government).

Our work also relates to the literature on transaction taxes, which has typically focused on advanced economies (e.g, Best and Kleven 2018, Kopczuk and Munroe 2015, Dachis, Duranton and Turner 2012), and has typically not estimated the importance of asset value under-reporting, presuming that third party reporting by mortgage lenders, real estate agents, etc. eliminate the ability of buyers and sellers to under-report transaction prices. For example, Kopczuk and Munroe (2015) finds no evidence of evasion regarding a mansion transfer tax in New Jersey.

Our work is also related to a small but growing literature studying under-reporting behavior in China. Fan, Wang and Zhang (2022) use data from Shanghai to estimate under-reporting, and Agarwal, Li, Qin, Wu and Yan (2020), and Agarwal, Kuang, Wang and Yang (2020) use similar data from individual Chinese real-estate brokerages to estimate under-reporting, doing so by cleverly utilizing data on underlying transaction prices collected by the brokerages which serve as the basis for brokerage commissions.

In the Mumbai setting, similar to many cities in developing countries, there is no administrative data recorded on true underlying transaction prices. This means that the method we develop, which combines third-party market value estimates with guidance and reported values, can be applied more broadly to detect under-reporting whenever the government sees reported and guidance values and can source estimated listing prices from analytics companies (as we do) or online listings portals.

Our paper connects to the literature on ongoing property taxation (i.e., taxes paid annually to governments as a percentage of the home value). In this context, we are the first to analyse under-reporting in a system of government-assessed values, a common approach to setting property tax bases around the world. The literature finds that assessed values for property taxes often diverge from recent transaction prices in systematic ways, with important distributional consequences (Avenancio-León and Howard 2022). A key difference in our context is the statutory obligation for homeowners to report the true market value, while in typical advanced economy property tax contexts, homeowners are required to pay tax on the government's assessed value, even if they know the assessed value is different from the market value (Amornsiripanitch (2020) reviews this literature).

We apply our method to estimate the impact of the 2016 demonetization policy, which allows us to understand the relationship between the broader availability of cash and tax evasion. This relates our work to early analyses of black money in Indian real estate, which studied the government's pre-emptive purchase provision. Under this provision, the central government tax authority was allowed to purchase any property at 15% above the reported value, creating strong incentives for accurate reporting.[7] In statute, the government was supposed to randomly select property transactions to determine whether to exercise this right. Whether this policy was followed is unclear, with most sources suggesting the sampling was not conducted randomly. National Institute of Public Finance and Policy (1995) estimate 44.8% under-reporting using a small sample of Mumbai transactions under this policy,[8] and the same study conducts a survey of real estate brokers and concludes from this evidence that approximately 60% of true transaction values were under-reported (for earlier small-sample estimates see, Tandon 1987, Gopalakrishnan 1986). In a survey article on Indian transaction taxes, Panchapagesan (2017) notes that there have been no more recent aggregate estimates of black money in Indian real estate transactions.

While our mortgage results are based on a limited sample of transactions that we were able to match to mortgages, they highlight the general point made in Basu (2015) that under-reporting of home values can have potentially important macro-prudential implications. This connects our work to the broader literature on housing collateral value misrepresentation during the global financial crisis (Piskorski, Seru and Witkin 2015, Griffin and Maturana 2016), as well as to Montalvo, Piolatto and Raya (2020), who estimate transaction tax evasion in Spain, focusing on a buyer's trade-off between under-reporting to avoid transaction taxes and over-reporting to obtain greater mortgage credit. Finally, our results on mortgage bank ownership and under-reporting also connect our work to studies of credit screening differences across public and private sector banks (La Porta, Lopez-de Silanes and Shleifer 2002, Mishra, Prabhala and Rajan 2022).

[7] This system appears to have been proposed in the economics literature by Harberger (1965), although Taiwan had a similar, largely unsuccessful, system implemented around the same time period Chang (2012). See Posner and Weyl (2019) for other examples of such "self-assessment" based mechanisms. A challenge to these systems is that those in charge of implementing the policy may be bribed to avoid exercising the government's right on certain properties.

[8] National Institute of Public Finance and Policy (1995) does not directly report the sample size for this estimate, however Table 3.1 in that study counts 46 properties purchased in Mumbai under this program.

2 Institutional Background

Our paper focuses on the valuation of property for taxation purposes. Systems for property valuation vary around the world, and can be broadly classified into two types. The first type of system is one in which taxation authorities generate "decentralized" or property-specific assessed values. Such decentralized systems can vary in the way that property-specific values are created. In some cases, assessors determine property valuations using some combination of site visits and comparable analysis to determine a valuation. In other cases, assessors determine broad features of the property (i.e., whether it is residential or commercial, square footage, and so on) which are then used as inputs into a hedonic model to determine the guidance value. (For example, the Danish system of tax assessor valuation at different points in time adopted both property-specific and model-based property assessments, see, e.g., Andersen et al. (2021)).

Using assessors to periodically evaluate individual properties has the benefit of producing more accurate assessments, which is particularly valuable given the inherent heterogeneity of (and therefore potential for unobserved quality in) real estate even within small regions. There are two major challenges with individual assessments, however. First, they are costly to implement, given the large number of assessments needed, and the relatively small number of qualified assessors. Second, property owners may bribe assessors to lower their assessments, as carefully documented in Khan, Khwaja and Olken (2016).

Typically, in systems with decentralized assessment, the statutory tax base is the government's assessed value, meaning that the owner has no obligation to report the true market value. An alternative method of decentralized valuation is the so-called "self-assessment" method as proposed in Harberger (1965). These systems encourage truthful self-assessments of property values by giving the state the right to purchase the property at the property owner's self-assessed value. Chang (2012) argues that even with these incentives property owners in Taiwan greatly under-reported property values, because the probability of the state actually exercising the right was too low.[9]

Centralized systems of property valuation for tax purposes are present in many jurisdictions, such as the Indian context we study here, as well as Brazil, Colombia, Mexico, Thailand, Indonesia, Philippines, and New Zealand, among others, as we document in Appendix A.1. In such systems, the authorities set location-specific as-

[9] We are not aware of any jurisdictions that has successfully implemented a "self-assessment" system of property valuation.

signed valuation bases, which are periodically updated as market prices evolve, and set these (usually per square-foot) valuations as a lower bound tax base for all properties physically located in specific areas. Such centralization costs less than decentralized property-specific valuation, and lowers the probability of captured assessors. However, centralized valuations increase the potential for mismatch between assessed values and market valuations, especially if the valuation bases are infrequently updated. In such systems the statutory requirement is that owners report the true market value of their property, with the tax base set as the maximum of the government's assessed value and the owner's reported value. The owner faces a penalty, typically a multiple of the amount of tax avoided if they report a value lower than the true market value, but generally have some form of (costly) recourse available to prove that their lower valuation can be justified.

Conversations with market participants, and reports from regulators in the Indian centralized tax assessment system suggest that under-reporting typically occurs as follows.[10] The buyer (usually an individual or household) and seller (either an individual/household or a real estate developer) of an apartment agree on a transaction price. If this price is higher than the guidance value, to avoid taxes, they also agree to under-report the transaction price on the registration document, often reporting exactly the guidance value.[11] To prevent detection of the under-reporting by paper/digital trail, the gap between the transaction price and the reported price is paid in currency notes. In this way, the corresponding bank records of the transaction will also be in agreement with the reported value on the registration. (This is important, as it is required that buyers and sellers report their tax-identification numbers, so reported real estate values can be easily linked to the transactors' bank account information.) Typical methods of obtaining large sums of cash include accrued currency from operations of a cash business, relatively small withdrawals taken from a bank account over time, or writing a check as a "gift" to a friend or relative in exchange for the cash. In some cases, such funds are sourced or earned completely outside the tax net, and often referred to as "black money" in the Indian context. Interviews with market participants suggest there are very few audits of reported values as long as the reported value is greater

[10] See, for example, `https://timesofindia.indiatimes.com/blogs/law-street/black-money-does-the-devil-lie-in-real-estate/`.

[11] While the law states that guidance values should be formulaic, following centrally assigned "circle rates," it is possible that the tax authority manually enters a valuation at their discretion. While this practice was not mentioned in our interviews with market participants, Comptroller and Auditor General of India (2016) discusses a few large transactions where the guidance value formulas were not followed. As we have administrative data on guidance values, our method detects under-reporting arising from such inspector discretion.

than or equal to the guidance value.[12]

Although market participants note that under-reporting is common, the general sentiment is that market prices are well understood, especially by developers and real estate brokers. Brokers are even known to quote market prices based on how much one is willing to under-report (i.e there is a lower overall price if a greater amount is paid in cash). However, it is generally not an easy task for the government to create systematic data on market prices (perhaps by interrogating brokers) to conduct the kind of comparisons we do. During our sample period there are also websites that host housing listings, but implementing our approach would require scraping and painstakingly matching apartment buildings from such websites, a task that to our knowledge the tax authority has never under-taken. (Major listing websites include `www.housing.com`, `www.99acres.com` and `www.proptiger.com`.)

3 Economic Framework

A household purchases a property for market price m and then chooses the amount to report to the government r. The property has an associated government assessed value c, and τ is the transaction tax rate. The incentives for reporting r differ based on whether the market price m is greater than or less than c. We first discuss the case when $m > c$, and subsequently explore the setting when $m < c$.

When $m > c$, the transaction tax liability is $\tau \times \max(r, c)$, i.e., the tax base is the maximum of the reported value and the government assessed value. If we include no other incentives in the model, then we would expect all households to report $r = c$ to minimize their tax burden. It is fair to say that in the empirical setting, the tax implications of under-reporting are more complex. The ongoing property tax is also a function of the reported value, providing an additional incentive to under-report. In theory, there may also be an incentive to "over-report" to reduce future capital gains taxes. We discuss the implications of these issues later in the paper.

Let π_1 be the perceived probability that the reported value r is verified by the tax authorities, either through a physical audit or other mechanisms. We assume for now that the probability of verification/detection is orthogonal to the reported property value, and results in the perfect discovery of the market price m. The penalty in the case of detection for under-reporting is that the buyer must pay n times the amount

[12] We have submitted a right-to-information request to obtain data on the number of audits, which are not publicly available.

in transaction taxes that he/she avoided. In our empirical setting, $n = 4$, i.e. the maximum penalty is four times the original transaction tax amount.

When $m > c$, the buyer chooses r to minimize the expected tax burden:

$$\min_r \tau[(1 - \pi_1)r + \pi_1(r + n(m - r))] \tag{1}$$

This formula assumes risk-neutrality. If households are risk averse over expected tax and penalty payments, this would move them towards reporting higher amounts, because the variance of any tax payments will be lower if they do so. We also note that this penalty term is independent of the government assessed value c, as it is in the empirical context in Mumbai.

The minimization problem in equation 1 has two corner solutions. If the expected marginal cost of under-reporting $\pi_1 n\tau < \tau$, the household reports $r = c$. This is because reporting an additional rupee has a marginal cost τ, while under-reporting this rupee has an expected marginal cost $\tau\pi n$. If $\tau\pi n < \tau$ then agents will under-report all the way down to the government assessed value, and all transactions will bunch at the government assessed value. In contrast, if $\pi n > 1$ then the agent reports the full market value $r = m$. In this case the expected penalty cost of under-reporting the marginal rupee is greater than the tax rate saving, which is the benefit of under-reporting.

Now, let π_2 be the perceived probability of a successful appeal to the authorities that $m < c$. In this case, the buyer's true market value is less than the guidance value c, and they will have to appeal to tax authorities that the tax burden is $\tau \times r$ rather than $\tau \times c$. We assume that t is the transaction costs of engaging in the appeals process, and that the total cost increases with the distance that r lies below c. The buyer once again chooses r to minimize the expected tax burden:

$$\min_r \tau[(1 - \pi_2)r + \pi_2(r + t(c - r))] \tag{2}$$

If the expected marginal cost of appeal is low enough (in practice, π_2 is small), buyers have an incentive to report $r = m$ and appeal, and if it is large, buyers simply report $r = c$.

3.1 Estimating Under-reporting from Bunching

On the face of it, this simple framework suggests that we should be able to infer households' perceived π values by inspecting the distribution of r around c. However, as the model shows, accurate inference also depends on the relationship between c and the true market value m. If government assessed values are perfectly aligned with market values, and reported values are completely truthful, we would still see bunching exactly at c even without any under-reporting. This is our main identification challenge, i.e., the need to distinguish between bunching arising from under-reporting and bunching because government-assessed values coincide with market values.

To make progress we need a measure of m. If such a measure were somehow available for each transaction, we could directly estimate under-reporting at the transaction level with no need for a bunching strategy. Even so, understanding how a bunching-based strategy to detect under-reporting would work with perfect observation of m is helpful, as it sets a benchmark for applying such a strategy in the more realistic scenario of observing a noisy proxy of m.

To fix ideas, let θ be the fraction of households that believe that $\pi_1 n < 1$. We refer to such households as "under-reporters," who report c irrespective of the true transaction price m; the remaining $(1-\theta)$ "truthful-reporting" households believe that $\pi_1 n \geq 1$, and report $r = m$. Our goal is to extract a measure of θ from observed bunching behaviour.

Let $f^j(m)$ be the probability density function of market values for all households of type j, where $j \in \{\text{under-reporter,truthful-reporter}\}$. We do not observe a household's type, but household $i's$ triple (r, c, m) reveals its type perfectly. If $r = c$ and $m > r$, then $j =$ under-reporter. If $r = m$, then $j =$ truthful reporter.

We can estimate θ by calculating the fraction of households with $r = c$ and $m > r$, divided by the total number of households. The *aggregate* amount of under-reporting is then the difference in aggregate m and aggregate r for those who report $r = c$. In order to compute aggregate under-reporting, we do not have to condition on type because aggregate under-reporting for truthful reporters is zero.

Figure 1a simulates the distributions of reported and market values around government assessed values for the case in which $\theta = 40\%$. The x-axis in this plot shows $\frac{r-c}{c}$ for the r distribution, and $\frac{m-c}{c}$ for the m distribution. For the purposes of the simulation we assume that the distribution of m around c is normal with a mean of 1 and standard deviation of 10, and that the underlying distribution of c is normal with a mean of 10 and standard deviation of 1. In this case, there is substantial bunching of r around c, with 40% of households with $m \geq c$ choosing to report c.

It appears, therefore, that the underlying θ parameter can be backed out, therefore, by inspecting how the bunched distribution of r around c differs from the smoother distribution of m around c. As discussed earlier, however, the extent to which the government-assessed values c track market values m is a key confound. In order to assess how important this is, we require an independent measure of m which is neither r nor c. For now, Figure 1a is drawn under the assumption that we have access to a measure p which is a perfect estimate of market value m. More realistically, p is a noisy measure of m, as we discuss below.

3.2 Household Leverage, Cash Payment, and Preferences vs. Beliefs

Thus far we have focused on buyers' subjective assessments of detection probabilities and penalties for under-reporting as the main drivers of reporting decisions. There are, of course, other costs and benefits that also influence household reporting decisions. For one, banks will typically only lend up to the reported officially registered value of the property. A buyer in need of funds therefore has an incentive, on the margin, to report more than the government-assessed "circle" rate value to obtain additional mortgage credit. This buyer will trade off the higher transaction tax associated with a higher reported value against the marginal benefit of obtaining more mortgage credit. This marginal benefit from relaxing the leverage constraint could include, for example, the ability to buy a more expensive property, or to use the borrowed funds for other investment or consumption purposes. We later explore this motivation by extending our economic framework, and empirically conditioning bunching behavior on whether the transaction has an associated mortgage.

Second, under-reporting buyers require cash to pay the difference between the reported value and the market value. There are likely significant differences in households' ability to obtain cash funds. Households who do not report sources of income to the tax authority may already have low-cost access to "black money". In contrast, salaried employees at large companies (i.e. information technology workers) may have difficulty accessing cash given that their income will be directly deposited into banks, meaning that large cash deposits and withdrawals are likely to attract additional scrutiny. The same goes for sellers—anecdotal evidence suggests that builders that make payments to workers and other suppliers in cash are more willing to accept large currency payments. This is in contrast with individual salaried sellers who have a more difficult time explaining or ultimately depositing/laundering large cash receipts without attracting government scrutiny. We therefore also later condition bunching behavior on whether the seller is a corporate entity or an individual.

Third, households or other property sellers may also prefer to be honest either for their own ethical motivations, or because there are pecuniary returns to being ethical.

One potential implication is that the "beliefs about detection" frame that we outlined is not the only possible interpretation of the mass of under-reporters. (A way to conceive of such ethical motivations in our setup would be to model a higher perceived value of the penalty n incurred in the event of detection to capture the experienced costs of ethical failure for truthful reporters and vice-versa.) Another potential implication is that some real estate developers may condition sales on the requirement that buyers report the true market price. Anecdotal evidence suggests that large developers require buyers to report honestly because they prefer sales records to reflect the true income received from these sales. One reason for this is that revenues can be transparently reported for audit certification and other accounting purposes. For example, a publicly listed developer, or a developer planning an IPO will have strong incentives to have all income transparently recorded because of significant public scrutiny (see, e.g., Coffee Jr (2002) and references therein). While such motivations might affect the interpretation of the underlying mechanisms driving under-reporting behavior, the main estimation ideas go through regardless of whether preferences or beliefs are at work.

We now turn to describing and analyzing the data.

4 Data

4.1 Registrar Data

Our dataset comprises real estate transaction registration documents from the Inspector General of Registration and Controller of Stamps (IGR), Department of Revenue, Government of Maharashtra, India. For our analysis, the important information in these documents is: 1) the reported property value; 2) the government assessed value; 3) the transaction tax paid; 4) the property's floor space area; 5) information (such as whether they are corporate or individual) about buyer and seller; 6) the transaction date; 7) the registration date. These data are publicly available from the registrar's website, and cover all available transactions over the period from 2013 through 2018. See Appendix A.1 for a complete description of this data and example documents.

We augment the IGR administrative data with data provided by Propstack Analytics, a for-profit real estate firm that uses transactions to provide pricing and ownership information via its Zapkey data platform.[13] In addition to the IGR variables described above, Propstack Analytics also provides us 1) an indicator for whether a property was

[13] See `https://www.propstack.com/` and `https://www.zapkey.com/` for details.

sold by the developer directly (a "primary" sale) or sold by an individual 2) the number of buyers 3) the unit number 4) the floor number of the apartment 5) the name of the real estate project associated with the transaction, and 6) the latitude and longitude of each project location in Mumbai. Propstack covers the universe of all transactions reported in the IGR admin data, thus resulting in no loss of information for our analysis. We further confirm that the overlapping data from the two sources are identical, using a 10% random sample of Propstack transactions between 2013–2018 which we find perfectly match IGR reports. Figure A.12 presents a comparison of aggregate counts of transactions and tax revenue from the IGR data, Propstack Analytics, and the aggregate numbers of transactions reported by IGR for the Mumbai Metropolitan Region—a region larger than the coverage of our sample but the closest level of aggregation for comparison.[14]

Between the IGR data and Propstack Analytics, we cover data from Mumbai and Mumbai suburban areas from 2013–2022 worth US$106.92 billion. This region is the most important metropolitan area in the state of Maharashtra, a state that generates a quarter of India's GDP. Our region of study remits approximately 30% of the state's total stamp tax revenues.[15] This region is also one with the most comprehensive spatial information that we leverage on for our empirical analysis.

4.2 Propequity Data

Our independent source of price data, i.e., our empirical measure of p, comes from Propequity, a real estate analytics firm that maintains a subscription real estate information portal for the Indian real estate market.[16] Propequity is a for-profit analytics firm that primarily earns revenue by selling access to its data products. The subscribers are real estate public and private equity investors, banks and real estate developers. The primary use case is to understand trends in local prices and quantities for new residential projects being developed.

Propequity aims to provide data on all new real estate projects in India with potential revenues over 10 million rupees (roughly US$ 200,000), with coverage varying over locations. Over the time period 2013 to 2022 this dataset includes information

[14] The government assessed value is determined by multiplying the "circle" rate (set on a per square meter basis for a given sub-zone $\times$ year within the city) by the area of the property. Additional adjustments to the circle rate are made based on other features, such as the floor on which the property is located or whether a parking spot is included with the property.

[15] Region-wise stamp tax revenue sourced from `https://igrmaharashtra.gov.in/dashboard_Data_ArticlewiseAndYearwise.aspx?GvData=maharashtra`.

[16] `https://www.propequity.in/`

on approximately 11,930 real estate projects (each such project has multiple apartment buildings which in turn have multiple apartment units) from the Mumbai and Mumbai suburban regions. For each project we observe the following time-invariant characteristics: longitude, latitude, a masked developer ID, the number and format of apartment units and amenity information, date project units started being sold, date project construction was completed, luxury status, and a few other features, in addition to an estimated current sales price of the apartments in the project, which is the main variable for our purposes. The quarterly price data are reported as the price per square foot for a "base" level apartment in the building (i.e. excluding optional amenities like parking spaces, higher floor levels, etc.).

These price data are sourced through two major methods: 1) physical visits to developers' sales offices to collect pricing sheets for projects 2) collecting developer-emailed advertisements of projects, which report prices. Developers typically market their apartments at a per apartment price. The data provider converts this to a per square foot price using the developer's reported "carpet" area per apartment, which is the area of usable space within the apartment including interior walls, but excluding exterior walls, outdoor spaces such as balconies, and any public spaces within the building. The provider uses these data to conduct valuations for banks that make mortgages, and note that in this context they are asked to provide an estimate of the total value of the apartment, including both the reported value and any under-reporting to avoid the transaction tax. For what it is worth, the provider reports that their prices are likely to be over-estimates of transaction prices, by roughly 2-4% given developers' incentives to obtain the highest possible prices for their apartments (which may go lower following negotiations with buyers).[17]

We map these price data to p, our estimate of m, the true market value for transactions in the administrative data, and investigate any effect of measurement error in p in greater detail below. To do so, we match Propstack transaction level data to Propequity data by the name of the project the transaction belongs to, and the location (by latitude and longitude) every quarter. Of the total of 260,614 transactions recorded in Mumbai and Mumbai suburban regions between 2013 and 2022, 60.01% or 156,645 transactions are matched to Propequity information on the same project to which the transaction belongs to. We match the remaining 40% of transactions to the nearest project available in Propequity. Figure A.9 documents the match quality in the data. Overall, 95% of all

[17] Propequity also reports an estimated number of units sold within the building in a given quarter. As we have administrative data on number of sales based on registration documents, we do not use this information in our main analysis.

transactions in the data are matched to Propequity transactions within 500 meters of the latitude-longitude of the purchase transaction. We eliminate from our analysis successful matches that are more than 1 kilometer away from the location of the associated Propstack transaction. Figure A.10 shows the spatial distribution of the transactions in our final sample for study.[18]

Table 1 presents means, medians, and counts by year for our primary analysis variables. 71% of transactions are sales made by a corporate entity (typically a real estate developer), with the rest being made by individuals, and the average property area is 76 square meters (818 square feet). The average reported value over the sample is US$ 321,770. The average government assessed value is 19% lower, at US$ 261,860, but the average p value is higher, at US$ 367,290. Based on these raw averages we would estimate an under-reporting rate of 12.4%, although we have yet to confirm that our estimated p values are sensible proxies for true transaction prices.

Appendix figure A.11 shows a binned scatter plot of p, r and guidance values c from our main sample of 260,614 transactions. The p and r values are highly correlated with c, suggesting that guidance values c are set to match geographic variation in market prices. Reported values r are lower on average than estimated market values p, but also strongly correlated with them. The strong relationship between p and c lends credence to the idea that government assessed values can be useful proxies for market values in developing country cities where high quality individual transaction market price data is not available (e.g. see Anagol, Ferreira and Rexer (2021), Tsivanidis (2019), Gechter and Tsivanidis (2020), Harari, Wong et al. (2018)). The high correlation between p and c is also consistent with the finding of a high correlation between assessed property values and high quality individual transaction data reported in Ahlfeldt, Redding, Sturm and Wolf (2015), who study Berlin. We now turn to documenting bunching in the data, and to our measures of property value under-reporting using these matched data.

5 Baseline Results

Figure 3 plots the empirical distribution of $\frac{r-c}{c}$ around zero. This plot is, to our knowledge, the first analysis of reporting behavior around government assessed values in this or other markets. The figure reports the number of transactions within 2% bins, with the bin around zero ranging from -1% to +1%.

[18] As we describe later, for our mortgage analysis, we use IGR data between 2013 and 2018 augmented with registered mortgage transactions. Appendix Section A.1 describes sample construction in detail.

First, the plot reveals that 14.7% of households report below the government assessed value, with an average (median) of 24.9% (14.7%) lower than the guidance value. This is further evidence of the imperfect tracking of m by c, and suggests that a non-negligible fraction of households is willing to pay verification costs to certify that $m < c$.

Second, there is a clear spike in r at government assessed c. As noted earlier, such bunching could be due to either a highly peaked distribution of true transaction prices (and primarily truthful reporting), or a flatter distribution of market values and a substantial amount of under-reporting. The figure also shows how our proxy for market values p varies around the circle rate, which points towards the latter interpretation being correct, to the extent that p is a reasonable proxy for m.

We now analyze the distribution of reported values relative to government assessed values in light of our simple model that only included a trade-off between tax savings from under-reporting and penalties associated with detection. To do so, we first assume that households act as if detection/audits are random, and that all households have the same beliefs about the probability of detection. Under these assumptions, a household will report r equal to c (i.e., bunch) if reporting one more rupee in value costs them more than the expected penalty: $\tau > \tau\pi_1 n$. Given that the statutory penalty for under-reporting is four times the avoided stamp tax ($n = 4$), we have that non-bunching households must believe that $\pi_1 > \frac{1}{4}$. Conservatively classifying only households that report 10% above the government assessed value as non-bunchers, this reporting behavior suggests through the lens of the model that 63% of households believe the audit probability is over $\frac{1}{4}$. While there is no official data on the audit rate of transactions, anecdotal evidence suggests that the rate of detection/audit is much lower than 25%, suggesting there are likely other non-penalty incentives driving behavior.

We now turn back to more closely analyzing the green line with triangles in Figure 3, which shows the estimated set of market values based on the Propequity dataset (p). To the left of zero, the distribution of $\frac{p-c}{c}$ is quite close to the distribution of $\frac{r-c}{c}$, which is consistent with buyers of properties with $m < c$ truthfully reporting market value. To the left of zero, a small fraction of transactions appear to have significantly lower values than guidance values. Consistent with the insights from the model, the blue line with circles appear lower than the green line, suggesting that there are incentives to bunch at zero even for those transactions whose $m < c$. The $\frac{p-c}{c}$ distribution shows no bunching at zero, but instead a smooth distribution centered around 0.3 above zero. To the right of the bunching region, the market value distribution appears to be a rightshifted version of the $\frac{r-c}{c}$ distribution.

5.1 Measurement Error

As discussed earlier, the patterns documented in Figure 3 are consistent with under-reporting, as shown in Figure 1a. These patterns are also potentially consistent with *no* under-reporting if p, the proxy for m, is measured with substantial error.

To develop intuition, Figure 1 shows two (extreme) simulations from the model, both of which show substantial bunching of r at c, but caused by different levels of under-reporting. In the "high under-reporting, without measurement error" case, as before, $\theta = .4$ (i.e. 40% of households under-report) and market prices are observed without error (i.e., $p = m$); this is the baseline case discussed in Section 3. The excess bunching mass in this case is only due to under-reporting.

In contrast, when there is "no under-reporting, with measurement error" ($\theta = 0$), the bunching arises from market values m equal to government assessed values c. Such bunching could result if the tax authority sets c values to perfectly match m, or indeed, if buyers and sellers anchor on c when negotiating m. However, we do not directly observe m, only p, a noisy proxy of transaction prices. This leads to a confound since the excess bunching mass is not due to under-reporting in this case; rather it comes from the incorrectly measured counterfactual p which differs from m. In our simulation, we assume that the distribution of reported values around government assessed values ($\frac{r-c}{c}$) takes exactly the same form in both cases to emphasize the point that bunching at c alone does not perfectly identify under-reporting behavior. Appendix B describes the inputs to the simulation in greater detail.

This confound is important since it is essentially impossible to obtain administrative data on "true" transaction prices given incentives to under-report. This forces reliance on potentially noisy proxies such as the Propequity measure that we use. The impossibility of measuring true m also means that we cannot ex-ante determine the size of the measurement error problem. Hedonic models are not a solution here either, since they are generally estimated using r or some other price contaminated by reporting incentives. In the next subsection, we propose two ways to correct for this potential confound.

5.2 Estimating Under-reporting in the Presence of Measurement Error in p

5.2.1 Aggregated Reported and Market Values

The first approach to distinguish the case with measurement error from genuine under-reporting involves inspecting the behavior of *aggregated* reported and market values by $\frac{r-c}{c}$ bins. Figures 1b and 1d plot simulated *aggregated* reported and market values for two different cases, namely, measurement error in p and truthful reporting when m bunches at c; versus under-reporting of r at c with p perfectly tracking m. In these figures, we aggregate the transactions into $\frac{r-c}{c}$ bins of 0.02 width.

In both figures, the blue bar labeled 0 on the x-axis indicates the total reported value of transactions (i.e., the sum of all r) with $-.02 \leq \frac{r-c}{c} < .02$, and the green bar is the total market value of transactions (i.e., the sum of all m). In Figure 1d, we also plot red bars, which show the total measured value of all transactions (i.e., the sum of all p) within each bin. (The assumption in Figure 1b is that p perfectly measures m, so the green bars in that figure capture the behavior of both m and p.)

In Figure 1b, with high under-reporting, but without measurement error, we can clearly see under-reporting in the zero bin, where aggregated r is substantially lower than aggregated m. Figure 1d shows the corresponding aggregates for the zero under-reporting with measurement error case, where we see that aggregated m and aggregated noisy p are (approximately) the same within each bin. Intuitively, aggregation within bins smooths out (symmetric) measurement error in p, and allows us to observe differences between r and m, allowing us to distinguish between the two cases. Truthful reporting with measurement error will exhibit no mass difference between p and r, and under-reporting will result in a mass difference between p and r in the central bin.[19] Figure 1f shows that with both high under-reporting and measurement error, aggregation within bins smooths out measurement error and still allows for a direct comparison of the average differences between r and m.

This approach also allows us to answer the extent of measurement error needed in our p measure to produce Figure 3, i.e., the plots *without aggregation*, even with no under-reporting. To do so, we assume that $m_i = r_i$, and simulate Propequity prices $p_i = m_i + \epsilon_i$. We find that p that are on average 15% above m, with a standard deviation

[19] Wider bins can, to a point, help us to smooth measurement error even further and better identify under-reporting at the expense of moving away from the sharp point at $r = c$. Panels A and B of Appendix Figure A.1 present aggregated estimates with bin-width set to 4% and 8% respectively, and confirms this in the data.

of 15% can replicate Figure 3 even with no under-reporting. This result illustrates the challenge in estimating under-reporting based purely on inspecting bunching in r around c versus that in p.

Turning to the actual data, Figure 4 aggregates reported and Propequity values in the data by $\frac{r-c}{c}$ bins. The figure shows that the largest magnitude of unreported value (the differential between aggregated r and p) comes from transactions that bunch at c, which is consistent with significant under-reporting for r reported to be exactly equal to c. The green dots indicate the total amount of Propequity estimated value (p) transacted amongst transactions within a .02 $\frac{r-c}{c}$ bin. The blue dots indicate the total value reported (r) for the same set of transactions within the $\frac{r-c}{c}$ bin. It is worth noting that with the aggregation procedure, the overall overlap of the p and r distributions is tighter than in the pure bunching plot in Figure 3, with the sharpest deviation evident between the distributions at exactly c. This pattern is consistent with the "two-type" model that we develop, which predicts an atom of mass in the distribution of r at c and a smooth distribution that is closer to p otherwise (with the mass shift from both the right and left of c, but for different reasons, i.e., appeal transactions costs from the left, and under-reporting from the right).

For bins with $\frac{r-c}{c} < 0$ (i.e. transactions with reported values less than government assessed values) the blue circles aggregate c. This is because if $r < c$ the tax base is effectively c, i.e., the government assessed value, since the government assesses taxes at c pending a successful appeal. Figure A.2 replaces the tax base value within bins (blue circles) in Figure 4 with the total reported value within bins. To the right of zero, these figures are identical. However, to the left of zero, the blue "+" points do not aggregate the circle value c, but rather, sum the reported value r. The total aggregated r in the bins immediately to the left of zero are lower than that implied by p, consistent with a low probability of successful appeals (i.e., π_2 is low), leading buyers in this range to bunch at c.

The figure also reveals that there is a monotonic decline in the gap between m and r as r increases to the right of c. For buyers who report more than 50% above the government assessed value, there is no visible gap between aggregate reported value r and estimated market values p. This is consistent with the theoretical prediction that those reporting over the government assessed value will report ever-closer to the market value as r increases. Furthermore, the agreement of the Propequity and reported values for those transactions with $\frac{r-c}{c} > 0.5$ provides additional confidence that the level of Propequity values are reasonable proxies for true market prices. Taken together, when we aggregate the rupee value of under-reporting across all bins in Figure

4, we estimate ₹455.1 billion (US$9.1 billion) in under-reported real estate value in our sample over the period 2013-2022.

Figure 5 directly shows the fraction of estimated market value under-reported by $\frac{r-c}{c}$ bin. Transactions that reported between 0 and 50 percent less than the government assessed value have an estimated under-reporting rate of about 20%. Note that these transactions are measured as reporting the government assessed value (because when $r < c$ the tax base is assumed to be c).[20] We observe a sharp discontinuity at zero and a marked change in slope as we move to transactions that report above the government assessed value, with our under-reporting estimates dropping to 0% for transactions that reported values 50% above the government assessed value. For properties that report 75% over the government assessed value, we estimate negative under-reporting rates, although these estimates are based on a relatively small number of transactions (the counts per bin in this region are in Figure 3).

5.2.2 Stale Government Assessed Values

A second approach to dealing with the effects of measurement error in p relies on a more careful inspection of the process used to set government assessed values c. For measurement error to explain our results rather than under-reporting, c must perfectly track m across both time and space. As we describe below, however, c is set in a geographically coarse manner, and infrequently updated. Over the sample period, real estate prices have grown substantially, and in a manner that varies across regions, meaning that it is implausible that c perfectly tracks m given the process of setting c.

Even without considering time-variation, c is set at a relatively broad geographical level. Within regions, therefore, given the considerable spatial variation of property values, it is implausible that c perfectly tracks m. In the data, c per sq. meter values are very close or the same for all properties within each sub-zone (the average (median) subzone in the data is 686,818 (264,136) sqm), and sub-zones often contain multiple CTS within them—in our sample, the average (median) number of CTS within a sub-zone is 213 (64) (average (median) CTS in the data is 3,000 (365) sqm).[21] This means that a single guidance value for a large region is unlikely to be an accurate reflection of

[20] One caveat here about these lower values is that it may be that for some of these transactions, $m < c$ and buyers disputed the accuracy of the government-issued circle rate c, so the under-reporting rate at zero (i.e., $r = c$) is potentially more accurate as it is less subject to this confound.

[21] The guidance values do incorporate some adjustments for whether a building is categorized as luxury, the floor the apartment is located in, and whether a parking space is included; but these are all categorical adjustments that are essentially swamped by the price variation within locations across buildings.

the full distribution of the true value of the assets in that region at any given point in time.

Despite this issue, it might still be the case that c values are set carefully to match the first spatial moment of m. This leads to a second problem, since the mass of transactions happening at prices above (infrequently updated) c will rise or fall over time as house prices grow or shrink on average given regional and aggregate price variation. If reporting is truthful, with such time-variation we would expect to find the greatest bunching of r at c immediately *after* c values are updated (i.e., when c is closest to m), and a gradual decline in bunching as m drifts away from c before c is updated again. If counterparties anchor m at government-determined c, a similar prediction obtains, as the accuracy and relevance of c might be expected to be highest immediately after it is updated. Moreover, infrequent updates in the presence of anchoring can create incentives for sellers to wait for c values to increase, as it could allow them to negotiate for substantially higher prices (in the three years we observe guidance value changes the average increase were 14.4 %, 10.5 %, and 6.98 %, see Table A.2).[22]

In contrast, if buyers and sellers seek to under-report to evade taxes, given that the dates of c changes are publicly announced in advance, we would expect that bunching would be greatest immediately *before* any increases in the government-assessed value. Put differently, in our model, if c predictably increases, there is a predictable jump in the tax burden incurred by under-reporting the transaction value $r = c$ immediately after the rise in c as opposed to immediately before the rise in c, delivering a strong incentive to under-report prior to the change in c.

In the data, Figure 6b documents behavior that is consistent with the under-reporting explanation, and inconsistent with either the measurement error or anchoring explanations. The figure shows that the largest extent of bunching occur immediately prior to scheduled guidance value changes, as c becomes "stale". This pattern is more consistent with the notion of buyers timing transactions to ensure that any potential change in c does not result in increased transactions taxes for the same asset. It is not consistent with c being adjusted to perfectly match m, which predicts greater bunching after c changes versus before. Nor is it consistent with buyers and sellers anchoring on c—which would predict no variation in the degree of bunching over time. Figure

[22] Even if sellers believe demand will be lower after c increases, they should be able to obtain some of the surplus generated by transacting at lower c values in the present by waiting and transacting at higher future c values tomorrow. Such arguments depend on discount rates and demand and supply elasticities and there are possibly constellations of parameter values that can deliver greater bunching prior to c value changes under truthful reporting. Ultimately, Occam's razor suggests that such arguments are potentially less plausible than under-reporting being higher amongst buyers who backdate transaction times to take advantage of infrequent updates.

A.3 uses the observed agreement date and the registration date for each transaction in the data to check for backdating behavior. The figure shows that there is a strong pattern of backdating agreement dates, further lending support to the under-reporting explanation.

Under-reporting appears to be higher in months prior to the circle rate changes especially in the early part of the sample, but it is difficult to draw strong conclusions relative to the month by month variation overall from pure visual inspection.[23] Table 3 estimates a regression model in which we explain the under-reporting rate using a time-trend and month of year fixed effects, and check whether the under-reporting rate varies in months prior to assessed value changes. The table shows that under-reporting rates are indeed approximately 6 % higher in months prior to changes in assessed values c. This is a large increase in under-reporting relative to the sample average under-reporting rate of 6%.

High-Frequency (Daily) Reporting Behavior

Figure 8 studies *daily* reporting behavior around scheduled guidance value changes (as indicated by the green vertical lines in 6a). Note that the dates of scheduled guidance value changes have moved around over time (i.e., they are not always on January 1st each year), which reduces concerns that these dates are scheduled for days with other policy announcements. Figure 8a shows a large spike in registered transactions on the day directly before the scheduled guidance value change. Figure 8b shows that the bunching rate for the large number of transactions registered right before the guidance value increase is approximately 10% higher. Figure 8c shows the fraction of transactions registered under the guidance value; this also shows a moderate increase in the days prior to the guidance value change. These results suggest that even non-bunching transactors expect the new guidance values to increase their tax burden (for example if the new guidance values are above what buyers planned to report).

5.3 Robustness to matching

To assuage concerns about whether IGR/Propstack and Propequity cover different types of properties, we also conduct all these tests restricting our sample for analysis to the 60% of transactions with an exact project match between the two datasets. This restricting the set of properties to those where we measure p only for those projects where we observe r. Appendix Figures A.13,A.14 and A.15 present counterpart results

[23] The circle rates were increased in 2014, 2015 and 2016. The Maharashtra government chose not to increase the guidance values in 2017 and 2018. See Table A.2.

to all of our aggregate analyses in Figures 3, 4 and 5. We find that the patterns and magnitudes are strongly consistent with the full sample, and the estimates are more precise, owing to a reduction in the extent of measurement error.

5.4 Formal Elasticity Estimation

In addition to our estimates of θ, and the corresponding under-reporting rate, we formally estimate the elasticity of reported value r to the transaction tax rate τ. We split our sample in to ten deciles based on guidance value c, and then separately estimate the elasticity for each group, thus capturing heterogeneity in the elasticity across the distribution of property prices. Appendix D describes the underlying method in detail.

Figure 7 presents estimated elasticities and associated confidence intervals, Table 2 presents the tabulated version of this figure, and Figure A.5 shows the corresponding bunching plots. For the lowest decile of c (median c in this decile is US$26,000), observed bunching translates to an elasticity of approximately 8, suggesting a very strong sensitivity to transaction tax rates for low value transactions (1 % increase in τ is associated with an 8% decrease in reported r). This elasticity means that an increase in τ from 0% to 5% would decrease reported values by 40%, consistent with the high under-reporting rates found in previous small sample estimates and common anecdotally. The second through fourth deciles, with guidance values between 26 and 46 thousand dollars have estimated elasticities between 1.87 and 3; transactions with guidance values above 64 thousand dollars have estimated elasticities less than 1.3.

These estimated elasticities are substantially larger than the structural elasticities for self-employed and wage earners using tax notches in Pakistan, where the estimated elasticities range from 0 to 0.28 (Kleven and Waseem 2013).[24] Our estimated elasticities are also large relative to medium and large sized South African small businesses, with estimated elasticities ranging between 0.23–0.27 (Anagol et al. 2022); the smallest businesses in the Anagol et al. (2022) sample have an estimated elasticity of 1.75, which is in the range that we find for medium and large size real estate transactions. These larger estimated elasticities may reflect the relative ease of under-reporting in our context (obtaining cash for the transfer and simply writing down a lower reported value) as opposed to corporate or personal income contexts where many transactions have to be adjusted to under-report overall income (and wage earner contexts where a worker must change jobs or adjust hours worked).

[24] We compare to Kleven and Waseem (2013)'s preferred estimates that exploit the fraction of taxpayers that report incomes in the "dead" zone above notches to correct elasticity estimates for frictions.

We note the following caveats with this analysis. First, we assume that the transaction tax rate is the main driver of behavioral responses; if capital gains or other taxes are also important drivers of r, then our elasticity estimates will be upward biased (as some of the bunching response would come from capital gains incentives as opposed to the transaction tax rate). Second, if there are non-tax benefits of under-reporting property values (i.e. money laundering incentives) then we will over-estimate the elasticity to transaction tax rate changes. Third, we do not consider extensive margin responses here (i.e. changes in the number of transactions occurring), which could also lead to a higher total elasticity of reported values to tax rate changes.

6 Applying the Method and Uncovering Heterogeneity in Under-reporting

Having developed our novel approach to detecting under-reporting, in this section, we apply the method in various ways. We begin by estimate the impacts of major government policy changes, most notably demonetization, on under-reporting rates. We then explore whether our baseline results are different when we focus on different segments of the residential property market in Mumbai. These results also allow us to better understand the underlying drivers of under-reporting behavior, including the interaction between financial constraints and under-reporting behavior.

6.1 Impact of India's Demonetization on Under-Reporting

We first apply our method of measuring under-reporting to evaluate a major program, with an important stated motive to eliminate unaccounted "black money" cash hoards, believed to frequently be employed in under-reported real estate transactions. On November 7, 2016 the Indian Prime Minister declared the ₹500 and ₹1,000 currency notes as no longer legal tender; these notes together comprised 86% of the nation's currency notes. Citizens would have approximately three months to deposit any of these currency units in banks; outside this window, these notes would be worthless. Banks were required to conduct audits on any deposits over ₹250,000, where the depositor was required to report the source of such cash holdings. A simple measure of the success of the policy, as reviewed in Lahiri (2020) and elsewhere, is that nearly 100% of the outstanding ₹500 and ₹1,000 rupee were ultimately deposited in banks. This suggests that the policy was largely unsuccessful at expropriating wealth from cash hoarders.

Chodorow-Reich et al. (2020) estimate the policy caused a 2% decline in GDP in the fourth quarter after implementation, with dissipated impacts after that.

As discussed earlier, many have argued that an important mechanism for under-reporting real estate values is the payment of cash from buyer to seller. Post-demonetization, therefore, with less cash available, we might expect a drop in under-reporting behavior. In the short-run, this could occur because transactions are delayed or abandoned completely as the cost of obtaining cash rises sharply. In the longer-run, agents may worry that a demonetization-like policy could be implemented again, or that demonetization sends a strong signal that the government will be cracking down on cash transactions. We might also therefore expect the trend of under-reporting behavior to change.

We begin by studying how the aggregate number and value of transactions changed in the weeks surrounding demonetization. We do not see a strong pattern of reduced transaction numbers or values in the weeks surrounding demonetization, despite the large macroeconomic shock it occasioned. The solid red line in Figure 6 indicates the demonetization month.

To further explore the impact of demonetization we test for heterogeneous impacts of demonetization along a dimension motivated by theory. We test whether the policy change differentially affected under-reporting for sales made by developers versus sales made by individuals. This analysis is motivated by the theory presented in Kleven, Kreiner and Saez (2016), which argues that it is more difficult for large organizations to maintain the collusive agreements that under-pin tax evasion because the probability of one "disgruntled" agent revealing the collusion increases with organization size. This model is apt for our setting, where developers would need to maintain the collusive agreement with many buyers, while re-sale transactions require a collusive agreement to be maintained between one buyer and one seller. We acknowledge that estimating the demonetization effects separately by developer versus re-sales is a heterogeneous treatment effect analysis where developer/re-sale status is not randomly assigned; in this sense these results should be interpreted as descriptive evidence on the heterogeneity of demonetization effects.

Figure 9b plots the monthly estimated under-reporting rates for re-sale transactions. As before, in a given month, the under-reporting rate is calculated as the difference between the total estimated market value of the transactions in that month (i.e., aggregate p) and the total tax base value (i.e., the monthly sum of $\max(r, c)$ for each individual transaction) expressed as a percentage of estimated market value. Once again, we aggregate to the month level to help with measurement error issues arising from the estimated market value appearing in the denominator. The figures shows that for re-sale transactions, the under-reporting rate dropped from approximately 10% prior to demonetization (leftmost, red vertical line) to about 5% afterwards. The drop in under-reporting rates amongst re-sale transactions could come from two sources: 1) demonetization increased reporting for similar transactions (an "intensive" margin effect) and 2) demonetization changed the composition of transactions towards properties more likely bought/sold by those less likely to under-report. Figure 9a plots the same under-reporting rates for developer sales; we observe no similar drop in under-reporting rates for these transactions.

6.2 Developer Regulation and Under-Reporting

There are two other major economic reforms during our sample period that this analysis allows us to investigate. The first is the passage of the Real Estate Regulation Act (RERA), which was a national law that was implemented ("notified") on April 18, 2017 (five months after the demonetization policy was announced). The broad goal of this act was to improve the functioning of the market for newly built apartment homes; the main provisions included requiring developers to set aside money in an escrow account to complete the building of real estate projects 2) requiring any project using 500 meters or more of land space or selling 8 or more units to register and provide updated data on completion times to MahaRERA (the newly established real estate regulator in Maharashtra), 3) procedures and time-lines for developers to respond to customer complaints.[25] Overall, the purpose was to curtail the ability of developers to sell properties before or during construction and then delay/abandon projects without reasonable compensation to buyers.

RERA did not include any specific provisions regarding the reported values of transactions within projects. We argue there are at least two plausible reasons RERA may affect under-reporting. First, if "fly-by-night" developers are also likely to be the ones who engage in collusive deals to under-report property values, it is plausible this regulation would reduce the amount of under-reporting through a change in the composition of developers selling properties. Second, RERA gave buyers stronger recourse in the case where a developer failed to deliver on time; to the extent that the financial compensation a developer pays is linked to the reported value, buyers have a stronger incentive to report the true values to obtain the full protective value of RERA. We expect RERA to primarily affect the market for new home sales; projects completed by

[25] For a full description of RERA provisions see: `https://maharerait.mahaonline.gov.in/PDF/FAQMergedPDF.pdf`.

May 1, 2017 were not affected by RERA.[26] Nonetheless, in Figure 9b panel (b) we do not see any major changes in estimated under-reporting rates in April/May of 2017 (demarcated by the rightmost, purple vertical line) for developer sales over and above the demonetization impacts.

A second potentially relevant policy change was implemented on July 1, 2017, namely, the introduction of the national Goods and Services Tax (GST), which centralized value-added-tax system administered at the central level.[27] It is possible that this policy created "input tax credits" for major construction supplies such as steel and cement relative to the previous fragmented VAT system, which means developers should report the cost of these inputs to the government. This paper trail of input costs may make it more difficult to under-report new sales of real estate. Again, this incentive primarily affects the market for new home sales. Figure 9b panel (b), however, once again does not show major changes in under-reporting rates around the GST reform (middle, green vertical line).

We next turn to more broadly evaluating the under-reporting behavior of different sub-groups of agents and properties in the Mumbai real estate market.

6.3 Heterogeneity in Reporting Behavior

6.3.1 Developer vs. Resales

As mentioned earlier, the data allow us to classify transactions into primary sales made by developers, and those that occur in the secondary market, i.e., resale transactions. As before, given that firms interact with many buyers, there is greater risk that one buyer may whistle-blow the under-reporting behavior leading to an audit or detection (Kleven, Kreiner and Saez (2016) argues that under-reporting of wages is lower amongst large firms because of the greater risk of a disgruntled employee revealing the under-reporting collusion).

Figures 10a and 10b show the count of developer sales and resales by $\frac{r-c}{c}$ bin respectively. Figures 10c and 10d show the bin aggregated reported and market values for developer sales and resales. Figure 10f shows under-reporting rates by $\frac{r-c}{c}$ bin for developer sales and resellers. Both types of sales show substantial bunching of reporting behavior at c, and a similar pattern of under-reporting by $\frac{r-c}{c}$ bin. While our overall under-reporting rate is not solely driven by resale transactions, Figure 11e shows that

[26] Completion was determined by whether the project had received an "occupancy certification" from the local housing authority.

[27] See Panigrahi (2021) for a more detailed policy description.

there is substantially more bunching for resale than for developer driven transactions, consistent with Kleven, Kreiner and Saez (2016).

6.3.2 Heterogeneity by Transaction Amount

Are under-reporting rates higher for larger value transactions? Figure 11 presents corresponding figures on densities, bin-specific aggregates of reported versus market values, and under-reporting rates. We find that smaller sized transactions show a more substantial jump in the under-reporting rate when $r = c$, consistent with the idea that larger value transactions are harder to under-report.

6.3.3 Financial Constraints, Mortgages, and Under-reporting

Nearly 60% of all transactions in the Mumbai sample report values extremely close to our proxy p for market value. Translating this observation using our simple economic framework, we interpret this as non-bunching households believing that the probability of audit π_1 is larger than 25% (given that $n = 4$). This number appears high relative to anecdotal accounts of the rate of audit by the registrar, or indeed, the income tax authority. We posit therefore, that there are likely other important incentives that drive truthful reporting.

One important non-penalty driven incentive in our context is the desire to alleviate financial constraints when borrowing to fund a house purchase. Mortgage lending policies in India (and indeed in many other jurisdictions) often take the official reported value of a property into account when undertaking credit-screening, thus linking the decision to under-report with the extent of financial constraints. More specifically, in our context, many banks in India, including the largest national lender, the State Bank of India, are only prepared to lend up to the reported value r. One important incentive to truthfully report, therefore, is generated by the desire to unlock greater mortgage financing. To the extent that this incentive operates, we are likely to observe less under-reporting for borrowers that are financially constrained and require mortgage financing, and thus a reduction in the extent of observed bunching if our proxy is an accurate reflection of under-reporting behavior.

To pursue this intuition more formally, we extend our basic model framework to incorporate a penalty that increases with the tightness of the mortgage constraint, as in Andersen, Badarinza, Liu, Marx and Ramadorai (2021). Consider a bank that is willing to lend $(1-\gamma)r$, where γ is the down-payment constraint. Now consider a potential buyer of a property, who requires funding of $(m-d)$ where d is their available liquidity and m is the market value of the property. The shortfall $[(m-d)-(1-\gamma)r]$ can be overcome at a cost, which for simplicity, we model below as quadratic in the extent of the shortfall (in the Danish housing market, Andersen, Badarinza, Liu, Marx and Ramadorai (2021) find that this functional form offers a good approximation). Note that the down-payment constraint affects transactions irrespective of whether true $m > c$ or true $m \leq c$. In this simple setup, borrowers must report more (or scale down the desired house size) in order to alleviate the financial constraint, which scales with the size of the financial shortfall.

In the case where $m > c$, the borrower (potential buyer) chooses r to maximize the following problem:

$$U = -\tau\left[(1-\pi_1)r + \pi_1\left(r + n(m-r)\right)\right] - \frac{\mu}{2}\left[(m-d)-(1-\gamma)r\right]^2 \tag{3}$$

Here, the parameter μ determines the tightness of the financial constraint. Setting $\frac{dU}{dr} = 0$, and solving for r^*:

$$r^* = \frac{(m-d)}{(1-\gamma)} - \frac{\tau(1-n\pi_1)}{\mu(1-\gamma)^2} \tag{4}$$

Scaling both sides by m, we obtain:

$$\frac{r^*}{m} = \left(\frac{1}{1-\gamma}\right)\left(1-\frac{d}{m}\right) - \frac{\tau(1-n\pi_1)}{m\mu(1-\gamma)^2} \tag{5}$$

Note here that $1-\frac{d}{m}$ is the True LTV ($TLTV$) on the mortgage. Empirically, $\frac{d}{m}$ can be estimated using the proxy p for m. This relationship suggests that the incentives to alleviate the down-payment constraint can enable potential buyers to report r more truthfully, thus presenting a countervailing force to the tax incentives at play.

To test the predictions of this simple model, we require data on mortgages. In our setting, mortgages also have to be reported to the registrar, but to make progress, we must match transactions and mortgages, which are separately reported to the registrar. Using the administrative IGR data, we undertake this matching exercise, and we are able to match roughly 31,000 reported transactions to a mortgage (we describe this matching process in detail in Appendix Section C). While this permits us to analyze the relationship between r, c and mortgage values, unfortunately the intersection between this matched sample and the Propequity data (from where we obtain our p values) is low. Appendix Figure A.4 restricts the sample to the 8,913 IGR transactions for which we have both mortgage information and the Propequity value p. Nevertheless, when we plot the empirical relationship between $\frac{R}{P}$ and the true LTV bins as estimated using the Propequity value p, we confirm that the upward sloping relationship between reporting and $TLTV$ predicted by the model is clearly observed in the data. The estimated slope that we estimate is 0.65, and the regression has a high r-squared of 42%.

While the overlap with p is low, our analysis thus far has been useful at equating bunching behavior of r at c with under-reporting, and makes alternative explanations for observed bunching less likely. Using the larger 31,000 observation mortgage-transaction matched sample, therefore, we split out the bunching behavior of reported values based on whether the transaction has an associated mortgage in the registrar data. Figures 12a-12c show how the bunching of r at c varies mortgage and lender characteristics.

Figure 12a shows how bunching behavior varies with the loan-to-value ratio we estimate based on a set of matched mortgage transactions. The figure shows that transactions with progressively higher loan-to-value ratios tend to exhibit less bunching, and that this relationship is monotonic. This finding is consistent with the model's prediction that incentives to relax credit constraints (since higher reported values lead to the possibility of greater mortgage loans) cut against incentives for tax evasion. The magnitudes are sizeable— low loan-to-value loans are approximately 10 percentage points more likely to bunch than transactions associated with a high loan-to-value mortgages. Given that many of these transactions may have an (unmatched) mortgage, this difference is likely a lower bound. We view this result as consistent with the fact that most banks will only lend up to the amount that the buyer reports on the sales deed, so a buyer seeking credit has an additional incentive to report as each extra rupee reported allows for an additional rupee of borrowing. We also note that this result could also be driven by a negative correlation between preferences for tax evasion and credit constraints.

In Figure 12b, we split the matched mortgage-registered transaction sample based on the organizational structure of the lending bank (we observe the identity of the bank in the administrative data). We find that transactions associated with mortgages from cooperative banks demonstrate the greatest bunching, followed by banks which we were unable to perfectly classify, followed by public sector banks, and finally, the lowest levels of bunching are observed in private and foreign banks. This heterogeneity can be attributed to both self-sorting of different types of borrowers to different types of banks, as well as by borrowers under-reporting more or less depending on the lending bank's credit-screening policies. Sorting of borrowers across banks could be driven by borrowers choosing different types of banks, or by banks having different lending rules which ex-post lead to selection in the type of borrowers at different banks. Conditional on a borrower matching with a bank, the bank may also have different rules which encourage different reporting amounts for the same borrower. For example, some banks will only lend up to the reported value on the sales deed, while others will lend based on their own assessment. Overall, these results are consistent with bank culture being correlated with borrower type (see, for example, Mishra, Prabhala and Rajan (2022), who show that public-sector Indian banks appear to have laxer credit screening standards and slower technology adoption than private sector Indian banks, and link this finding to differences in organization culture).

Finally, Figure 12c presents bunching behavior for the same group of 31,000-odd transactions with mortgages, but splits the sample by the average non-performing loan (NPL) rate (over the period 2013-2018) of the lending bank to focus more closely on a possible credit-screening channel. The figure shows that loans issued by banks with the highest NPL rates also exhibit the greatest amount of bunching. This correlation could be generated by borrowers that are more likely to default also being the types who under-report, other types of selection correlated with banks' differential lending process, or borrower selection into banks that have more lax screening of borrowers and collateral.

6.4 Estimated Aggregate Revenue Losses Over Time

Before we conclude, we note that government stamp tax revenues are directly related to the under-reporting rate. If the true value of a property is v, and with a 13% under-reporting rate the reported value is $0.87v$, a stamp tax of 5% on the true transactions price effectively translates into 4.35% of this true price being received by the exchequer, i.e., a reduction of 13% in stamp tax revenues.

Using data on the aggregate tax revenues obtained from property registrations, Figure 13 estimates the loss of revenues attributable to under-reporting. If we take our estimated under-reporting rates from Mumbai as representative for the entire state of Maharashtra, this translates into a reduction of about US$450 million in tax revenues from property transactions (1% of total revenue) for the government for the year 2018. While this back-of-the-envelope calculation imposes several assumptions and is not precise, it provides a sense of the broader economic magnitude of tax losses arising from our under-reporting estimates.

6.4.1 High-Frequency (Daily) Reporting Behavior

Figure 8 studies *daily* reporting behavior around scheduled guidance value changes (as indicated by the green vertical lines in 6a). Note that the dates of scheduled guidance value changes have moved around over time (i.e., they are not always on January 1st each year), which reduces concerns that these dates are scheduled for days with other policy announcements. Figure 8a shows a large spike in registered transactions on the day directly before the scheduled guidance value change. Figure 8b shows that the bunching rate for the large number of transactions registered right before the guidance value increase is approximately 10% higher. Figure 8c shows the fraction of transactions registered under the guidance value; this also shows a moderate increase in the days prior to the guidance value change. These results suggest that even non-bunching transactors expect the new guidance values to increase their tax burden (for example if the new guidance values are above what buyers planned to report).

7 Conclusion

We develop a new method to detect under-reporting of asset values for tax purposes, and apply it to a large administrative dataset of reported property values for transactions tax purposes in Mumbai, India between 2013 and 2022. Our method reveals an overall 13% under-reporting rate over the period. Using monthly under-reporting rates, our back-of-the envelope calculation is that annual losses to the state government are on the order of 1% (2021 estimate) of total revenues from lost stamp duties on property value under-reporting. These estimates are meaningfully large for the public exchequer, but are nonetheless substantially lower than prior estimates using smaller data samples, and lower than anecdotal priors related to the purported prevalence of tax evasion and "black money" in emerging market real estate transactions. Over and above these empirical results, we view the methods that we develop as more widely applicable to a range of contexts in which the tax basis is set on the basis of government-assessed asset values.

We also find that under-reporting rates are significantly higher immediately before pre-announced changes in government-assessed values, which add credence to calls to index such policy instruments, which are frequently-used in emerging economy contexts, to more frequently updated measures of market value (Campbell, Ramadorai and Ranish 2015).

We also find evidence of a strong correlation between the degree of bunching of reported values at guidance values and features of the mortgage contract such as the LTV ratio on the loan, and the identity and financial health of the bank issuing the loan. This relationship is intriguing, and suggests a link between the quality and extent of financial screening and household incentives for tax evasion, a link we believe should be explored more carefully and fully going forward.

References

Agarwal, Sumit, Keyang Li, Yu Qin, Jing Wu and Jubo Yan. 2020. "Tax evasion, capital gains taxes, and the housing market." *Journal of Public Economics* 188:104222.

Agarwal, Sumit, Weida Kuang, Long Wang and Yang Yang. 2020. "The Role of Agents in Fraudulent Activities: Evidence from the Housing Market in China.".

Ahlfeldt, Gabriel M, Stephen J Redding, Daniel M Sturm and Nikolaus Wolf. 2015. "The economics of density: Evidence from the Berlin Wall." *Econometrica* 83(6):2127–2189.

Amornsiripanitch, Natee. 2020. "Why Are Residential Property Tax Rates Regressive?" *Available at SSRN 3729072* .

Anagol, Santosh, Benjamin B Lockwood, Allan Davids and Tarun Ramadorai. 2022. "Diffuse Bunching with Frictions: Theory and Estimation.".

Anagol, Santosh, Fernando V Ferreira and Jonah M Rexer. 2021. Estimating the Economic Value of Zoning Reform. Technical report National Bureau of Economic Research.

Andersen, Steffen, Cristian Badarinza, Lu Liu, Julie Marx and Tarun Ramadorai. 2021. "Reference dependence in the housing market." *Available at SSRN 3396506* .

Artavanis, Nikolaos, Adair Morse and Margarita Tsoutsoura. 2016. "Measuring income tax evasion using bank credit: Evidence from Greece." *The Quarterly Journal of Economics* 131(2):739–798.

Avenancio-León, Carlos F and Troup Howard. 2022. "The assessment gap: Racial inequalities in property taxation." *The Quarterly Journal of Economics* 137(3):1383–1434.

Basu, Kaushik. 2015. *An economist in the real world: the art of policymaking in India*. MIT Press.

Best, Michael Carlos and Henrik Jacobsen Kleven. 2018. "Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK." *The Review of Economic Studies* 85(1):157–193.

Braguinsky, Serguey, Sergey Mityakov and Andrey Liscovich. 2014. "Direct estimation of hidden earnings: Evidence from Russian administrative data." *The Journal of Law and Economics* 57(2):281–319.

Campbell, John Y, Tarun Ramadorai and Benjamin Ranish. 2015. "The impact of regulation on mortgage risk: Evidence from India." *American Economic Journal: Economic Policy* 7(4):71–102.

Casaburi, Lorenzo and Ugo Troiano. 2016. "Ghost-house busters: The electoral response to a large anti–tax evasion program." *The Quarterly Journal of Economics* 131(1):273–314.

Chang, Yun-chien. 2012. "Self-Assessment of Takings Compensation: An Empirical Study." *The Journal of Law, Economics, & Organization* 28(2):265–285.

Chetty, Raj, John N Friedman, Tore Olsen and Luigi Pistaferri. 2011. "Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records." *The quarterly journal of economics* 126(2):749–804.

Chodorow-Reich, Gabriel, Gita Gopinath, Prachi Mishra and Abhinav Narayanan. 2020. "Cash and the economy: Evidence from India's demonetization." *The Quarterly Journal of Economics* 135(1):57–103.

Coffee Jr, John C. 2002. "Racing towards the top: The impact of cross-listing and stock market competition on international corporate governance." *Colum. L. Rev.* 102:1757.

Comptroller and Auditor General of India. 2016. Report of the Comptroller and Auditor General of India on Revenue Sector for Year Ended 21 March 2015. Technical report.

Dachis, Ben, Gilles Duranton and Matthew A Turner. 2012. "The effects of land transfer taxes on real estate markets: evidence from a natural experiment in Toronto." *Journal of economic Geography* 12(2):327–354.

Fan, Ziying, Xiaxin Wang and Hang Zhang. 2022. "Understanding Misreporting: New Evidence from Responses to a Housing Transaction Tax.".

Fisman, Raymond and Shang-Jin Wei. 2004. "Tax rates and tax evasion: evidence from "missing imports" in China." *Journal of political Economy* 112(2):471–496.

Gechter, Michael and N Tsivanidis. 2020. "Spatial spillovers from urban renewal: evidence from the Mumbai Mills Redevelopment." *Penn State University (mimeo)* .

Genesove, David and Christopher Mayer. 2001. "Loss aversion and seller behavior: Evidence from the housing market." *The quarterly journal of economics* 116(4):1233–1260.

Gopalakrishnan, S., Das-Gupta A. 1986. "Stamp Duties and Registration Fees in West Bengal." Unpublished Manuscript.

Gordon, Roger and Wei Li. 2009. "Tax structures in developing countries: Many puzzles and a possible explanation." *Journal of public Economics* 93(7-8):855–866.

Griffin, John M and Gonzalo Maturana. 2016. "Who facilitated misreporting in securitized loans?" *The Review of Financial Studies* 29(2):384–419.

Harari, Mariaflavia, Maisy Wong et al. 2018. Slum upgrading and long-run urban development: Evidence from indonesia. In *2018 Meeting Papers. Society for Economic Dynamics*.

Harberger, Arnold C. 1965. *Issues of tax reform for Latin America*. PAU, Department of Economic Affairs.

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

Kleven, Henrik J and Mazhar Waseem. 2013. "Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan." *The Quarterly Journal of Economics* 128(2):669–723.

Kleven, Henrik Jacobsen. 2016. "Bunching." *Annual Review of Economics* 8:435–464.

Kleven, Henrik Jacobsen, Claus Thustrup Kreiner and Emmanuel Saez. 2016. "Why can modern governments tax so much? An agency model of firms as fiscal intermediaries." *Economica* 83(330):219–246.

Kleven, Henrik Jacobsen, Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen and Emmanuel Saez. 2011. "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark." *Econometrica* 79(3):651–692.

Kopczuk, Wojciech and David Munroe. 2015. "Mansion tax: The effect of transfer taxes on the residential real estate market." *American economic Journal: economic policy* 7(2):214–57.

La Porta, Rafael, Florencio Lopez-de Silanes and Andrei Shleifer. 2002. "Government ownership of banks." *The Journal of Finance* 57(1):265–301.

Lahiri, Amartya. 2020. "The great indian demonetization." *Journal of Economic Perspectives* 34(1):55–74.

Mishra, Prachi, Nagpurnanand Prabhala and Raghuram G Rajan. 2022. "The Relationship Dilemma: Why Do Banks Differ in the Pace at Which They Adopt New Technology?" *The Review of Financial Studies* .

Montalvo, José G, Amedeo Piolatto and Josep Raya. 2020. "Transaction-tax evasion in the housing market." *Regional Science and Urban Economics* 81:103526.

National Institute of Public Finance and Policy. 1995. "Black Money in the Real Estate Sector: A Study.".

OECD. 2007. "Report on tax fraud and money laundering vulnerabilities involving the real estate sector.". OECD Centre for Tax Policy and Administration.

Panchapagesan, V., Karthik N. 2017. "Revenue Neutral Approach to Lower Stamp Duty and Registration Charges for Affordable Housing." Unpublished Manuscript.

Panigrahi, Piyush. 2021. "Endogenous Spatial Production Networks: Quantitative Implications for Trade and Productivity." *Available at SSRN 3982029* .

Piketty, Thomas. 2015. "Capital and wealth taxation in the 21 st century." *National Tax Journal* 68(2):449–458.

Piskorski, Tomasz, Amit Seru and James Witkin. 2015. "Asset quality misrepresentation by financial intermediaries: Evidence from the RMBS market." *The Journal of Finance* 70(6):2635–2678.

Pissarides, Christopher A and Guglielmo Weber. 1989. "An expenditure-based estimate of Britain's black economy." *Journal of public economics* 39(1):17–32.

Pomeranz, Dina and José Vila-Belda. 2019. "Taking state-capacity research to the field: Insights from collaborations with tax authorities." *Annual Review of Economics* 11:755–781.

Posner, Eric A and E Glen Weyl. 2019. *Radical Markets*. Princeton University Press.

Saez, Emmanuel. 2010. "Do taxpayers bunch at kink points?" *American economic Journal: economic policy* 2(3):180–212.

Slemrod, Joel. 2007. "Cheating ourselves: The economics of tax evasion." *Journal of Economic perspectives* 21(1):25–48.

Stiglitz, Joseph E. 2015. "The origins of inequality, and policies to contain it." *National Tax Journal* 68(2):425–448.

Tandon, N.K. 1987. "Under-reporting of Real Estate Transactions in India." Unpublished Manuscript.

Tsivanidis, Nick. 2019. "Evaluating the impact of urban transit infrastructure: Evidence from bogota's transmilenio." *Unpublished manuscript* .

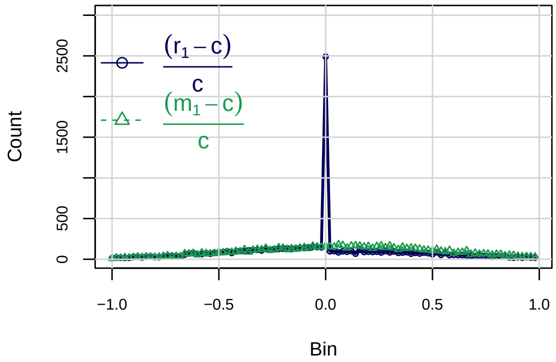

(a) High under-reporting, without measurement error

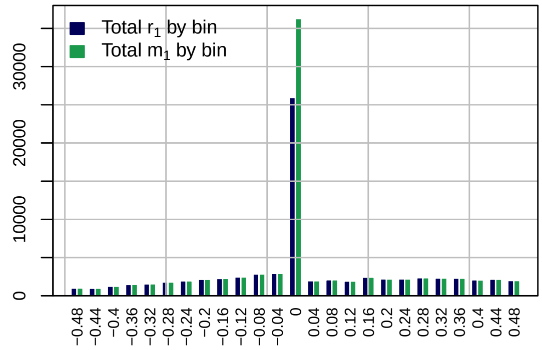

(b) High under-reporting, without measurement error

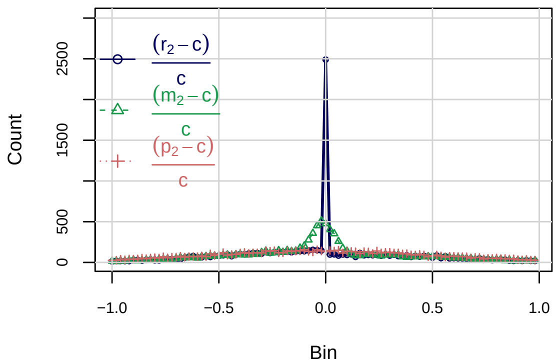

(c) No under-reporting, with measurement error

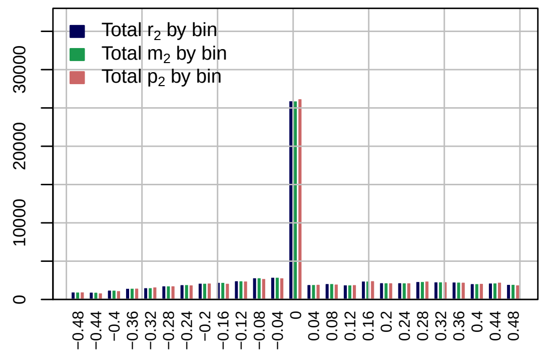

(d) No under-reporting, with measurement error

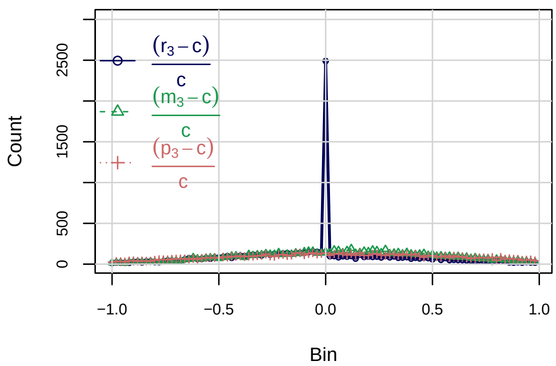

(e) High under-reporting, with measurement error

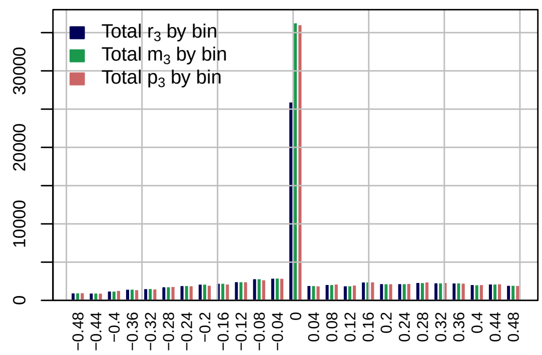

(f) High under-reporting, with measurement error

Figure 1
Simulation Results

r_1 and m_1 are reported and transaction values in the high under-reporting case. r_2, m_2 and p_2 are reported, true transaction, and noisily measured transaction price variables for the no under-reporting case. r_3, m_3 and p_3 are reported, true transaction, and noisily measured transaction price variables for the high under-reporting case with measurement error. c is the government assessed value. Measurement error refers to noise in our estimates of market prices relative to the true unobserved market price.

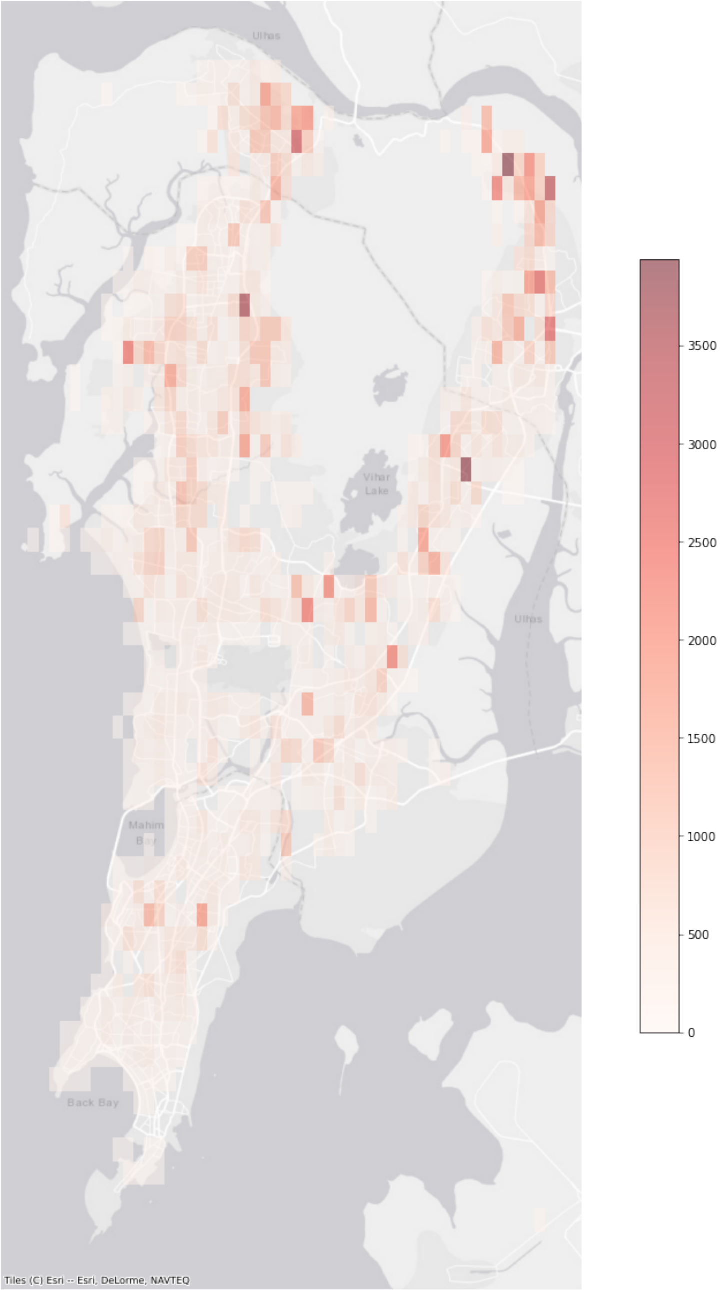

Figure 2
Geographic Distribution of CTS in our sample

The areas marked in red on the Mumbai map marks those CTS in our matched sample with the Propequity data.

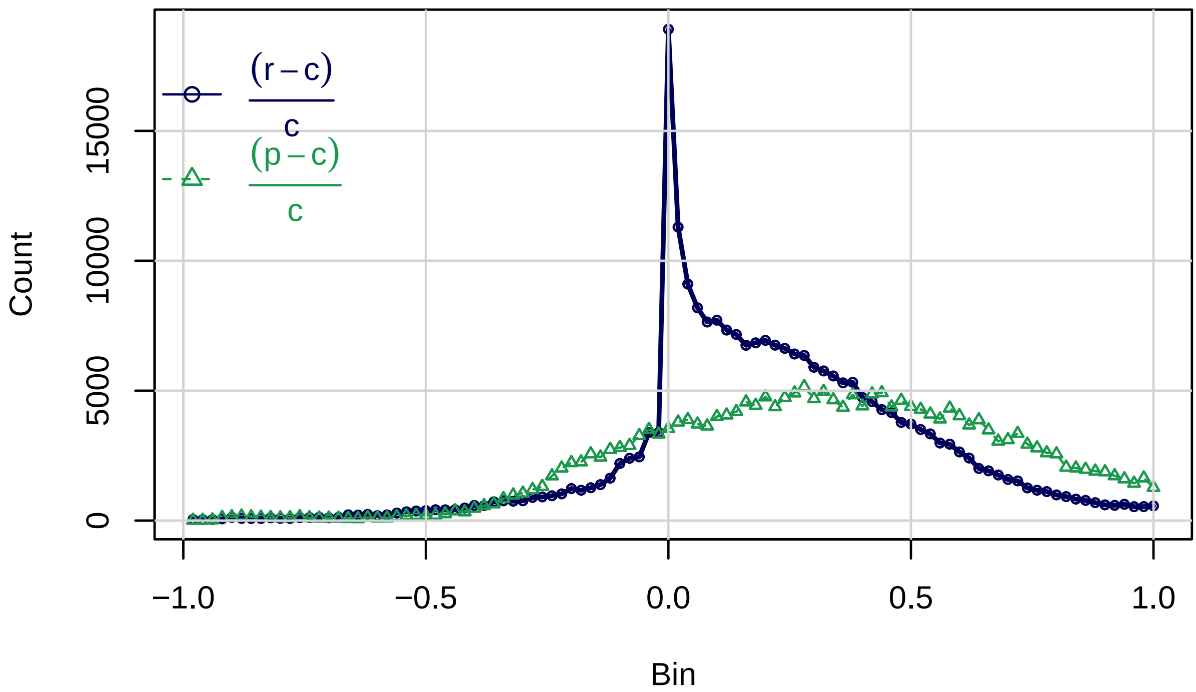

Figure 3
Bunching of Reported and Propequity Values Around Circle Values

The blue line shows the distribution of reported values across 2% reported value bins, where a reported value bin is measured as a deviation from the government assessed value. The green line shows the distribution our noisily measured estimated of the market price (the Propequity values) for the same underlying set of transactions reported in the blue line.

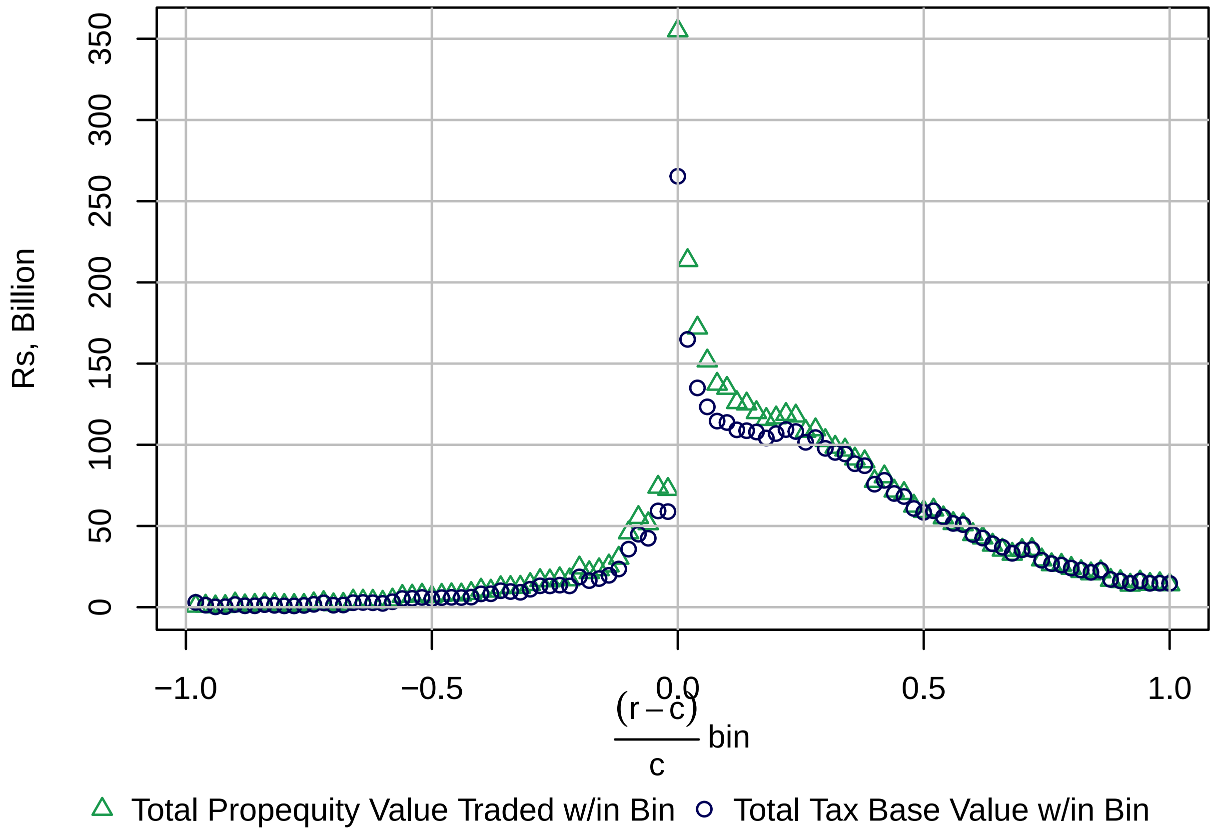

Figure 4

Aggregate Taxbase and Propequity Values by Reporting Behavior Bins

The blue circles show the aggregated taxbase value within 2% reported value bins, where a reported value bin is measured as a deviation from the government assessed value. The green triangles show the aggregate noisily measured estimated of the market value (the Propequity values) for the same underlying set of transactions reported in the green triangles.

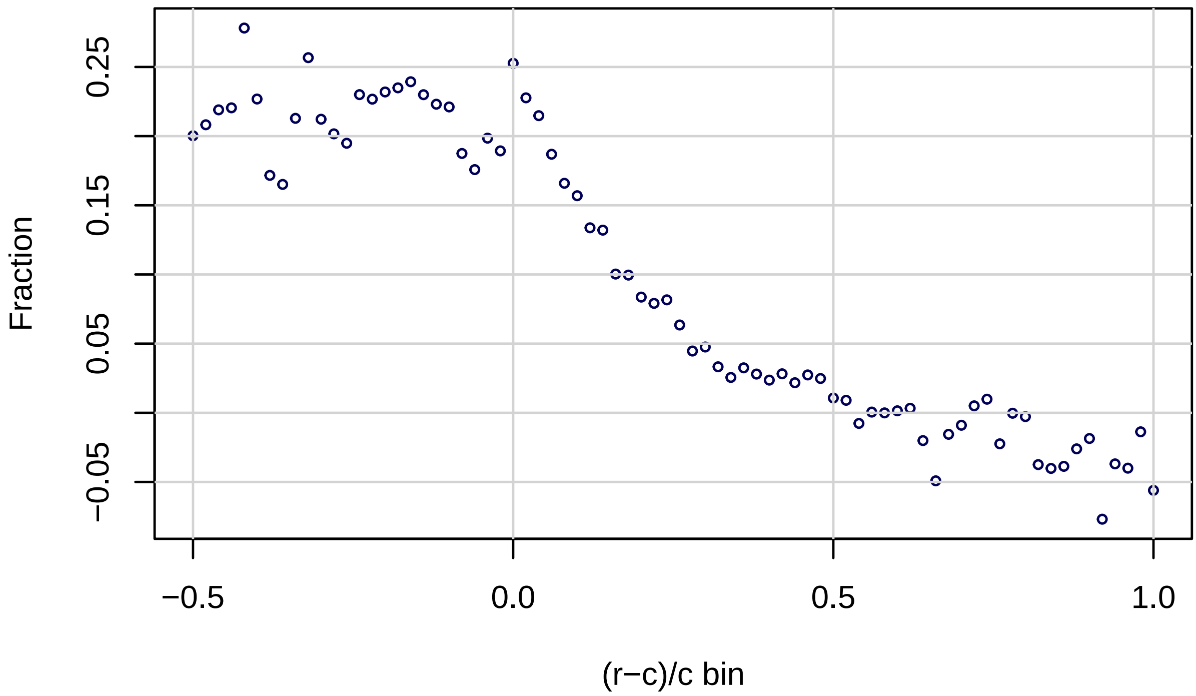

Figure 5
Under-Reporting Rate by Reporting Behavior Bins

The blue circles show the estimated under-reporting rate within 2% reported value bins, where a reported value bin is measured as a deviation from the government assessed value.

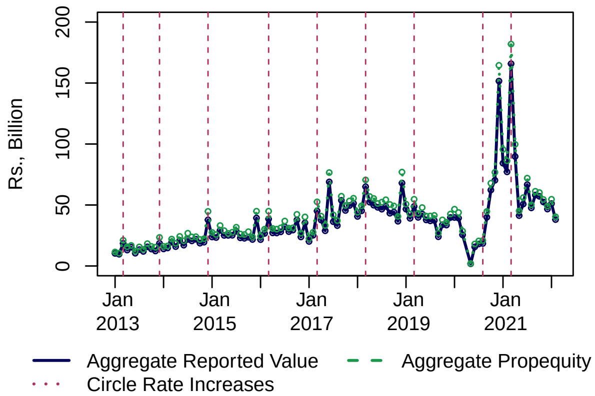

(a) Monthly Aggregates

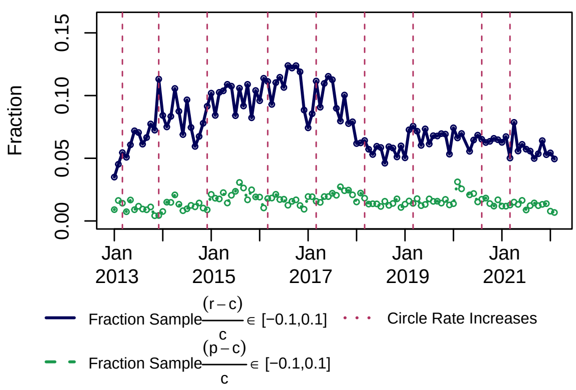

(b) Monthly Bunching

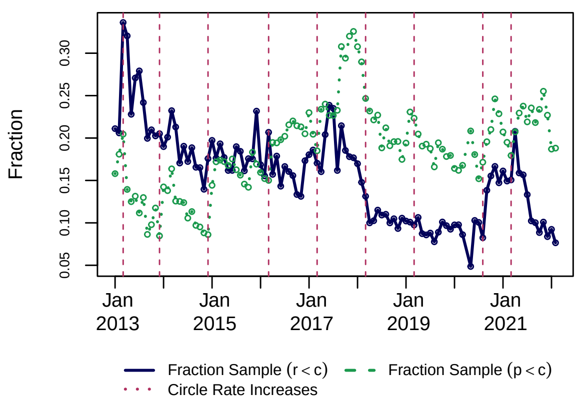

(c) Reporting/Market Values < Guidance Values

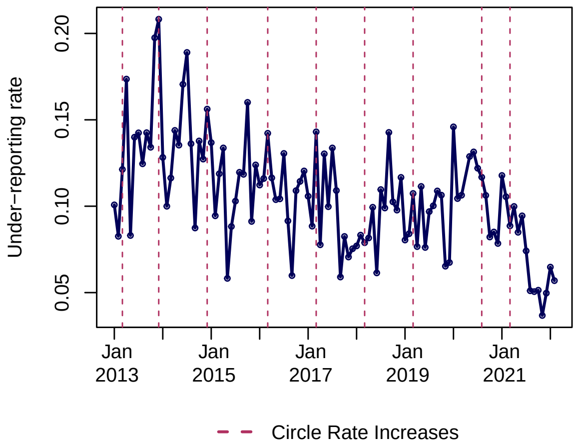

(d) Monthly Under-Reporting Rate

Figure 6
Heterogeneity Over Time

Red vertical dashed lines refer to scheduled circle rate increases. The circle rates were increased in 2013, 2014 and 2015; The government kept circle rates the same in 2016, 2017 and 2018. It was brought down during the COVID-19 relaxation of the transaction tax rate and then subsequently increased in 2021.

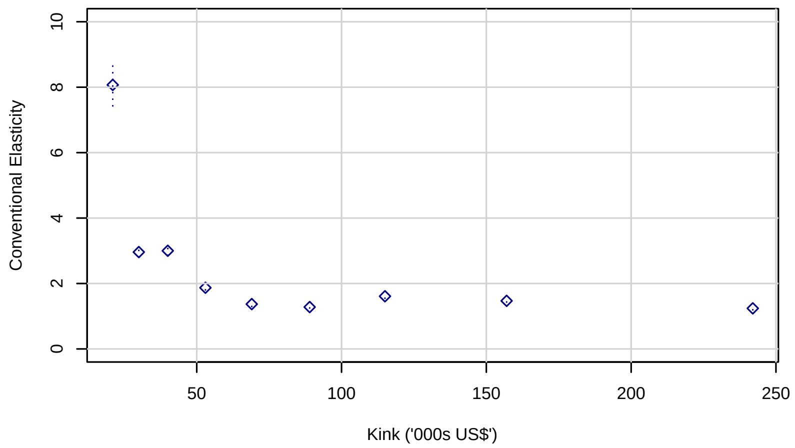

Figure 7
Reported Value Elasticity to Transaction Tax Rate

This figure plots the reported value elasticity to transaction tax rate by deciles of the guidance value distribution. These estimates are presented in Table 2 of the paper.

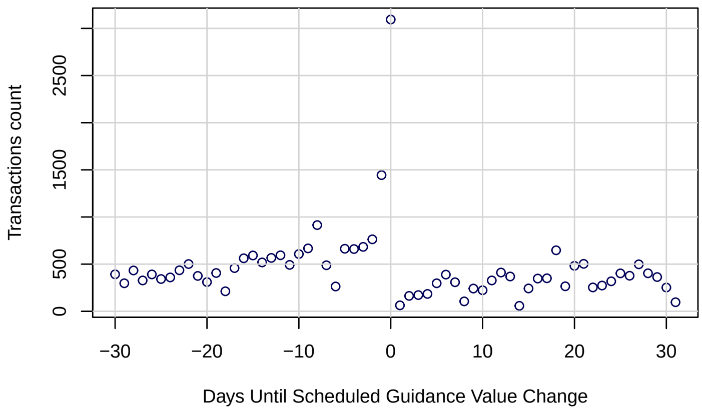

(a) Transaction Counts

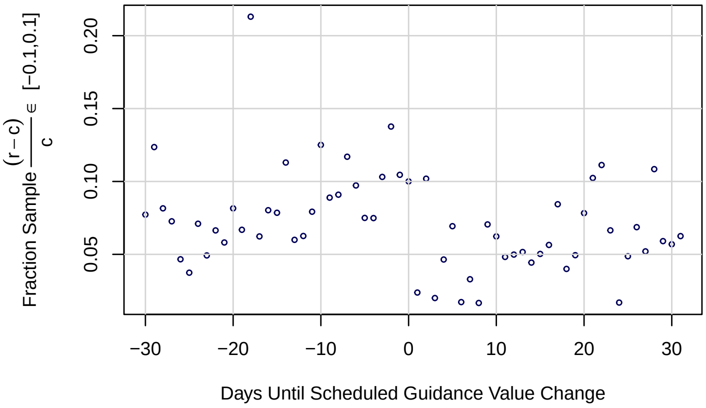

(b) Fraction Bunching

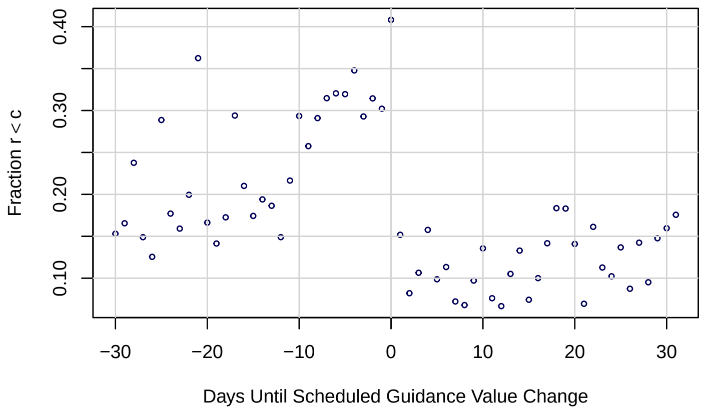

(c) Fraction $r < c$

Figure 8

Reporting Behavior in Days Around Scheduled Guidance Value Increases

Panel (a) shows counts of transactions made per day within a 30 day window of all the scheduled guidance value changes that occurred over our sample period. Panel (b) shows the fraction of the transactions on the given event-day that "bunchers", i.e. had a reported value within 1% of the guidance assessed value. Panel (c) shows the fraction of transactions that had a reported value less than the guidance assessed value.

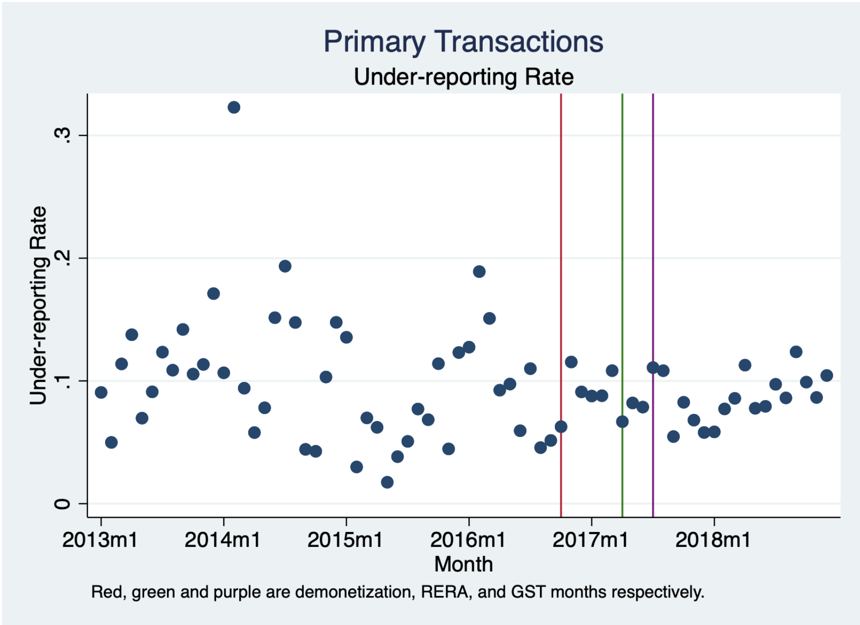

(a) Primary Transactions

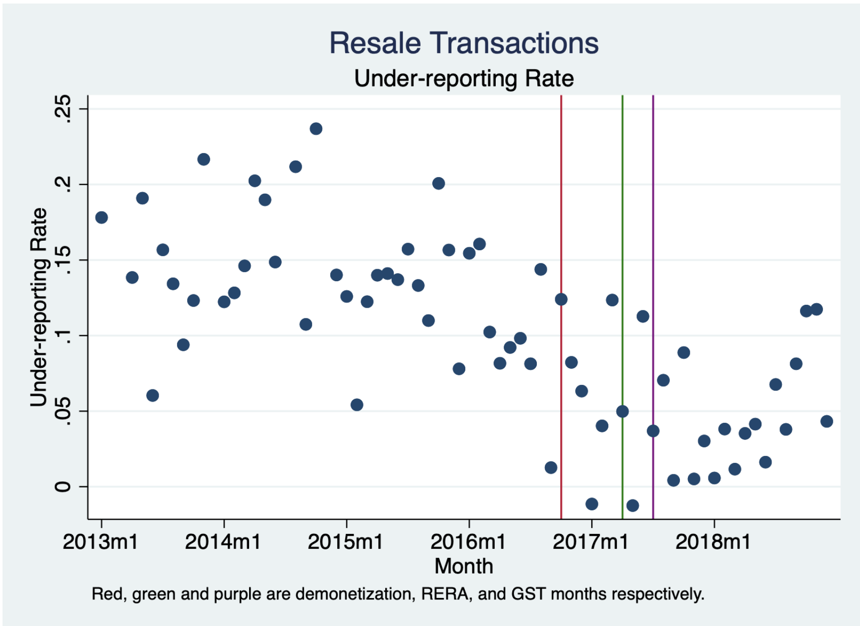

(b) Resale Transactions

Figure 9
Demonetization Effects

This figure plots monthly under-reporting rate for primary and resale transactions in panels (a) and (b) respectively. The red, green and purple lines denote demonetization, introduction of the real estate regulatory authority and goods and services tax in the state of Maharashtra, respectively.

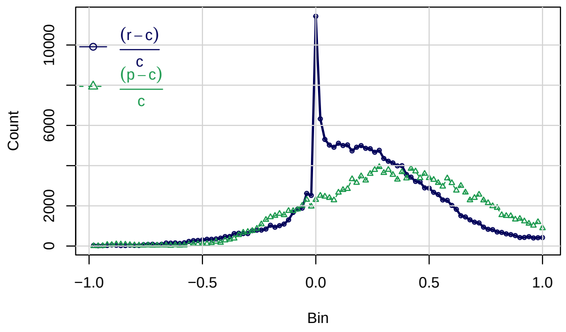

(a) Developer Sales

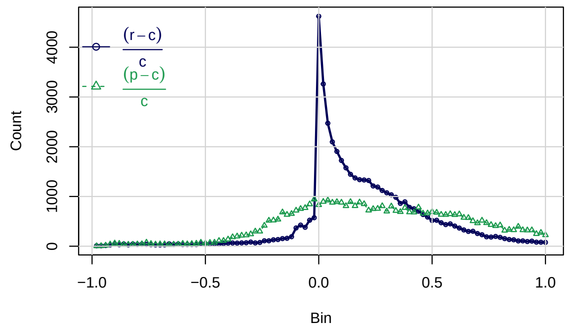

(b) Resale

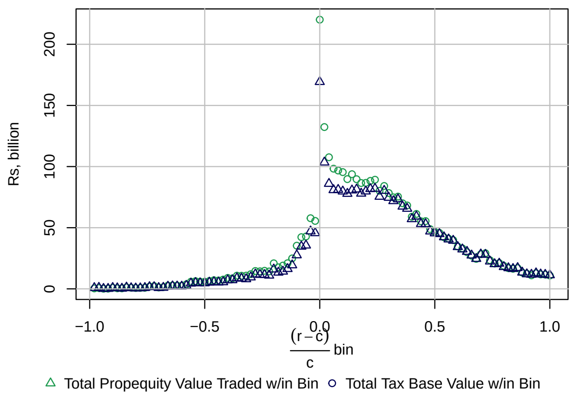

(c) Developer Sales

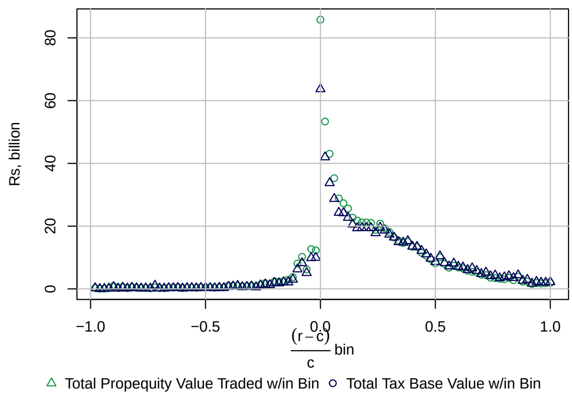

(d) Resale

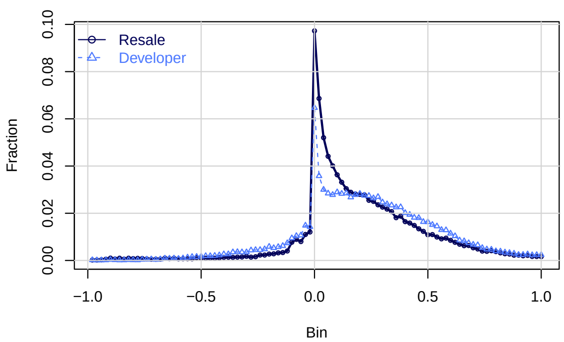

(e) Reported Counts (Density) Comparison

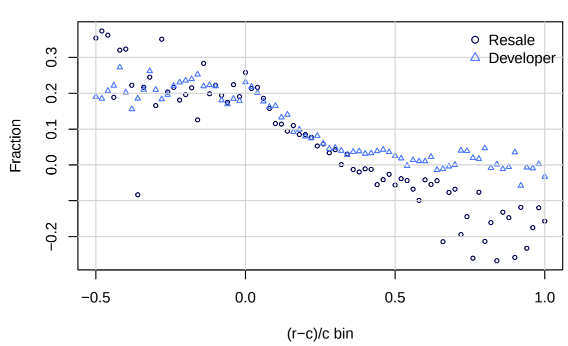

(f) Under-Reporting Rates

Figure 10
Firm vs. Non-Firm Seller Heterogeneity

See Figures 3, 4 and 5 for detailed descriptions.

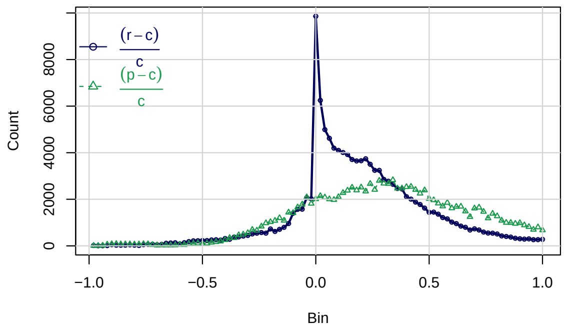

(a) Above Median Circle Value

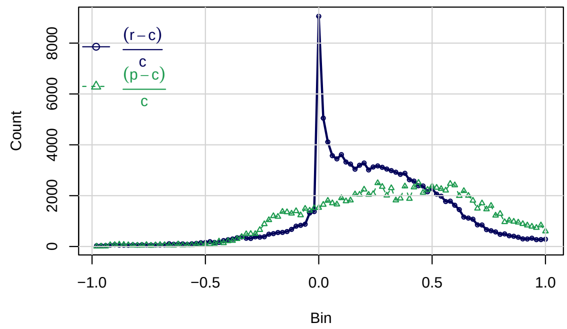

(b) Below Median Circle Value

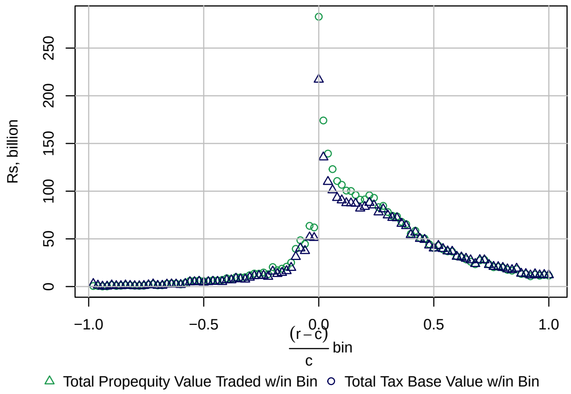

(c) Above Median Circle Value

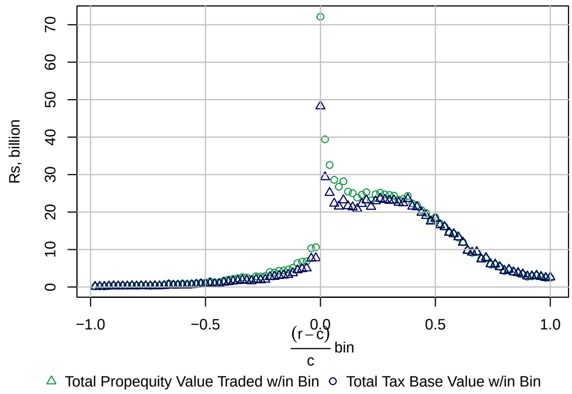

(d) Below Median Circle Value

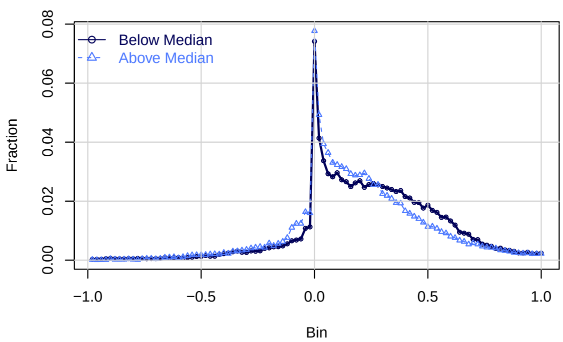

(e) Reported Counts (Density) Comparison

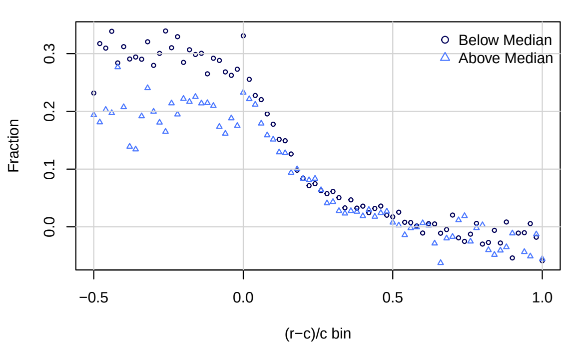

(f) Under-Reporting Rates

Figure 11
Above vs. Below Median Circle Value Heterogeneity

See Figures 3, 4 and 5 for detailed descriptions.

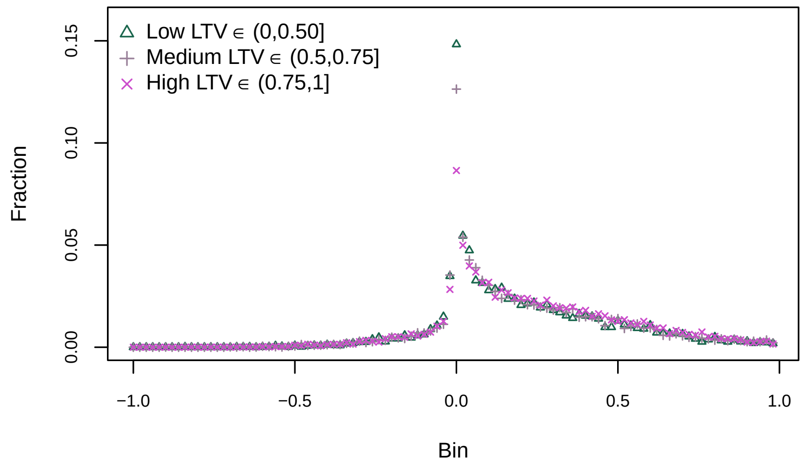

(a) Loan-to-Value Bunching Heterogeneity

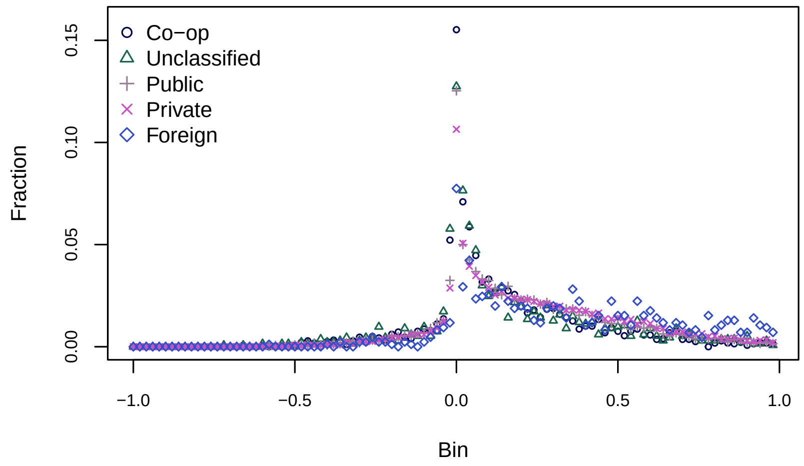

(b) Bank Type Bunching Heterogeneity

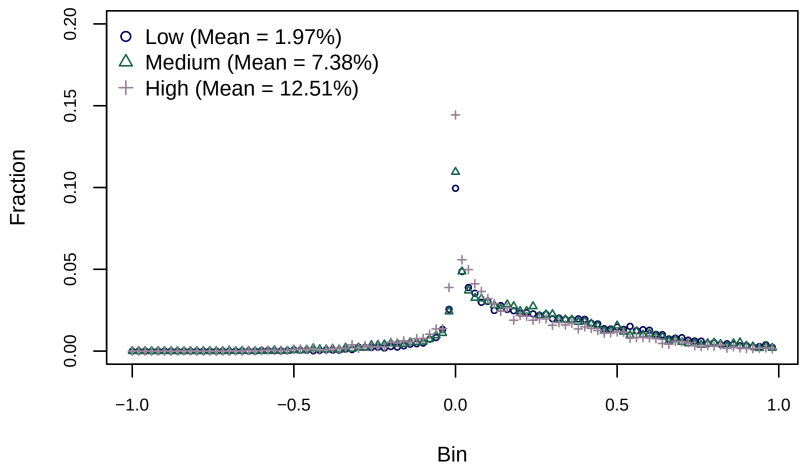

(c) Non-Performing Loan Heterogeneity

Figure 12
Heterogeneity by Mortgage Status

Panel (a) shows the distribution of transactions across 2% reported value bins, where reported value is measured as the percentage deviation from the guidance assessed value. The sample in (a) includes 187,999 transactions in Mumbai and Mumbai suburban districts. Panel (b) shows the distribution of transactions by the ownership structure of a bank making an associated mortgage - the sample here is 32,166 transactions where we were able to successfully match a mortgage. Panel (c) uses the same sample as panel (b) but presents distributions based on the terciles of the average non-performing loan rate of the associated bank.

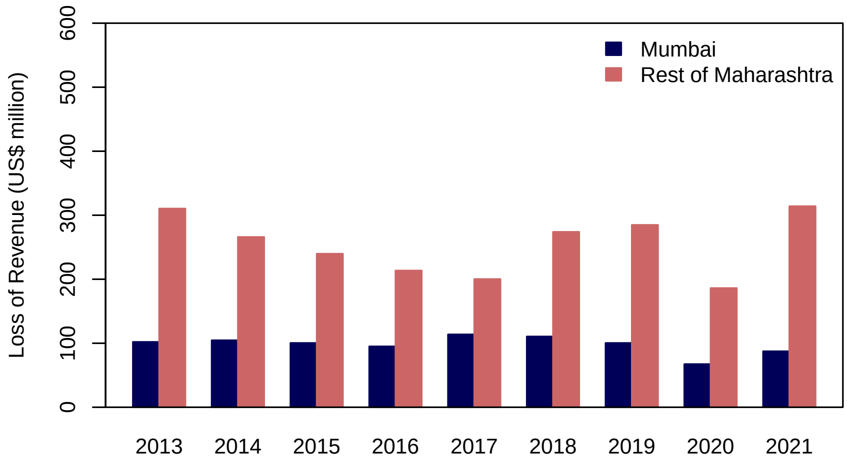

Figure 13
Estimating Total Revenue Loss for the Government

This figure presents the annual total revenue loss for the government as a result of underreporting, from the Mumbai Metropolitan Region and the rest of Maharashtra.

Table 1
Summary Statistics on Transactions

	Reported Value '000s USD		Guidance Value '000s USD		Propequity Value '000s USD		Primary Transaction = 1		Area (sq M)		No. Obs.
	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median	
2013	275.66	173.01	208.71	146.04	326.32	205.23	0.66	1	85.55	76.58	13,648
2014	313.27	195.57	234.15	158.61	362.44	225.50	0.69	1	88.75	75.14	17,213
2015	319.13	203.41	251.66	175.52	362.99	243.08	0.70	1	83.16	72.46	20,615
2016	315.70	205.50	258.22	175.88	363.41	241.19	0.70	1	79.13	69.90	23,803
2017	337.22	220.00	289.13	194.27	386.11	258.64	0.74	1	78.00	68.82	31,104
2018	324.03	212.86	264.82	175.48	367.51	243.07	0.76	1	73.53	65.04	38,228
2019	315.59	215.04	254.19	171.50	360.73	245.33	0.73	1	71.02	62.97	30,602
2020	320.36	210.00	266.78	172.50	371.71	241.19	0.72	1	71.35	62.34	30,289
2021	334.36	212.52	274.85	177.39	374.59	246.06	0.68	1	72.21	62.15	49,663
2022	324.45	220.00	257.77	176.44	355.16	240.91	0.63	1	68.14	60.57	5,449
Total	321.77	210.00	261.86	174.42	367.29	242.02	0.71	1	76.06	66.28	260,614

The table reports summary statistics for the set of transactions that is either matched to the same project from Propequity or to the nearest Propequity project.

Percentile of Guidance Value Distribution	Lower End of Guidance Region ('000s US$)	Kink ('000s US$)	Upper End of Guidance Region	Conventional Elasticity	Standard Error of Conventional Elasticity
5–15	16	21	26	8.07	0.328
15–25	26	30	35	2.96	0.078
25–35	35	40	46	3.00	0.070
35–45	46	53	61	1.87	0.035
45–55	61	69	78	1.37	0.038
55–65	78	89	101	1.28	0.016
65–75	101	115	134	1.61	0.032
75–85	134	157	190	1.57	0.020
85–95	190	242	365	1.24	0.024

Table 2
Elasticity of Reported Value to Transaction Tax Rate

The table reports the formal estimates of the elasticity of reported value to the transaction tax rate.

Table 3
Under-reporting Before Circle Rate Changes

Dep Var: Under-reporting Rate	(1)	(2)	(3)	(4)
Month Before Policy Change	0.061***	0.054***	0.064***	0.068***
	(0.000)	(0.000)	(0.000)	(0.000)
Mean Dep Var.	0.06	0.06	0.06	0.06
Intercept	Yes	Yes	Yes	Yes
Time-trend	No	Yes	Yes	Yes
Month of year FE	No	No	Yes	Yes
Year FE	No	No	No	Yes
No. Obs.	260,614	260,614	260,614	260,614

The table reports the regression results estimating the average under-reporting rate the month before circle rate changes.

A Bad Bunch: Asset Value Under-Reporting in the Mumbai Real Estate Market

Internet Appendix

Santosh Anagol Vimal Balasubramaniam

Tarun Ramadorai Antoine Uetwiller

A Data

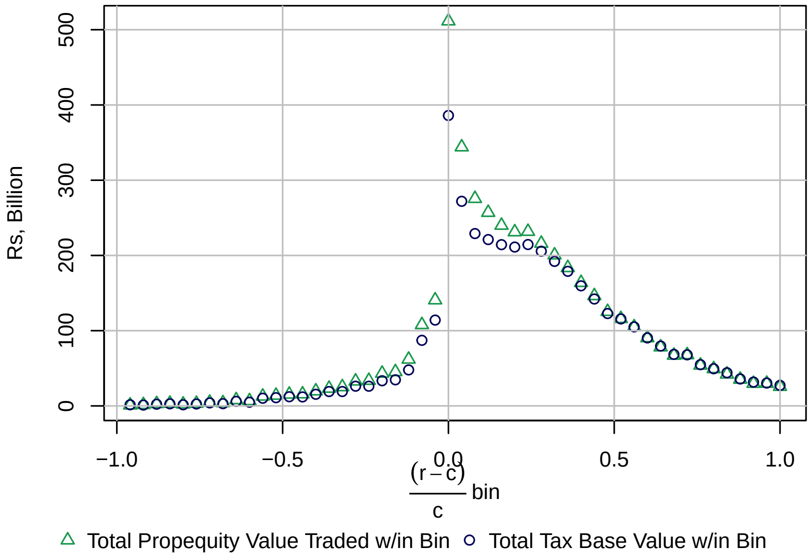

(a) Bin width of 4%

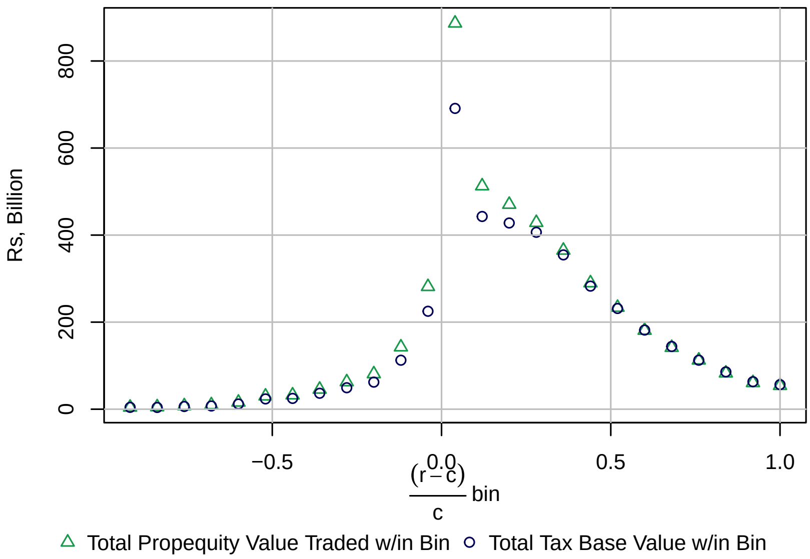

(b) Bin width of 8%

Figure A.1
Robustness: Aggregate Reported and Propequity Values by Reporting Behavior Bins

Panel A reports aggregated values with a bin width of 4% and Panel B with a bin width of 8%. See Figure 3 for details.

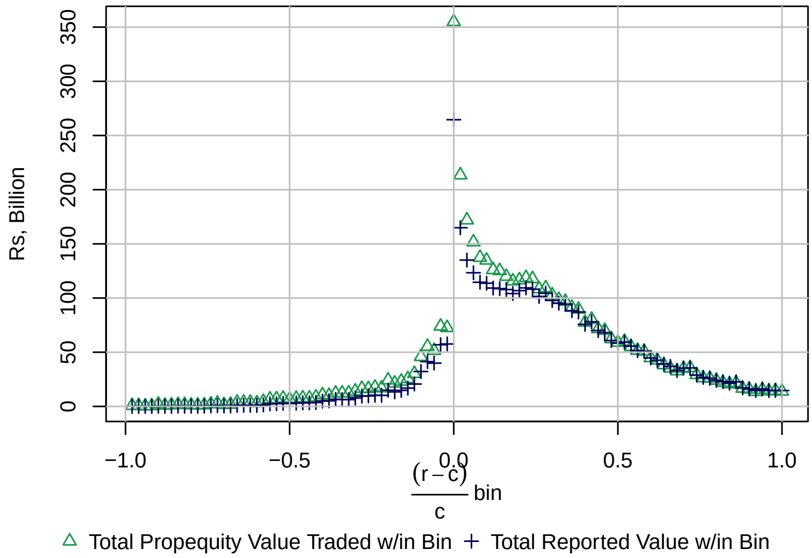

Figure A.2
Aggregate Reported and Propequity Values by Reporting Behavior Bins

This figure reports the aggregated reported and propequity values by reporting behavior bins. See Figure 4 for details.

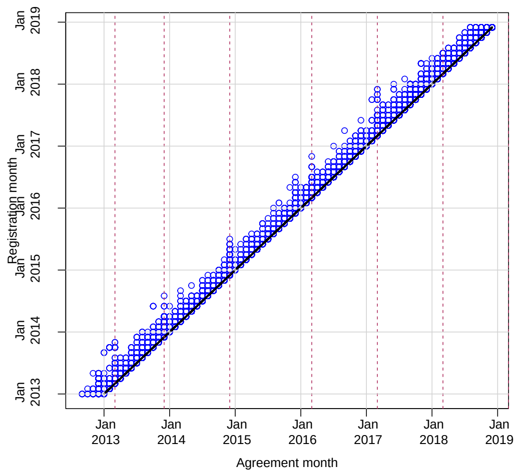

Figure A.3
Agreement Date Backdating

This figure plots the agreement month on the x-axis and the registration month on the y-axis.

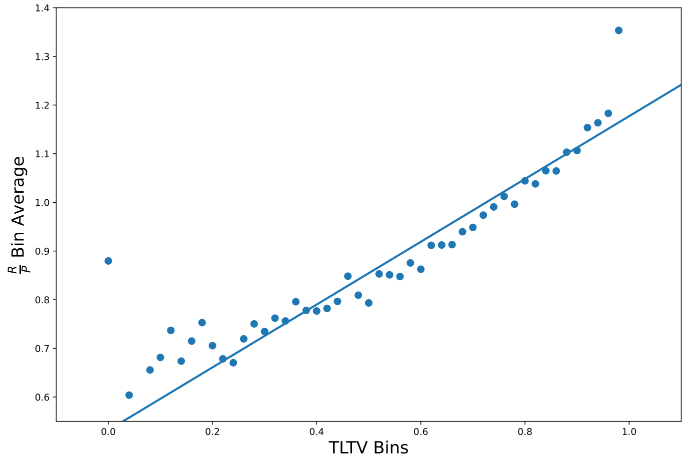

Figure A.4
Equilbrium Relationship between $\frac{R}{P}$ and $TLTV$

This figure plots the average fraction of reported value to the propequity value $\frac{R}{P}$ (y-axis) for bins of the true LTV ($TLTV$) as implied by the propequity value p.

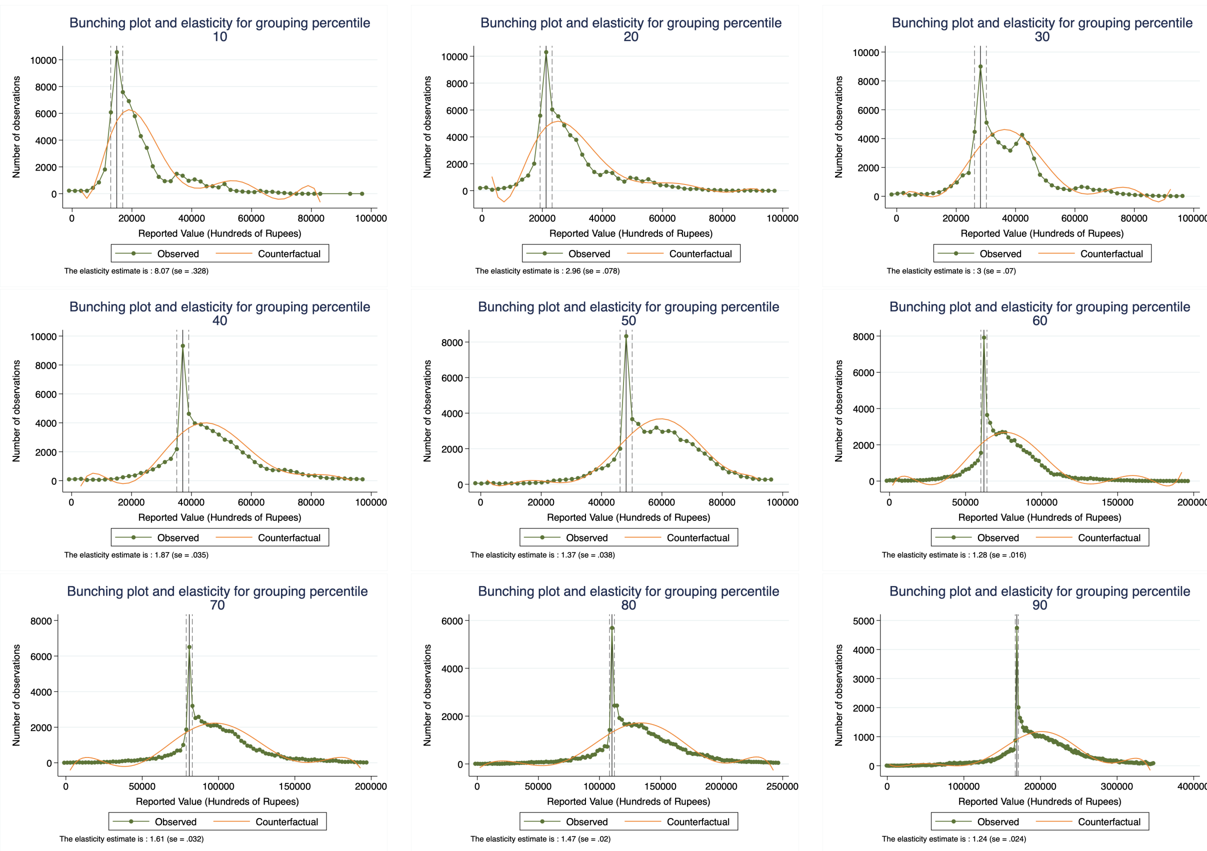

Figure A.5
Bunching Estimates for Guidance Value Bins

This figure presents the bunching estimates for various Guidance Value bins.

Table A.1
Guidance Value Systems in Cities around the World

City	Country	Transaction Tax Rate	Additional tax	Assessment Method	Registration Fee
Delhi	India	6% men, 4% women, 5% joint (stamp)		Centralized	1%
Sao Paulo	Brazil	3%		Centralized	0.75%
Mumbai	India	5% men, 4% women (stamp)		Centralized	1% or 30,000 rupees, whatever is lower.
Buenos Aires	Argentina	3.6%; if the property is for residential use, valued under ARS 975,000, and the client's first purchase, stamp duty is waived.	10.5% VAT for residential buildings, 21% for other buildings	Centralized	0.20%
Kolkata	India	4%-5% stamp		Centralized	1%
Lagos	Nigeria	2% stamp	8% consent fee; 5% VAT	Decentralized	3%
Rio de Janeiro	Brazil	3%		Centralized	0.75%
Moscow	Russia	0.3% land tax	20% VAT	Centralized	0.1%-1%
Paris	France	For properties more than 5 years old, stamp duty is 5.8%, or 5.09%. For properties less than 5 years old, stamp duty is 0.7%.	For properties more than 5 years old, additional 20% VAT	Decentralized	5.10%

Continued on next page

Table A.1 – continued from previous page

City	Country	Transaction Tax Rate	Additional tax	Assessment Method	Registration Fee
Bogota	Colombia	1%	Transaction tax is called a registration tax (impuesto a registrar)	Centralized	0.50%
Jakarta	Indonesia	5%	1% deed tax	Centralized	0.20%
Chennai	India	7% stamp		Centralized	4%
Lima	Peru	3%		Centralized	0.81%
Hyderabad	India	4% stamp		Centralized	0.50%
London	United Kingdom	Progresssive stamp duty. 3% higher in each bracket if buyer owns another residence.		Centralized	Progressive fixed fee.
Tehran	Iran	10% for new buildings. 3%-5% otherwise	0.5% stamp	Decentralized	0.10%
Chicago	United States	3.75 dollars per 500 dollars		Decentralized	First registration 250-500 dollars for vacant buildings. Semiannual fee of similar ammount afterward
Ho Chi Minh City	Vietnam	No transfer tax	5% VAT	Decentralized	0.5% and VND20,000

Continued on next page

Table A.1 – continued from previous page

City	Country	Transaction Tax Rate	Additional tax	Assessment Method	Registration Fee
Luanda	Angola	2% IPT. Does not apply if transacted property is first property owned by buyer and used for personal and long-term residential purposes	0.3% stamp (on true value)	Decentralized	AOA 105,600 (for certain commercial real estate; unclear for other types of properties)
Ahmedabad	India	4.9% stamp		Centralized	1% for properties exceeding Rs. 30 Lakh. Otherwise, women pay 0%
Kuala Lumpur	Malaysia	1-4% progressive stamp		Decentralized	MYR100
Hong Kong	China	15% BSD	5-20% SSD; 1.5-8.5% AVD	Decentralized	230-450 dollars
Riyadh	Saudi Arabia	5%		Decentralized	
Surat	India	4.90%		Centralized	1% (men only)
Madrid	Spain	6.00%	new property: 1.5% stamp duty + 10% VAT if the seller is a company)	Centralized	0.02%-0.175%

Continued on next page

Table A.1 – continued from previous page

City	Country	Transaction Tax Rate	Additional tax	Assessment Method	Registration Fee
Pune	India	5-6% men; 4-5% women	1% metro + 1% Local Body Tax	Centralized	For properties below Rs 30 lakh - 1% of the property value. For properties above Rs 30 lakh - Rs. 30,000
Toronto	Canada	0.5-2.5%	Ontario land transfer tax	Decentralized	
Belo Horizonte	Brazil	3%		Centralized	
Singapore	Singapore	1-3%	5-15% additional buyer stamp duty dependent on citizenship	Decentralized	SGD70 (US$52)
Philadelphia	United States	3.47%	additional 1% tax for commonwealth	Decentralized	$25
Atlanta	United States	$1 per $1,000		Centralized	$100
Barcelona	Spain	7%-11%	10% VAT (new house)	Centralized	400 to 750 EUR per deed; stamp duty (registration) 1.5%
Saint Petersburg	Russia	0.3% land tax	20% VAT	Centralized	0.1%-1%

Continued on next page

Table A.1 – continued from previous page

City	Country	Transaction Tax Rate	Additional tax	Assessment Method	Registration Fee
Washington Met.Area	United States	1.1 % of consideration or fair market value for residential property transfers less than $400,000 and 1.45% of consideration or fair market value on the entire amount, if transfer is greater than $400,000.		Centralized	

A.1 Transactions Data

The main data source for this project is the publicly available individual property transaction reports released by the Office of the Inspector General of Registration and Controller of Stamps (IGR), Department of Revenue, Government of Mahrashtra, India. This state apparatus plays an important role in collecting state government revenues from across the state using various fiscal instruments available in the state government's toolkit. The state is split into 8 regional divisions and we obtain data for the Mumbai regional division which is comprised of Mumbai City and Mumbai Suburban districts. Our study area currently covers 437 square kilometers out of the 6,640 square kilometers Mumbai Metropolitan Region. We currently focus on this region because we can reliably obtain transaction data that can be mapped to geo-spatial information relevant for our study.

The *eSearch* facility set up by the IGR enables access to transaction-level data for all properties transacted in Greater Mumbai. Every transaction report is in Marathi, the most commonly spoken language in Maharashtra. Figure A.6 presents an example of the original document downloaded from the IGR *eSearch* facility. Figure A.7 presents the transaction report translated into English using Google's translation services. The details available in each transaction report provides a consistent information set for all real-estate transactions for Greater Mumbai. This information set also serves as the basis for the government to make policy decisions on real-estate transaction taxes.

Each transaction report obtained from the *eSearch* facility begins with document number, and the name of the registrar office (the local IGR office for a region). The more substantive information is in the form of a table starting with the name of the local village where the property is located[28]. The first row of the table in Figure A.7 lists the type of transaction. All real-estate transactions in Maharashtra are classified as "Agreement", "Agreement to Sale", "Sale deed" and "Transfer Deed" types. We filter all downloaded transaction reports to these deed types to form our core data set.

The second row lists the reported price at which the transaction took place. In this case, the reported transaction price is ₹7,500,000. The third row lists the price as per the government issued guidance value, known as the *policy circle rate* that is determined annually by a legally predetermined process. The policy circle rate determines the floor price at which the government will deem this property to be sold for taxation purposes. The value of this property according to government determined circle rates is ₹4,434,062. Row 12 provides the computed stamp tax paid on this transaction of ₹375,000, determined as the prevailing stamp tax rate, in this case 5%, on the reported transaction value. The circle rate plays an important role in that it sets the lower bound for the stamp tax revenue generated for a property of this type. In the event this property's reported transaction price is below ₹4,434,062, the stamp duty payable will be 5% of this guidance value, after which the law facilitates a process by which the related parties can file for revision. The fourth row of this table provides the property address, and other measurement details in terms of the area of the house, the land registry survey number, and other information relevant for determining the circle rate. The fifth row of this table provides us with the property area, and the next few rows provide details of the two parties to the transaction. Row 9 reports the transaction date for the document, and Row 10 the actual date of formal registration for the sale. These two dates can be different as the law allows for a grace period of 3 months from the actual transaction date during which time they are legally bound to register the sale with the registrar. The last row provides data on the registration fee paid which is capped at ₹30,000 or 1% of the reported value, whichever is lower.

We validate the coverage of our transaction reports data from the IGR *eSearch* facility by matching the total real-estate stamp duty revenue generated in each year from

[28] Historically, the Mumbai region was formed of seven islands or fishing villages, which then expanded rapidly over time. The village tag to geographies is more of an artefact of historical documentation than a reference to the economic or social conditions of different regions in Mumbai.

our transactions data to the official aggregate numbers. Figure **??** presents this comparison. The top panel of the figure shows the official aggregate tax revenue collected from stamp duty in each month in orange and the aggregate stamp revenue from our transaction data. The two time-series are very highly correlated, especially in the second half of our sample period. The bottom panel of Figure **??** presents the estimated aggregate stamp revenue as a percent of the official figures. The total revenue figures from both sources include stamp revenues and the registration fees for all transactions in a given month.

Although we capture a majority of the transactions in Greater Mumbai, the differences between our aggregate revenue numbers and the official figures arise primarily due to two reasons. The official figures for Mumbai also includes a suburban area of Navi Mumbai, which we do not include in our sample. Moreover, we count revenues in the month the transaction was registered, and this may not necessarily be the same as the official approach, especially for transactions that may be executed at the end of the month, but fulfilled in the early periods of the following month. We therefore plot the three-month and twelve-month rolling average in orange and green respectively and they suggest that we cover nearly 80% of the aggregate revenue in our transactions data.

A.2 Matching Transaction Documents to Propequity Buildings Via Location

A major part of our methodology is estimating market values for the transactions we observe by matching transactions to new building prices from the Propequity database. Matching transaction documents to the Propequity dataset requires us to obtain geolocation information for our transaction documents. Our main approach is to match transaction documents to the "CTS" level. A CTS is the smallest administrative geospatial unit in Greater Mumbai for the IGR.[29] This CTS number is important for property registration, mortgages, and determination of the stamp duty to be payable when a property is being bought or sold. We primarily use the CTS of a sales transaction to geo-locate the project. We obtain GIS information for CTS in the Mumbai division from the Urban Development Research Institute (UDRI). We extract the shapefiles for each CTS using ArcGIS, and use these polygons to identify the CTS location for each

[29] CTS stands for the Chain and Triangulation Survey Number in the Mumbai suburban district, and the Cadastral Survey Number in Mumbai division. A set of CTS numbers form a sub-zone, and then aggregate upwards to the Mumbai division of the IGR.

of our IGR transactions.

Panel A of Appendix Table A.3 presents the details of sample attrition resulting from the process of geo-tagging each transaction report. We begin with the set of transaction reports that contain non-missing information for reported value, guidance values, and the property area. We also exclude reported values less than ₹1000) and area of under 10 square meters. After these initial filters, we have 215,121 transactions in our sample period. 6,864 reports have no property description in the transaction reports, making it impossible to identify its location. Of the remaining 208,257 transactions, we identify the location for these properties using three different approaches. First, if the property description in the transaction report contains the CTS number, we use it to match to the CTS geo-location using our spatial polygons obtained from UDRI. Second, some properties mention several CTS numbers in their property description. This happens when large apartment blocks straddle multiple CTS, and in these instances, we map the property to the first CTS number in the property description. Lastly, if there are no CTS numbers available, we use the property description to obtain the latitude-longitude information from Google Maps or Bing, and then match it to the geo-spatial data to identify the CTS number for these properties. Using the three approaches to locate the transactions in Mumbai, we successfully match the data for 187,999 transactions, or 87.3% of the full sample of transaction reports obtained from the IGR website.

To complete the match to Propequity buildings, we next match Propequity projects to the CTS level based on the projects latitude and longitude. Finally, we use the average price of Propequity projects in a transaction's CTS to estimate the transactions "market" price.

A.3 Circle Rate Scheduled Changes

Figure 6 shows how reported values and under-reporting behavior evolved over time, particularly in reference to pre-determined dates when circle rates were changed. Circle rates are set at the sub-zone level, a geographic area of approximately .67 square kilometers on average. Table A.2 presents the summary statistics on the circle rate variation in Greater Mumbai for our sample period. At the start of our sample period we have 727 sub-zones, which increase to 747 sub-zones in 2015, and then stabilize at 734 for the remainder of our sample period.[30] In the early years of our sam-

[30] New sub-zones are formed by either dividing existing sub-zones into multiple new ones, or by fusing different parts of multiple sub-zones to form new ones. We keep track of all of the changes in the

ple, nearly all sub-zones underwent changes in circle rates. The average change in each year vary from 0% in 2018 to 14.4% from the previous year in 2015 (Column 3). The cross-sectional distribution is also large. At the lowest end of the distribution are sub-zones with ₹7330 as the circle rate per square meter of property area (in 2014), to ₹653,240 per square meter of property area in 2018. Figure A.8 presents the geo-spatial variation at the sub-zone level in circle rates at the beginning of our sample (Panel A) and at the end of the IGR sample in 2018 (Panel B). The circle rates have been re-scaled to the mean sub-zone, with the darker red indicating sub-zones with high circle rates and sub-zones in lighter shades of yellow indicating those with the lowest circle rates in Greater Mumbai.

B Simulation to Illustrate Effects of Measurement Error in p

Figure 1e shows the distribution of $\frac{r-c}{c}$ and $\frac{m-c}{c}$ for the no under-reporting with measurement error case. The $\frac{r-c}{c}$ blue distribution is exactly the same, by construction, as that in the high under-reporting with measurement error case. However, the green $\frac{m-c}{c}$ is constructed to reflect bunching in the market value distribution around circle assessed values as opposed to under-reporting behavior. Such bunching could result if the tax authority sets circle assessed values to correctly match market values, or indeed, if buyers and sellers anchor on circle prices when negotiating sale prices. The red curve plots p, which is assumed to be a noisy measure of m.

In the simulation shown in Figure 1e we set $p = m + \epsilon$, where ϵ is distributed normally with mean 0 and standard deviation 5. The key insight is that if we do not observe the true distribution of market values $\frac{m-c}{c}$, and cannot correct for measurement error, then it is difficult to distinguish the high under-reporting, no measurement error case from the zero under-reporting, high measurement error case by inspecting bunching behavior alone. The way we have set up the simulation, bunching in reported values is the same magnitude in both cases, and the distribution of our measured market prices p is simply a smoothed version of the bunching in both m and r around c.

geo-spatial files, thus identify which regions form to create the new sub-zones, and the old sub-zones they belonged to.

C Matching Transactions to Mortgages

We start with 215,121 transactions and 125,195 mortgage transaction documents in the Mumbai central and Mumbai suburban districts.[31] Table A.4 describes our matching procedure so far. We currently match mortgages to transactions using the PAN (tax identification number) of the property buyer (from the transaction document) and the borrower (form the mortgage document), as well as the area (in square meters) of the property. 121,410 of the mortgage documents have a usable PAN number and area information.

Matching directly based on PAN and area we are able to match 42,996 mortgage transactions to a transaction document. Of these 42,996 matching mortgages, 10,385 are duplicate or extraneous documents for the same transaction. For this set of mortgage to transaction matches, we take the chronologically first matching mortgage and drop the others. This leaves us with 35,706 transactions associated with a mortgage (where we only keep the first mortgage). Of these 35,706 transactions associated with a mortgage, 32,611 are transactions that had a property description and a CTS number (i.e. these are transactions that are in our transaction analysis sample of 187,999 transactions). Finally, amongst those 32,611 transactions with mortgages, 31,119 have usable loan information and therefore can be used to calculate a loan to value ratio. To summarize, we currently have 187,999 sales transactions. Of these, 31,119 are matched to a mortgage and are assigned a loan to value ratio based on that mortgage. 1,419 are matched to a mortgage but we do not observe the loan to value ratio. This leaves 187,999 - (31,119 + 1,419) = 155,461 transactions that are currently not matched to a mortgage and therefore assigned a loan to value ratio of zero.

We note that this matching procedure leaves 78,414 mortgages completely unmatched to transactions. This is a very large number and suggests that many of the transactions we are currently assigning a loan to value ratio of zero actually do have an (unmatched) mortgage and therefore a higher true loan to value ratio.

[31] Note for this matching exercise we start with the full number of sales transactions we downloaded for this region before removing observations without a property description and without a CTS.

Table A.2
Circle Rate - Summary Statistics

This table reports the summary statistics on variation in circle rates across sub-zones in Greater Mumbai. Column 1 reports the total number of sub-zones in each year of our sample. Column 2 reports the percent of sub-zones that witnessed a change in circle rates compared to the previous year. Column 3 presents the average percentage change in circle rates relative to the previous year. Columns (4–9) present the cross-sectional distribution of circle rates in 1000s of rupees per square meter of property area.

	Sub-zones			Cross-sectional Distribution (×1000₹)					
Year	#	% with Change	% Change	Mean	1%	25%	50%	75%	99%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
2014	727	-	-	139.59	7.33	81.35	109.00	161.85	580.89
2015	747	100.00	14.44	160.21	14.72	94.60	126.20	189.55	619.25
2016	734	97.49	10.54	172.75	30.04	103.92	134.45	201.05	652.03
2017	734	68.46	6.98	178.40	11.62	109.78	145.15	209.88	653.24
2018	734	0.00	0.00	178.40	11.62	109.78	145.15	209.88	653.24

Table A.3
Data Validation: GIS Tagging

This table reports the sample attrition due to GIS tagging of all transaction data from the IGR (Panel A), and validating the GIS tag by using circle rate information in the IGR transaction reports (Panel B).

Panel A: Sample Attrition

	Number of Observations
Registrar	215,121
Without Property Description	-6,864
With Property Description	208,257
Does Not Match to a CTS	-20,258
Final Sample	**187,999**
- Perfect CTS Match	48,351
- Match on the First Number of the CTS	86,139
- Google/Bing based Match	53,509

Figure A.6

Original Transaction Document: An Example

This figure presents the original transaction document obtained from the IGR *eSearch* faciilty `https://freesearchigrservice.maharashtra.gov.in`.

249451
18-04-2019
Note:-Generated Through eSearch Module,For original report please contact concern SRO office.

सूची क्र.2

दुय्यम निबंधक : सह दु.नि. बोरीवली 7
दस्त क्रमांक : 249/2017
नोदंणी :
Regn:63m

गावाचे नाव : 1) कांदिवली

(1)विलेखाचा प्रकार	करारनामा
(2)मोबदला	6499800
(3) बाजारभाव(भाडेपटटयाच्या बाबतितपटटाकार आकारणी देतो की पटटेदार ते नमुद करावे)	6259071
(4) भू-मापन,पोटहिस्सा व घरक्रमांक(असल्यास)	1) पालिकेचे नाव:मुंबई मनपाइतर वर्णन :सदनिका नं: 602, माळा नं: 6 वा मजला, इमारतीचे नाव: कांदिवली किशकांत को ऑप हौ सो लि, ब्लॉक नं: कांदिवली पश्चिम मुंबई 400067, रोड : दत्त मंदिर रोड,डहाणूकर वाडी, इतर माहिती: एकुण क्षेत्रफळ 471 चौ. फुट((C.T.S. Number : 933 ;))
(5) क्षेत्रफळ	1) 52.52 चौ.मीटर
(6)आकारणी किंवा जुडी देण्यात असेल तेव्हा.	
(7) दस्तऐवज करुन देणा-या/लिहून ठेवणा-या पक्षकाराचे नाव किंवा दिवाणी न्यायालयाचा हुकुमनामा किंवा आदेश असल्यास,प्रतिवादिचे नाव व पत्ता.	1): नाव[redacted]य[redacted] पत्ता[redacted]
(8)दस्तऐवज करुन घेणा-या पक्षकाराचे व किंवा दिवाणी न्यायालयाचा हुकुमनामा किंवा आदेश असल्यास,प्रतिवादिचे नाव व पत्ता	1): नाव[redacted]वय[redacted] पत्ता[redacted] 2): नाव[redacted]य[redacted] पत्ता[redacted]
(9) दस्तऐवज करुन दिल्याचा दिनांक	16/01/2017
(10)दस्त नोंदणी केल्याचा दिनांक	16/01/2017
(11)अनुक्रमांक,खंड व पृष्ठ	249/2017
(12)बाजारभावाप्रमाणे मुद्रांक शुल्क	325000
(13)बाजारभावाप्रमाणे नोंदणी शुल्क	30000
(14)शेरा	

मुल्यांकनासाठी विचारात घेतलेला तपशील:-:
मुद्रांक शुल्क आकारताना निवडलेला अनुच्छेद :- : (i) within the limits of any Municipal Corporation or any Cantonment area annexed to it.

Figure A.7

Translated Transaction Document: An Example

This figure presents the translated version of the original document in Figure A.6 using Google Translation services.

249451 18-04-2019 Note: -Generated Through eSearch Module	**List no. 2**	Sub Registrar with duni Borivali 7 Document number: 249/2017 Note: Regn: 63m

Name of the village: 1) Kandivali

(1) Type of document	Agreement
(2) Reward	6499800
(3) Quotes (The leaseholder leases the details of the rent that the sergeant should specify)	6259071
(4) Land measuring, portals and home number (if any)	1) Name of the corporation: Mumbai Manipatayar Description: House No. 602, Malala No. 6th Floor, Name of the building: Kandivali Kishakant to Op Hau Soil, Block No. Kandivli West Mumbai 400067, Road: Datta Temple Road, Dahanukar Wadi, Others Information: total area 471 square Foot (CTS Number: 933;))
(5) area	1) 52.52 sq.m
(6) When the levy or connection is given.	
(7) If the name of the party giving the name / address of the document or the order or order of the Civil Court, the name and address of the reply.	1): Name:[redacted] Age:[redacted] Address:[redacted]
(8) Name and address of the respondent, if there is a decree or order of the parties,	1): Name:[redacted] Age:[redacted] Address:[redacted] 2): Name:[redacted] Age:[redacted] Address:[redacted]
(9) Date of the date of the document	16/01/2017
(10) Date of registration of the document	16/01/2017
(11) Serial numbers, Volumes and Pages	249/2017
(12) Stamp duty as per market price	325000
(13) Registration Fee as per marketable	30000
(14) Remarks	

Details taken for the assessment: -	
Selected article on stamp duty: -:	(i) In the limits of any Municipal Corporation or any Cantonment area to annexed it.

1. Number of Mortgages in Mumbai Post 2012	127,195
1.a - With a PAN	124,773
1.b - With a PAN and SQM	121,410
2.a Number of Matching Mortgages	42,996
2.b Number of non Matching Mortgages	78,414
3. Number of Earliest Matching Mortgages	35,706
3.a - With Transaction in our Sample	32,611
3.b - With Transaction in our Sample and LTV ≤ 1	31,119

Table A.4
Mortgage Matching Summary Statistics

Figure A.8
Sub-Zones

Panel A presents a heatmap of the sub-zones at the start of our data, and Panel B at the end of our sample period. The circle rates are rescaled to the mean sub-zone with darker shades representing the sub-zones with the largest circle rates (in Southern Mumbai) and subzones in white the lowest (northern periphery of Greater Mumbai).

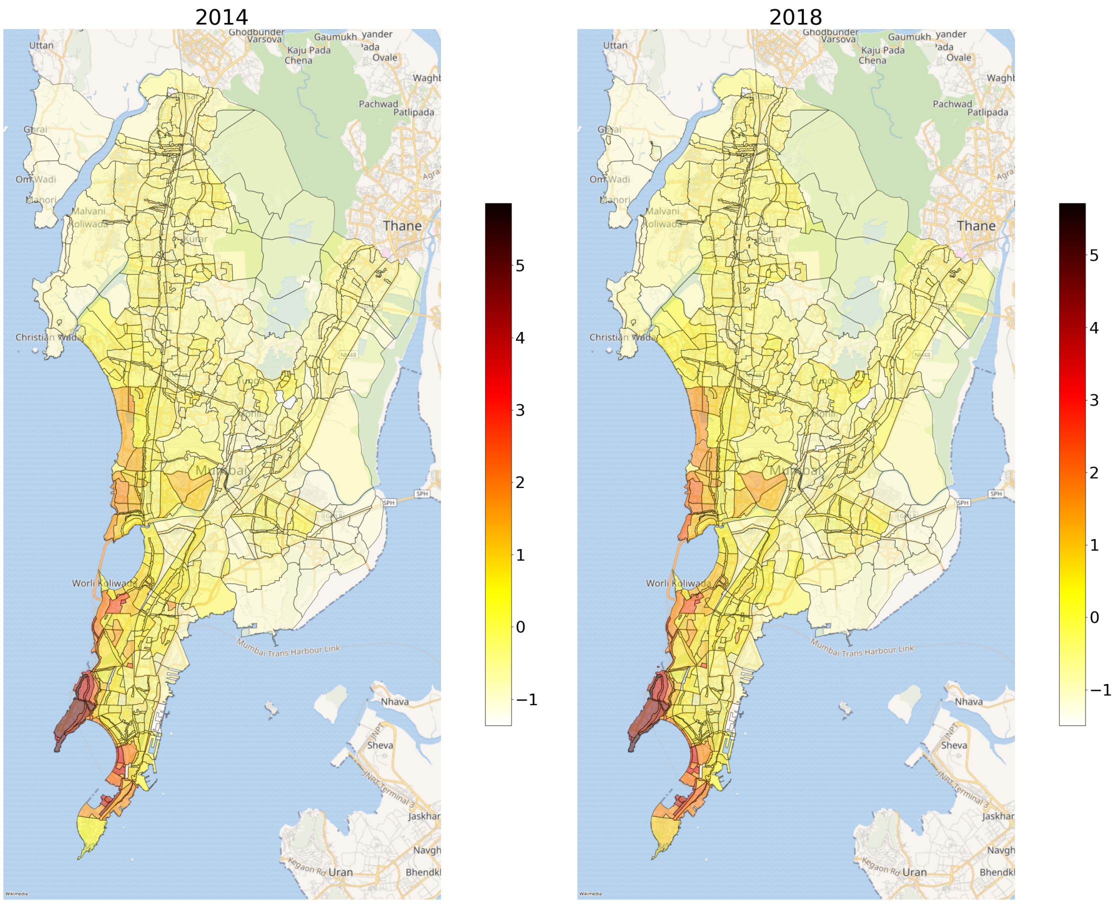

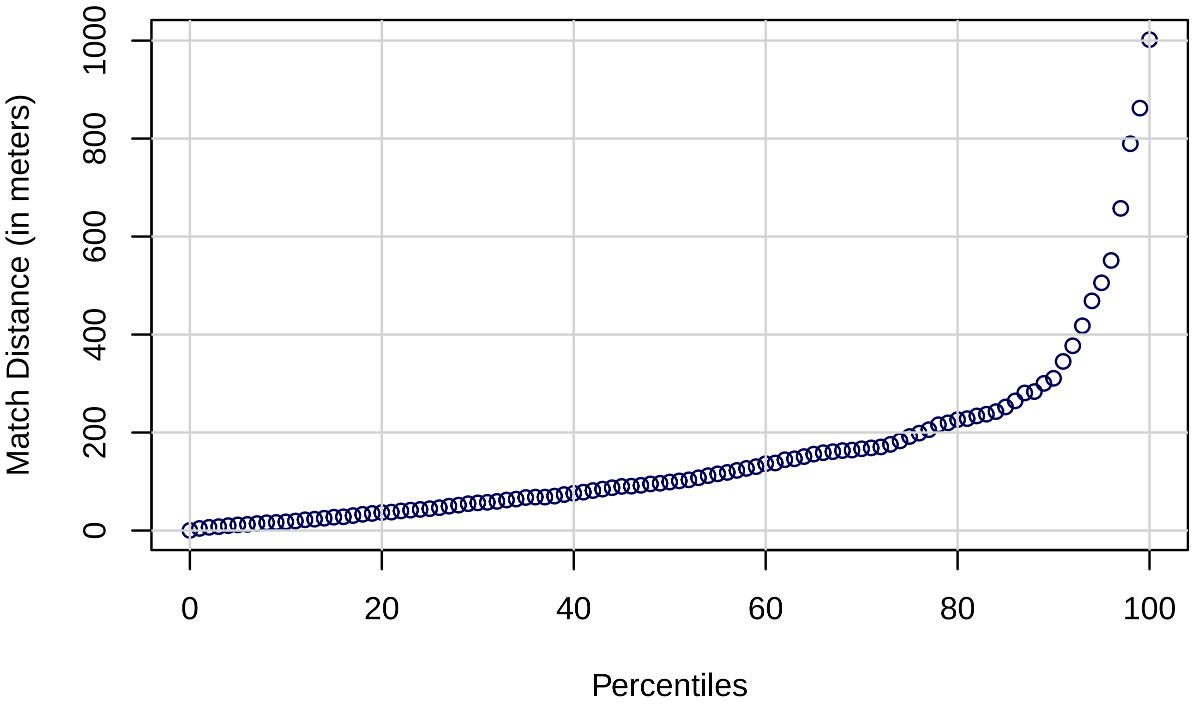

Figure A.9
Propstack-Propequity Match Quality

This figure plots the empirical distribution of the match distance (in meters) between a transaction in propstack data and the match from propequity data. 80% of the transactions matched to propequity information within 200 meters, and 95% of the transactions are matched to propequity transactions within 500 meters.

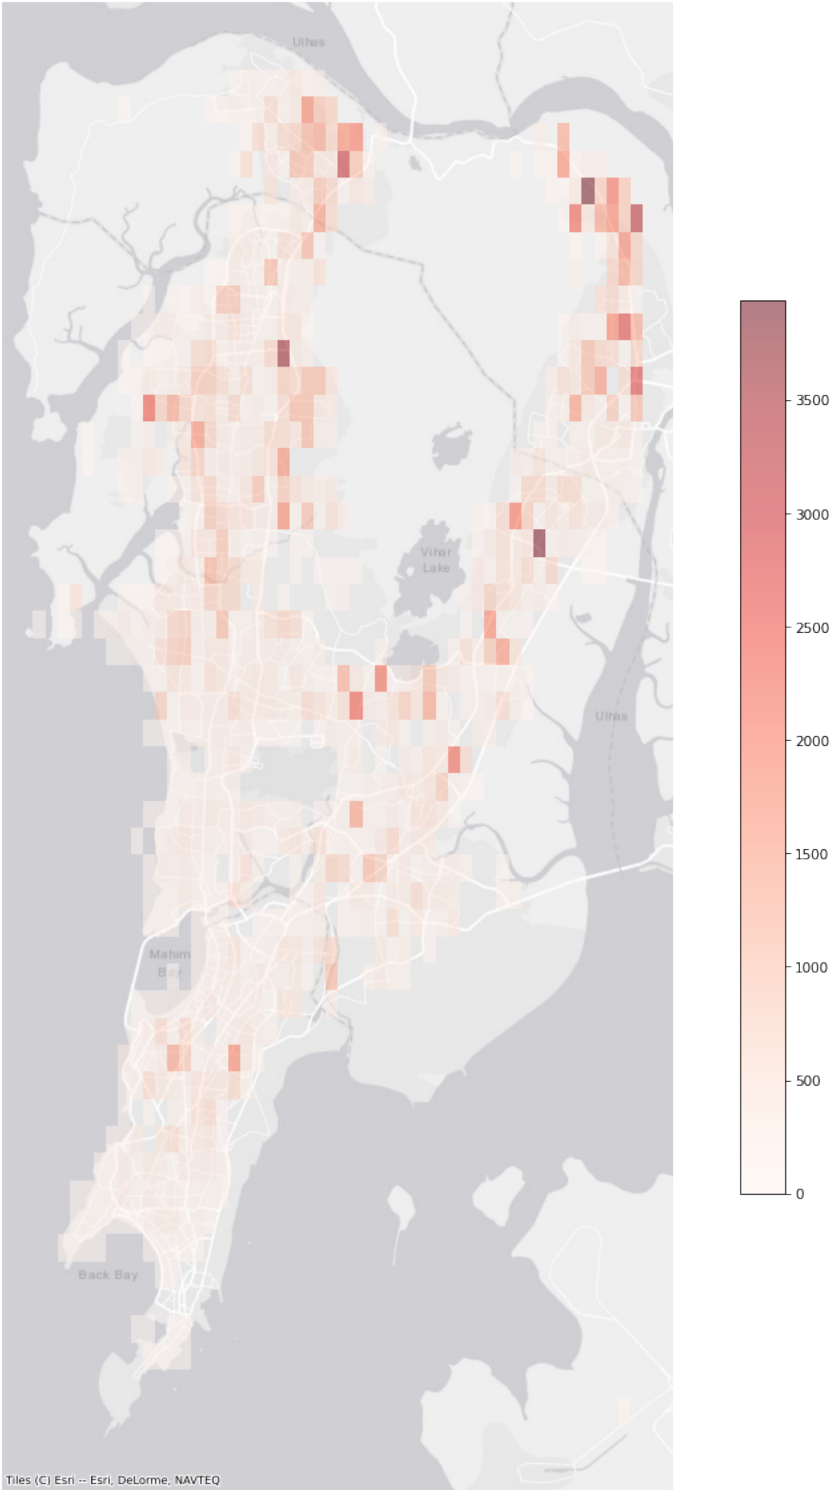

Figure A.10
Heatmap of Transactions in our Final Sample

This heatmap presents the spatial distribution of the final set of transactions in our sample in Mumbai and Mumbai Suburban regions between 2013–2022.

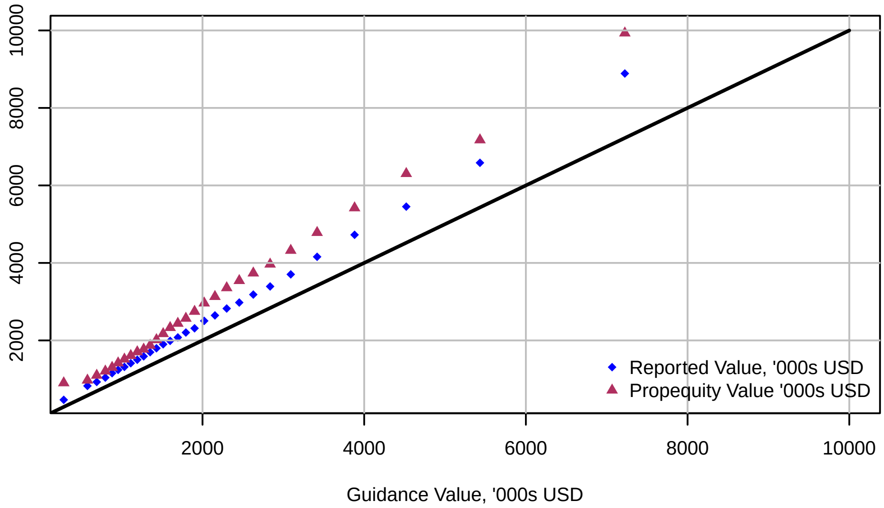

Figure A.11

Correlation between Reported Value, Propequity Value and Guidance Values

This figure presents a binned scatterplot of the average reported values (blue diamonds) and propequity value (maroon triangles) within guidance value bins.

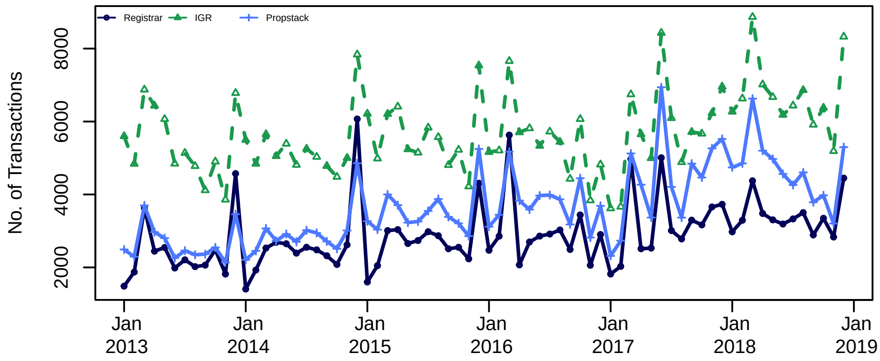

(a) Transactions

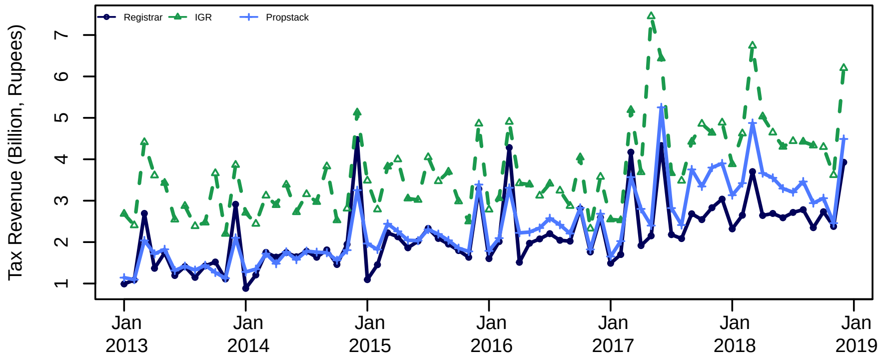

(b) Tax Revenue

Figure A.12
Sample Comparison to Aggregate Tax Revenues

This figure plots the monthly time series of the total number of transactions in panel (a) and the total tax revenue from these transactions in panel (b). The blue line with circles plot the numbers obtained from aggregating the extracted Registrar data, the green triangle is the sum reported by the Inspector General of Registrations for a region that Mumbai and Mumbai suburban areas belong to, that is larger than our sample, and the light blue line with "+" plots the aggregated information from Propstack analytics. The overlapping data sample period ends in January 2019, although our full sample is between 2013–2022.

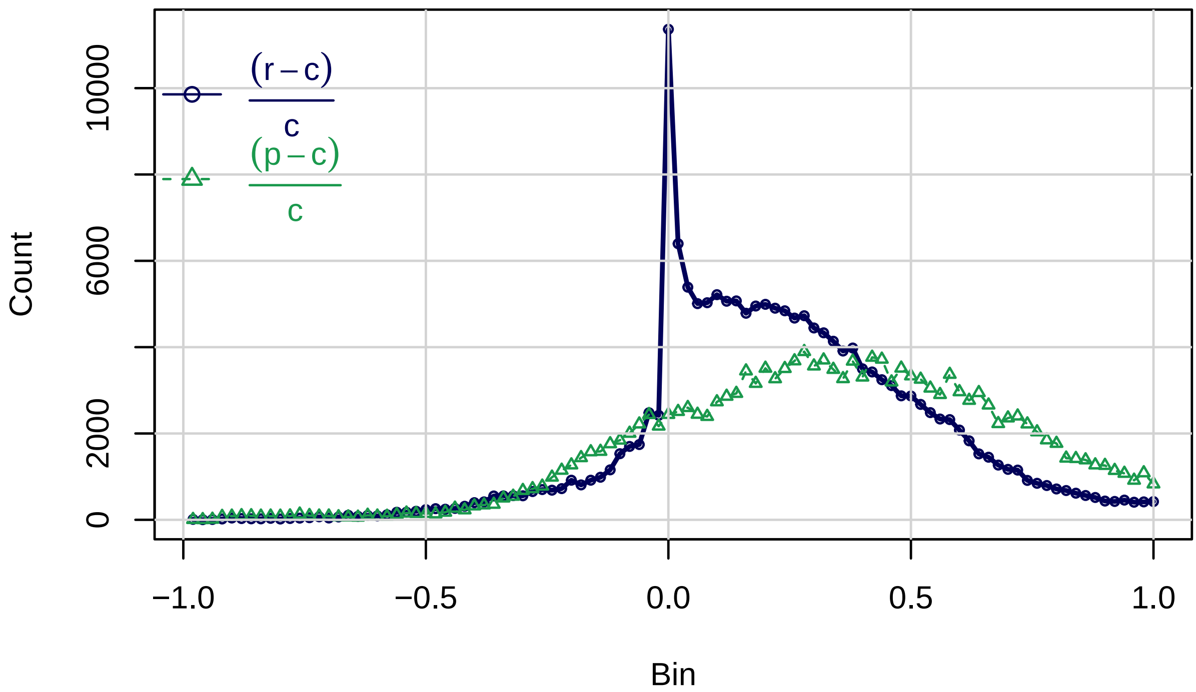

Figure A.13
Bunching of Reported and Propequity Values Around Circle Values

The blue line shows the distribution of reported values across 2% reported value bins, where a reported value bin is measured as a deviation from the government assessed value. The green line shows the distribution our noisily measured estimated of the market price (the Propequity values) for the same underlying set of transactions reported in the blue line.

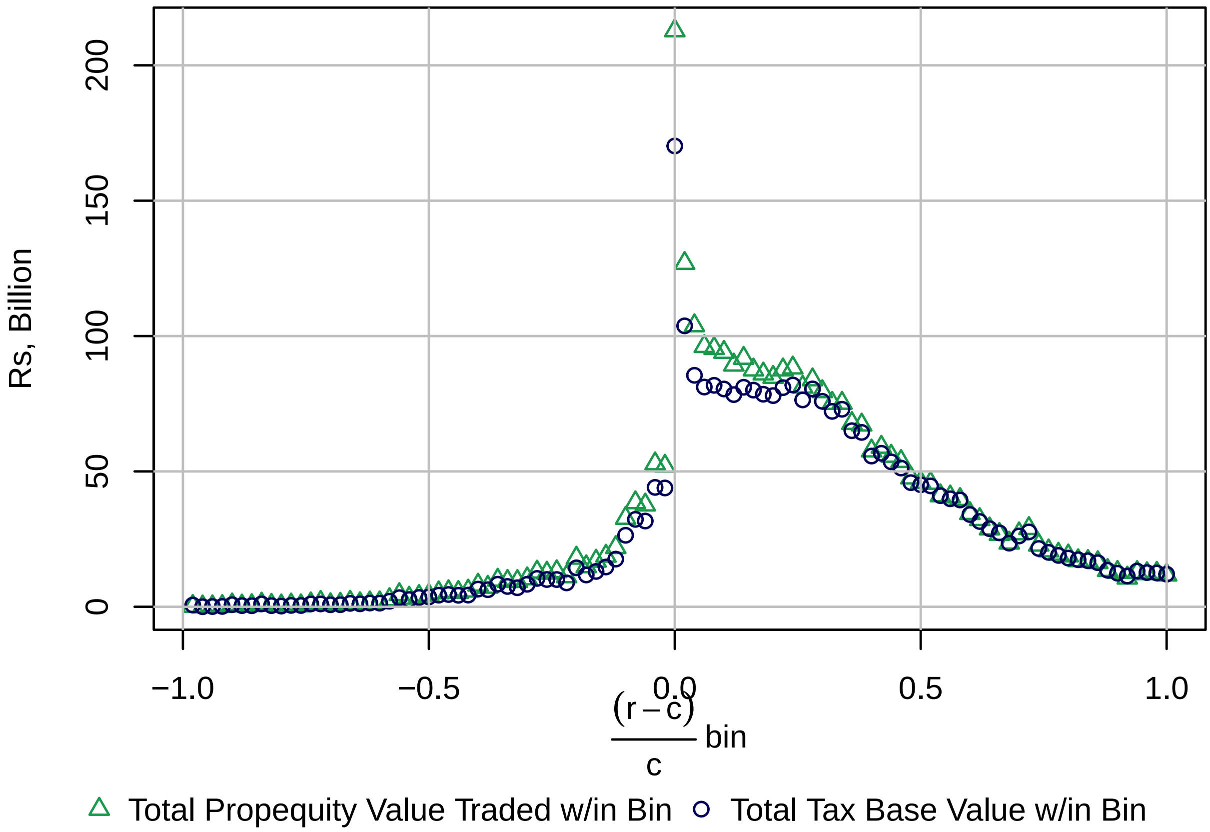

Figure A.14

Aggregate Reported and Propequity Values by Reporting Behavior Bins

The green triangles show the aggregated reported value within 2% reported value bins, where a reported value bin is measured as a deviation from the government assessed value. The blue circles show the aggregate noisily measured estimated of the market value (the Propequity values) for the same underlying set of transactions reported in the green triangles.

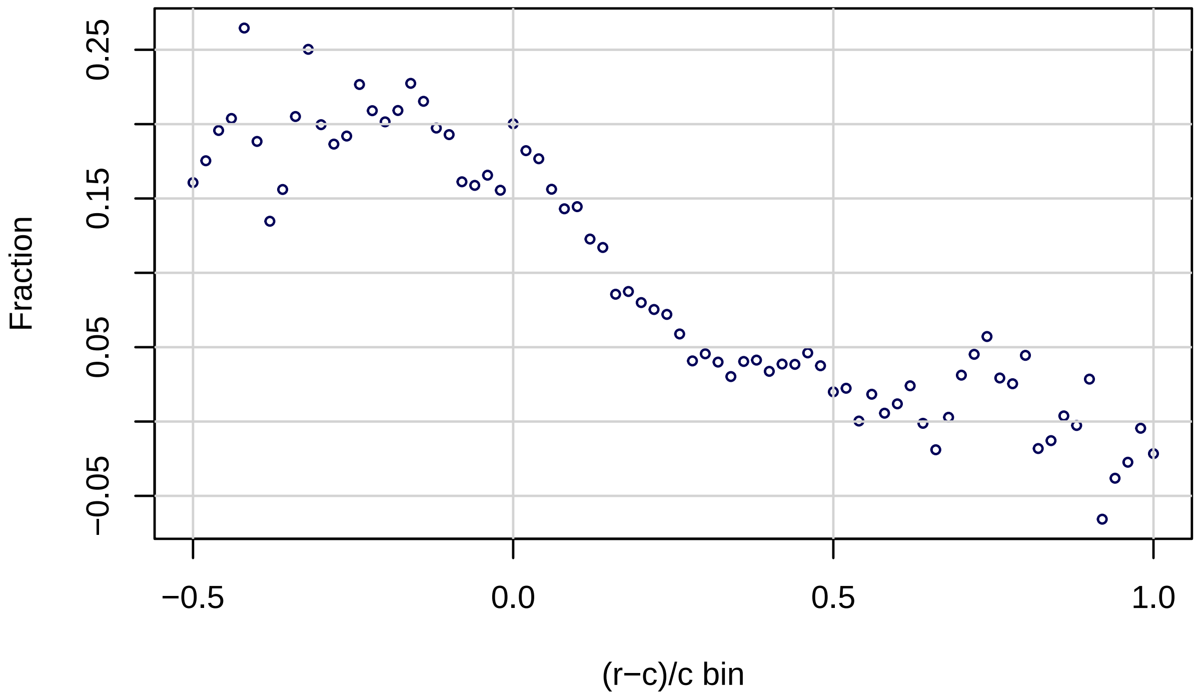

Figure A.15
Under-Reporting Rate by Reporting Behavior Bins

The blue circles show the estimated under-reporting rate within 2% reported value bins, where a reported value bin is measured as a deviation from the government assessed value.

D Bunching Elasticity Estimation

We estimate the elasticity of reporting property values with respect to the transaction tax rate. We employ the conventional method developed by Saez (2010). For our main estimates we assume the tax rate increases from zero below the kink (i.e the guidance value) to five percent above the kink. The formal setup underlying our elasticity estimation is as follows. A household maximizes the following utility function in choose how much to report r for a given house purchase:

$$\max_{r}(R - \tau r) + [r - \frac{r^{(1+\frac{1}{\epsilon})}}{(1+\frac{1}{\epsilon})m^{\frac{1}{\epsilon}}}]$$

where τ is the transaction tax rate (set to .05), m is set to the value the household would report in the absence of any transaction tax, and ϵ is the elasticity of reported value with respect to the transaction tax rate.

The first order condition yields:

$$r = m(1-\tau)^{\epsilon}$$

Substituting $\tau = 0$ we obtain the definition of m as the reported value when there is no transaction tax. This could be truthful reporting, or it could be lower than truthful reporting in the case where there are benefits of under-reporting beyond avoiding the transaction tax. As shown in Saez (2010), differentiating this gives the definition of the ϵ as the percent change in reporting due to a percent change in the tax rate.

Combining this first order condition with equality conditions from the marginal buncher and non-buncher (see Chetty et al. 2011), we have the following relationship between the underlying elasticity of reporting and the bunching mass B.

$$\hat{e} \approx \frac{\hat{B}}{z^* \cdot h_0(z^*) \cdot \log(\frac{1-t_0}{1-t_1})} \tag{6}$$

The estimation procedure involves two steps, first estimating a counterfactual income density based on the income density excluding data points near the kink, and then using the counterfactual density to estimate the excess mass from which the elasticity is recovered.[32] To estimate the counterfactual density, we fit a polynomial of a specified degree to the observed reporting density, excluding the data in a specified

[32] This description closely follows the implementation of the conventional bunching estimator discussed in Anagol et al. (2022).

window around the kink, using the following specification:

$$C_j = \sum_{i=0}^{q} \beta_i^0 \cdot (Z_j)^i + \sum_{i=R_l}^{R_u} \gamma_i^0 \cdot \mathbf{1}[Z_j = i] + \epsilon_j^0. \tag{7}$$

Here, q denotes the order of the polynomial, and R_l and R_u denote the lower and upper bounds of "bunching window" near the kink, which is excluded from the polynomial estimation. The convention in **?** is to set a symmetric bunching window, such that $R_l = -R_u$. Based on visual inspection of the plots we set the bunching window as one bin to the left of the kink, the kink bin, and one bin to the right of the kink. When estimating the polynomial regression, we follow **?** and impose an "integration constraint" such that the total count of observations across the empirical distribution equals the integral of observations under the counterfactual density across the plotted region.[33]

The second step is to compute the excess mass of reported values around the kink relative to this counterfactual density. Using equation (7), we compute the counterfactual mass in each bin within the bunching window, $\hat{C}_j^0$. Subtracting this predicted mass from the observed density yields the estimated excess number of individuals who report incomes near the kink relative to this counterfactual distribution:

$$\hat{B} = \sum_{i=R_l}^{R_u} C_j - \hat{C}_j^0 = \sum_{i=R_l}^{R_u} \hat{\gamma}^0. \tag{8}$$

Standard errors for $\hat{e}$ are estimated using a bootstrap procedure. We resample with replacement from the underlying distribution of transactions 10 times, re-estimating the elasticity each time, and defining the standard error as the standard deviation of the distribution of $\hat{e}$ estimates.

[33] Kleven (2016) notes that imposing an integration constraint may bias the elasticity estimate: "This approach may introduce bias, especially in relatively flat distributions in which interior responses do not affect bin counts (except at the very top of the distribution away from the threshold being analyzed). It would be feasible to implement a conceptually more satisfying approach that does not have this potential bias, but for the reasons stated above, it will matter very little in most applications."