Optimal Tax Policy with Misreporting: Theory, and Evidence from Real Estate

Santosh Anagol Vimal Balasubramaniam

Benjamin B. Lockwood Tarun Ramadorai Antoine Uettwiller*

March 20, 2025

Abstract

We develop a model of optimal taxation and enforcement when policymakers seek both welfare maximization and "tax accuracy," where taxpayers remit the amount statutorily owed under truthful reporting; and characterize the optimal combination of tax rate and enforcement stringency. We apply the framework to the Mumbai real-estate market, a setting with widespread misreporting and a transparent enforcement instrument: government-specified guidance values set the minimum required tax base. Using bunching in reported transaction values around the guidance value to identify under-reporting, we estimate elasticities governing optimal enforcement and recover the revealed-preference inaccuracy penalty that rationalizes observed policy. Mortgage-facilitated purchases—subject to additional reporting requirements—exhibit less misreporting, consistent with financial-market incentives complementing tax enforcement.

*Anagol: Wharton School, University of Pennsylvania. Email: anagol@wharton.upenn.edu. Balasubramaniam: Queen Mary University of London and CEPR. Email: v.balasubramaniam@qmul.ac.uk. Lockwood: Wharton School, University of Pennsylvania and NBER. Email: ben.lockwood@wharton.upenn.edu. Ramadorai: Imperial College London and CEPR. Email: t.ramadorai@imperial.ac.uk. Uettwiller: Queen Mary University of London. Email: a.uettwiller@qmul.ac.uk. This paper subsumes "A Bad Bunch: Asset Value Under-reporting in the Mumbai Real Estate Market". We thank Cristian Badarinza, Paul Carrillo, Edward Coulson, Sabyasachi Das, Raj Iyer, Venkatesh Panchapegesan,Tanner Regan, Shing-Yi Wang, Caroline Weber, and seminar participants at Imperial College London, UC San Diego, SITE Financial Regulation, Tilburg, Wharton, Johns-Hopkins SAIS, Princeton, WEFIDEV, 2022 Zurich Conference on Public Finance in Developing Countries, 2023 Syracuse-Chicago Webinar on Property Tax and Administration, MoFIR Conference on Banking, and the 2023 AREUEA National Conference for helpful comments. We thank Karthik Suresh, Karan Gulati, and Alexandru Zanca for able research assistance, and Raja Seetharaman at Propstack for data access. We acknowledge funding from Wharton Global Initiatives.

1 Introduction

The central goal of tax enforcement is to promote tax compliance and accuracy, so that taxpayers' actual remittances align with the amount that they legally owe. Inaccuracies arise in the form of "underpayments" if taxpayers pay less than their statutory obligations, e.g., due to evasion or avoidance activities. Inaccuracies may also take the form of "overpayments" in net taxes, e.g., if taxpayers fail to receive credits to which they are entitled, or if they are overbilled due to an error in tax administration. In this paper, we develop a model of optimal taxation and enforcement wherein tax authorities care both about maximizing welfare and about minimizing tax inaccuracies.

We apply the insights of this model to an empirical setting of property transaction taxes with self-reported transaction values. Many countries have attempted to control under-reporting by creating formulaic assessments of property value based on the physical location of properties, setting the tax base as the higher of this government-assessed value and the sales price reported by property buyers. Based on our review of transaction tax policies in the 82 largest cities in the world, 35 of these cities employ this specific system, and variants of this policy that involve some form of "model-based" appraisal of property values are even more widespread.[1] One such city is Mumbai, India, where we study the universe of residential property transactions between 2013 and 2022 using a large and granular administrative dataset. In this setting, home values serve as the basis for transaction taxes, capital gains taxes, and annual property taxes; overall, real estate transaction taxes constitute approximately 20% of state government revenues in India.[2]

To complement our theoretical work on tax inaccuracies, we develop a new empirical method to detect the under-reporting of property transaction values. We apply this method to the data, and estimate an 11% under-reporting rate on average over the sample period; this average estimate comes from underlying transactions which exhibit both under- and over-reporting. While substantial, this estimate is lower than the priors inferred from both Indian media and tax authority reports, which have for decades surmised that real estate buyers routinely and substantially under-report valuations as a way of laundering unreported income, such as cash earnings and bribes (so called "black money").[3] The desire to reduce such black money has motivated

[1] Appendix Table A1 presents the cities that employ this guidance value system along with their transaction tax rate. A detailed spreadsheet of valuation systems for the top 82 cities of the world can be found here. We briefly discuss the broader context of cross-country variation in systems of property appraisals for tax purposes later in the paper.

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

[3] See, for example, Indian Department of Revenue's "White Paper on Black Money," 2012, `https://dor.gov.in/sites/default/files/FinalBlackMoney.pdf`.

economically massive policy interventions, such as India's 2016 demonetization.

The statutory requirement in the Mumbai setting is that property purchasers report the true market transaction value to the government. If the buyer's reported value falls below the government assessed value (also known as the "guidance" value), then the tax base is set equal to the guidance value. An important mode of evasion in 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 serve as a structural foundation for our work on optimal taxation, we model agents' reporting incentives by adapting and extending the classic tax evasion model developed by Allingham and Sandmo (1972) and Srinivasan (1973) to our setting. In this framework, the choice of reported value results from trading off penalties in the event of detection against the tax saving from under-reporting. Buyers with a low enough expectation of detection probabilities in the event of an audit simply report the guidance value regardless of the true market transactions price, resulting in "bunching" of reported transactions prices at the guidance value.

This model prediction shows up clearly in the data, where we uncover prominent bunching of self-reported property transaction values at government-assessed guidance values: 7.3% of reported transaction values bunch within 1% of the guidance value, and an additional 13.9% report more than 1% below the guidance value. Interestingly, 78.8% of transactions (corresponding to 26.5% of transaction tax revenues) are reported at 1% or more above the guidance value, suggesting that penalties and/or moral concerns impede a large part of the transacting population from simply reporting the minimum government value. We show similar bunching of reported transaction values at guidance values in data from Sao Paulo, Brazil, in Appendix Figure A1, suggesting that our methods are applicable elsewhere.[5]

While bunching of reported values at government guidance values is consistent with under-reporting, it could also be consistent with truthful reporting for at least two reasons. First, while it is unlikely, infrequently updated 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.[6] If true underlying market values were observable, we could easily dis-

[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`.

[5] We replicate this plot based on exhibits originally shown to us by Thiago Scot as part of their working paper Rocha, Scot and Feinmann (2023, mimeo).

[6] While we are unaware of direct evidence that market prices anchor on guidance values Genesove and Mayer (2001) and Andersen, Badarinza, Liu, Marx and Ramadorai (2022) show evidence that prop-

tinguish between these alternative explanations, but the existence of under-reporting itself renders market values difficult to observe.

To address this challenge we match our administrative data on reported and guidance values to a third-party provided price dataset of new buildings developed and sold during our sample period.[7] These data serve as a proxy for market values, and are based on collecting pricing sheets and other marketing materials from developers. Unlike the distribution of reported property values around government-assessed guidance values, the distribution of this proxy for market values is smooth, with no bunching at guidance values. The visible difference between the distributions of self-reported and proxy market values constitutes the basis of our new technique to detect under-reporting.

While these third-party price data are purchased and used by banks, developers, and investors in the real estate space, we nevertheless expect that they are an imperfect proxy for market values. We therefore adapt our approach and conduct several additional checks to account for measurement error. First, we consider the case of "classical" measurement error, in which third-party listing prices are noisy but unbiased measures of true market prices. We show using simulations that mean-zero noise in the proxy can in theory rationalize both truthfully-reported values that bunch at guidance values and a smooth distribution of the proxy around guidance values, under the (strong) assumption that guidance values perfectly align with market values. We deal with this empirical confound using the insight that aggregation "smooths out" the noise in the market-value proxy. We show that if reporting is truthful and guidance values perfectly track market prices, there will be small differences between *aggregated* reported transactions values and *aggregated* (noisy) market-value proxy values. In contrast, as we find in the data, under-reporting delivers a distribution of aggregated reported values that lies below the distribution of aggregated market value proxies.[8] Second, we re-estimate our findings on a more restricted dataset (70% of the full sample) of "exact matches" across reported transactions and third-party price data and confirm the robustness of our estimates. Finally, as discussed more fully below, we inspect reporting behavior around revisions in government guidance values, an exercise which reveals additional evidence strongly consistent with under-reporting.

---

erty sellers anchor on original purchase prices, and Garmaise and Moskowitz (2004) suggest that assessed values may be informative about market prices.

[7] Multiple real estate analytics firms in India collect and sell these data, and we expect such data exists or will emerge in many developing country cities in the coming years.

[8] We extend this insight to the case when the market value proxy is both a noisy and biased measure of true market value. Here, we show that the form of measurement error required to rationalize the patterns in the data is implausible, requiring a very specific and unusual pattern of discrepancies between reported values and true market values.

In the data, we estimate under-reporting rates of approximately 25% for "buncher" transactions that report exactly equal to the guidance value, and roughly the same under-reporting rate for transactions with reported values less than guidance values. The remaining estimated under-reporting rates decline linearly with the extent to which buyer-reported values exceed the guidance value; for buyers reporting greater than 50% above the guidance values we estimate zero under-reporting. Overall, our estimated 11% under-reporting rate, when applied to Maharashtra state, roughly translates to an annual loss for the state government of US$ 475mn lower tax revenues from property transactions in 2021, which is approximately 1% of total state government revenue.

Governments have strong incentives to set accurate guidance values, but it is unlikely that they perfectly capture true underlying value, given the fairly broad geographical regions to which individual guidance values apply (this is especially true in our context, given the density of physical properties in Mumbai). Moreover, guidance values are infrequently updated, leading to staleness and inaccuracy in an environment of changing market prices. We use this insight to uncover additional behavioral responses using our approach, studying the revision of government-assessed values (which occur in three of the nine years in our sample) across multiple neighborhoods. We detect large spikes in the volume and value of registered transactions in the days and months immediately before scheduled guidance value changes, consistent with gaming behavior by agents rushing to register transactions right before these changes; we also find evidence of agents back-dating transactions to pre-change dates to exploit lower guidance values. Quantitatively, these spikes generate increases in under-reporting rates of approximately 6%-12% in months immediately preceding guidance value changes.

While we find significant under-reporting, our estimates also suggest that many agents report truthfully. Though anecdotal evidence suggests that actual audit probabilities are close to zero, the mere threat of audit is potentially sufficient to generate truthful reporting (Kleven, Knudsen, Kreiner, Pedersen and Saez 2011). To investigate this issue further, we use our workhorse model to interpret observed reporting behavior in terms of agents' underlying beliefs about detection probabilities. Viewed through the lens of the model, many agents in the economy believe that auditing is sensitive to their reporting behavior, which echoes experimental evidence in Bergolo, Ceni, Cruces, Giaccobasso and Perez-Truglia (2023) where the possibility of tax audits generates fear and induces probability neglect.

We find that under-reporting rates in secondary market resale transactions are higher than those in primary-market developer sale transactions. This finding is con-

sistent 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. Moreover, we find that under-reporting rates are lower for larger transactions, likely in response to higher perceived detection probabilities for more visible transactions with greater absolute economic penalties for under-reporting.

How do governments optimally set tax policy in this environment to balance their twin objectives of maximizing welfare-weighted tax revenues and minimizing inaccuracy? Raising guidance values will increase revenues (assuming small extensive margin elasticities, which we find in our empirical work), but will inevitably lead some individuals to overpay, thus raising inaccuracy. We develop this intuition further, writing down sufficient statistics for optimal tax policy in this dual-objective setting.

To evaluate this model in the data, we first structurally estimate the parameters that govern the taxpayers' misreporting, which we allow to vary across individuals in our simple model of reporting incentives. We find that a parsimonious model delivers a good fit to the data, i.e., we are well-able to match the empirically observed distribution of reported values, counts and the under-reporting rate to the corresponding model-implied values using the estimated structural parameters.

In a second step, we estimate the weight that the government places on its inaccuracy minimization objective. We do so using the estimated misreporting aversion distribution from the first step to predict reported values, and vary the unobserved weight on the inaccuracy penalty to match model-implied government guidance values with the observed government guidance values in the data across different areas of Mumbai.

We estimate that the government has a high aversion to inaccuracy, and is particularly averse to overcharging taxpayers. Our estimates reveal that the government is willing to forego ₹1.05 in tax revenues for each ₹1 of over-payment by taxpayers. To validate the model and these estimates, we generate out-of-sample predictions of government guidance value, and use these predictions to forecast the direction of revisions to the government guidance value in the data. More specifically, in areas where the model suggests that observed guidance values are too low (high), we later observed upward (downward) revisions to guidance values out of sample. This exercise confirms the ability of the model to captures the government's decision rule.

As a final exercise, we merge administrative data on mortgage values with our main dataset for a smaller subset of reported transactions. With the caveat that this process allows us to match roughly a third of (possibly selected) transactions, we find the inter-

esting pattern that reported values of transactions with low 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. We also estimate the greatest under-reporting for properties with mortgages from cooperative and public-sector banks, and least for those from private and foreign banks; mortgages from banks with high overall non-performing loans are also associated with high bunching behavior. As we demonstrate in a simple extension to our model, these patterns can be rationalized by borrowers trading off higher reported values to relax financial constraints associated with mortgage credit against the greater tax bill associated with doing so; we note that assortative matching between household evasion types and bank screening technologies can also rationalize these patterns.

The remainder of this paper is organized as follows. The remainder of this section reviews related literature. Section 2 sets up a simple model to guide our empirical work. Section 3 describes the institutional background for property valuation for tax purposes in our setting and elsewhere, and the data. Section 4 documents our empirical results, including a discussion of how measured under-reporting varies in settings and sub-samples with different economic incentives to under-report. Section 5 structurally estimates our model and validates it out-of-sample, and Section 6 concludes.

1.1 Related Literature

Our conceptual framework builds most directly on the literature on optimal tax administration and enforcement. Keen and Slemrod (2017), building on the early work of Sandmo (1981) and Mayshar (1991), present a conceptual framework in which a policymaker maximizes social welfare by adjusting both tax rates and tax enforcement parameters such as audit probabilities. This framework identifies as a key sufficient statistic the *enforcement elasticity of tax revenue*, which quantifies the change in revenues—accounting for behavioral responses of taxpayers—of adjustments in enforcement policy. We build on this work by accounting for policymakers' potential desire for compliance and accuracy in tax administration, as distinct from pure revenue considerations.

Our empirical application also contributes to the large positive literature documenting tax evasion and enforcement responses. 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, Knudsen, Kreiner, Pedersen and Saez (2011) review 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 apply 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. Hendren, Sprung-Keyser and Stuart (2023) estimate the welfare and revenue impacts of tax audits in the U.S.

Our empirical application also links our work to the bunching literature, developed by Saez (2010) and surveyed in Kleven (2016), which uses an optimizing model to translate bunching 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 features a "kink," in the sense that the marginal tax rate below the guidance value is zero, but then discretely increases to 5% (the transaction tax rate) for the marginal rupee reported above the guidance value. We later use this insight to estimate "evasion elasticities" using the standard approach.

In terms of possible remedies, Pomeranz and Vila-Belda (2019) survey research with tax authorities, focusing 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 relates to the literature on transaction taxes, which has mainly 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, and other market participants eliminates 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, and while Slemrod, Weber and Shan (2017) finds evidence of house price manipulation to avoid a higher average transaction tax rate in Washington D.C., they note that major tax evasion or avoidance behavior is unlikely in their setting given the relatively small change in average transaction tax rates studied. Sood (2020) highlights high transaction taxes as one component of land market frictions in India (and likely other developing countries) that may hinder productivity overall.

A related literature studies ongoing property taxation, finding that assessed values for property taxes can systematically diverge from recent transaction prices, with important distributional consequences (Avenancio-León and Howard 2022, Regan 2023). In this context, we are the first to analyse under-reporting in a system of government-assessed values. A key difference in our context is the statutory obligation for homeowners to report the true market value; in typical advanced economy property tax contexts, homeowners are required to pay tax on the government's assessed value, even if assessed value differs from market value (Amornsiripanitch 2020, reviews this literature).

A small but growing literature studies under-reporting behavior in China, though these papers do not study government guidance values. Fan, Wang and Zhang (2022), Agarwal, Li, Qin, Wu and Yan (2020), and Agarwal, Kuang, Wang and Yang (2020) use data on underlying transaction prices collected by Chinese real-estate brokerages (which serve as the basis for brokerage commissions). In the Mumbai setting, similar to many cities in developing countries, such administrative data are not recorded on true underlying transaction prices. Our approach can be applied more broadly to detect under-reporting whenever the government sees reported and guidance values and can source (as we do) measures of market value or listing prices from analytics companies or online listings portals—in line with the broader agenda of technology-enabled improvements to developing-country tax collection (Okunogbe and Santoro 2022).

Our analysis of property value under-reporting also connects to early analyses of black money in Indian real estate, which studied the government's pre-emptive purchase provision, under which the central government tax authority was allowed to purchase any property at 15% above the reported value, creating strong incentives for accurate reporting.[9] In statute, the government was supposed to randomly select property transactions to determine whether to exercise this right, though most sources suggest 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,[10] 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

[9] 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.

[10] 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.

1987, Gopalakrishnan and Das-Gupta 1986). In a survey article on Indian transaction taxes, Panchapagesan and Karthik (2017) notes that there have been no more recent aggregate estimates of black money in Indian real estate transactions.

Our mortgage results, though based on a limited sample of transactions, highlight the observation in Basu (2015) that home-value under-reporting 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 using data from one housing brokerage to focus on a buyer's trade-off between under-reporting to avoid transaction taxes and over-reporting to obtain greater mortgage credit. Our results on mortgage bank ownership and under-reporting also connect 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).

Finally, our (relatively low) estimate of the property value under-reporting rate is potentially useful for developing country studies on how neighborhood change, transportation infrastructure, and zoning reforms affect real estate prices, as these studies often use guidance property values in lieu of frequently unavailable market price data (Anagol, Ferreira and Rexer 2021, Tsivanidis 2019, Gechter and Tsivanidis 2020, Harari and Wong 2018).

2 Model

We model taxpayers who optimally report values that may differ from true underlying asset values. This model of misreporting serves as the foundation for our analysis of optimal tax policy when agents have incentives to misreport.

Canonical optimal taxation models typically operate under the assumption that tax authorities are solely concerned with maximizing social welfare (Ramsey 1927, Mirrlees 1971, 1986) when faced with individual agents who rationally optimize their economic decisions. Our goal here is to study how governments optimize policy when they are motivated *both* by welfare-maximization and "inaccuracy minimization" imperatives. This latter channel is a type of fairness criterion: Governments may not wish to obtain tax over-payments from taxpayers who truthfully report asset values, or to allow misreporting agents to substantially under-pay taxes relative to what they would owe if they truthfully reported.

2.1 A Model of Taxpayer Reporting Behavior

We adapt and extend the classic tax evasion model developed by Allingham and Sandmo (1972), Srinivasan (1973) and Yitzhaki (1979) to model individual taxpayer behavior. A taxpayer purchases an asset (in our empirical application, real estate) for market price m and then chooses the value r (which could differ from m) that they report to the government. The property has an associated government-assessed guidance value which we denote by c (for "circle rate"; the nomenclature used in the Indian context, which we describe in detail below), and τ is the transaction tax rate; the transaction tax liability is therefore $\tau \times max(r, c)$, i.e., the tax base for the asset is the maximum of the reported value and the guidance value unless the individual chooses to appeal. If the appeal is granted, the individual may pay his reported value with $r < c$.

In practice the guidance value c is computed on a property-specific basis using a coarse set of observables such as property size and geographic location, and there is a single policy parameter that serves to rescale all guidance values. To simplify notation in what follows—and in our empirical figures to follow—we restrict focus to a set of properties with a common guidance value, or equivalently, we interpret the variables c, r, and m as having been residualized on the observables that determine c. Thus we can treat c as a single policy-determined constant in our model analysis.

The taxpayer's optimal reporting amount after making a transaction is determined based on the trade-off between saving money in transaction taxes paid versus expected penalties of under-reporting and transaction costs of getting audited (which is more likely if they under-report). Here we present the specific functional form we use later to structurally estimate our model of under-reporting (which will be an important ingredient in ultimately estimating the government's value on tax accuracy).

The taxpayer's objective function is as follows:

$$\min_r \left\{ \tau r + \pi \left(\left(1 - \frac{r}{m} \right)^{\rho} n\tau(m - r)^{+} + t \right) \right\} \tag{1}$$

The buyer pays a tax on the reported transaction value, τr, with certainty. With probability π the buyer is audited, in which case the buyer incurs two types of costs. First cost is the expected cost of penalties if the audit actually detects under-reporting. The buyer believes detection occurs with probability $(1 - \frac{r}{m})^{\rho}$, in which case he incurs an additional penalty of $n \cdot \tau(m - r)$, where n represents the penalty multiplier applied to the evaded tax balance $\tau(m - r)$. Thus the parameters π and ρ together determine the probability that an audit detects the under-reporting by a buyer.

We also assume that the buyer must pay an audit transaction cost t - this incorpo-

rates lawyer and accountant fees as well as psychic costs associated with being audited for under-reporting.

When an individual reports $r \geq c$, we assume the audit probability $\pi_{r \geq c} < 1$. For now we assume that this audit probability does not vary across individuals who report $r \geq c$, although we relax this assumption later in this section. The lower bound of $\pi_{r \geq c}$ is governed by how the penalty multiplier interacts with the audit probability. Since $n\pi_{r \geq c} > 1$, the lower bound for $\pi_{r \geq c} = \frac{1}{n}$ – otherwise the agent would have an incentive to under-report because the expected penalty of under-reporting one rupee would be less than paying the tax on an additional one rupee reported.

In the case where a buyer reports $r < c$, we assume that this implies the buyer has chosen to appeal, which triggers an audit automatically. Therefore, $\pi_{r<c} = 1$. Figure 1a plots the shape of π for values of r over and under c. We make this assumption as it creates a strong discrete incentive for agents to report $r = c$, relative to $r < c$; this bunching pattern is observed in the data we discuss in the next Section. Note that given a buyer that reports $r < c$ is automatically appealing, they will not actually pay c in our framework, but instead pay τr plus the expected penalty $(1 - \frac{r}{m})^{\rho} n\tau(m - r)^{+}$ plus the transaction cost of being audited.[11]

In the instance where the buyer under-reports, i.e., $r < m$, their belief about the probability that an audit detects the actual transaction value m is:

$$\left(1 - \frac{r}{m}\right)^{\rho} \tag{2}$$

This parameterization implies that the detection probability ranges from 0 (under truthful reporting with $r = m$) to 1 (under full misreporting with $r = 0$). The parameter $\rho \in [0, +\infty)$ controls the buyer's beliefs about the detection probability between these extremes. Figure 1b plots the profile of this detection probability belief for various values of ρ. A low value of ρ—like 0.25—implies that buyers perceive the detection probability to be high unless the report r close to the truthful transaction value m. A higher value of ρ—like 5—implies that the detection likelihood is close to zero even for those who report r well below the truthful value of m.

Including a separate π (probability of audit) and ρ detection of under-reporting conditional on audit allows our framework to ultimately match to key facts in the data. First, we see a substantial amount of bunching at c and excess mass of reporting above (but not below) c. The probability of audit parameter π changing from 1 below c to

[11] We do not have administrative data on appeals and their outcomes. Anecdotally, our understanding is that one reason buyers report $r < c$ is that they intend to appeal a guidance value that is set higher than the property's market value. In theory it is possible that some buyers may report $r < c$ but then choose to not appeal, but instead just pay τc as their tax – for example they want the "truth" reported on the sales deed even if the system forces them to pay more.

less than 1 below causes agents in our model to bunch in this fashion. Second, our data includes many cases where buyers with the same guidance value c and market value m report different amounts. Individual heterogeneity in beliefs about under-reporting detection naturally deliver this kind of variation.

From equation (2), when the taxpayer's preferred reported transaction value r is interior it satisfies the following first-order condition:

$$r^* = m \times \left[1 - \left(\frac{1}{\pi n(1+\rho)}\right)^{1/\rho}\right] \tag{3}$$

The buyer's optimal choice of r^* depends on ρ, t and π. To illustrate this dependence, Figure 2a plots optimal reported transaction values normalized by the guidance value (i.e., r^*/c) for different values of ρ, assuming that $m = 2c$ and $t = 0$. The audit probability π jumps discontinuously from $\pi_{r \geq c}$ to $\pi_{r<c} = 1$ at the point $r = c$.

Diffuse Bunching: For a range of values of ρ, reported values bunch exactly at the guidance value, i.e., $r/c = 1$ on the y-axis. In reality, we might expect bunching to be less sharp than this frictionless model predicts. In order to allow for this formally, we allow for the possibility that a buyer i believes that the probability jumps from $\pi_{r \geq c}$ to $\pi_{r<c}$ at a value c_i that is near (but not exactly equal to) c, i.e., $c_i = c + \epsilon_i$. We then have:

$$\pi(r, c) = \begin{cases} \pi_{r<c_i} = 1 & \text{with } r < c_i = c + \epsilon_i \\ \pi_{r \geq c_i} < 1 & \text{with } r \geq c_i = c + \epsilon_i \end{cases} \tag{4}$$

If $m > c$, then in the presence of this added friction the discontinuity in π occurs at $c + \epsilon_i$, which is assumed to be greater than c. If instead $m < c$, then t buyers choose either to report report the truthful transaction value $r = m < c$ or to report a higher-than-truthful transaction value $r = c_i$ in order to avoid incurrent the audit cost t.

Although t does not appear in equation 3, it is particularly salient for buyers with $m < c$ but close to c. In other words, when buyers have a high t, but their true m is lower than c, they may choose to report at c to avoid incurring the t- costs, which we will highlight below. Figure 2b plots the r/c values on the y-axis for various values of ρ, with and without t. A small $t = 0.01$ increases the region for under-reporting, and thus plays a role in magnitude of bunching that arises from the mass under and over c.

2.1.1 The Effect of External Financial Constraints

We now highlight how taxpayer asset-value reporting behavior changes in the presence of external financing constraints. In many asset markets, collateralized loans take the official reported/registered value of a property into account at the point of loan issuance. This links the decision to misreport with the extent of financial constraints. In the real estate setting that we consider, the financial incentive to truthfully report (or potentially over-report) arises from the desire to unlock additional mortgage finance.[12]

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 (2022). Consider a bank that is willing to lend $(1-\gamma)r$, where γ is the down-payment constraint governed by macroprudential regulations, and/or by increases in lender-demanded mortgage credit premiums with loan-to-value (LTV) ratios. 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 we model for simplicity as linear in the shortfall.[13] In this simple setup, borrowers must report more, or incur penalties (e.g., reducing house size relative to desired levels, or coming up with additional, expensive unsecured/informal financing to bridge the gap) which scale with the size of the financial shortfall. The individual taxpayer's problem thus becomes:

$$\min_{r}\left\{\tau r+\pi\left[\left(1-\frac{r}{m}\right)^{\rho}n\tau(m-r)^{+}+t\right]+\mu\left[(m-d)-(1-\gamma)r\right]\right\} \tag{5}$$

Equation (5) modifies equation (1) to incorporate the down-payment constraint. The parameter μ determines the tightness of the financial constraint, which we later evaluate as a parameter that governments can influence, for example, by clamping down on unsecured/informal bridge financing. The minimization problem in equation (5) has the following solution:

$$r^{*}=m\times\left[1-\left(\left(1-\frac{\mu}{\tau}(1-\gamma)\right)\frac{1}{\pi n(1+\rho)}\right)^{1/\rho}\right] \tag{6}$$

Equation (6) shows that the financial constraint creates an incentive for the borrower to report more than the guidance value to obtain additional mortgage credit. The buyer trades off the marginally higher transaction tax associated with a higher re-

[12] More specifically, in our empirical 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.

[13] The down-payment constraint affects transactions irrespective of whether true $m>c$ or true $m\leq c$. For brevity, we focus on the $m\geq c$ case in this section.

ported value against the marginal benefit of increased mortgage credit. An increase in the permissible loan-to-value ratio (raising $(1-\gamma)$) or tightening the constraint on obtaining funds outside the formal mortgage system (increases in (μ)) can both induce more truthful reporting behavior. This financial market channel complements (and takes pressure off) the use of tax policy instruments to encourage truthful reporting.

The empirical method that we later develop measures misreporting as the joint effect of all of these incentives. While the primary focus of the paper is on the relationship between tax policy instruments and misreporting, we also later provide illustrative evidence of the financial constraint channel later in the paper.

The tax implications of misreporting are in reality more complex than this simple economic framework. For example, there may be additional incentives for buyers to under-report to reduce ongoing taxes on property/asset values.[14] And there is potentially an opposite incentive, for asset buyers to "over-report" to increase their current cost basis and reduce future capital gains taxes. Buyers that focus on minimizing the current transactions tax burden have under-reporting incentives that are well-aligned with sellers who wish to report depressed sale values to minimize capital gains taxes. Financially-constrained buyers and forward-looking buyers that care about the capital gains tax basis more than transactions taxes, however, may have opposing incentives to sellers, thus reducing opportunities for collusion to report low asset values. While we do consider the effect of financial constraints, we do not consider such capital gains tax incentives for buyers in what follows, primarily because the horizons over which housing transactions occur are very long, depressing the relative importance of this channel.

2.2 Optimal Government Policy

The government chooses values for the tax rate τ and the guidance value c to maximize its objective, taking as given the taxpayer's optimization behavior over the reported value r^* described by the model above.

The selection of the optimal guidance value amounts to a choice of optimal enforcement along the lines of Keen and Slemrod (2017).

To illustrate the trade-off that the government faces, Figure 3 shows two different levels of c (c low in the top panel, and high in the bottom panel), superimposed on the distribution of underlying market values m. Although the government may be

[14] One appealing feature of the institutional setting for our empirical work is that annual property taxes are only a function of the government's assessed value, not the reported value, so concerns about annual property taxes do not affect household's reporting behavior at the time of transaction. See `https://housing.com/news/guide-paying-property-tax-mumbai/` for details.

tempted to set c high to increase tax revenue, this comes at the cost of forcing taxpayers with $m < c$ to either appeal or overpay taxes, paying $\tau \times c$ rather than $\tau \times m$. We model how the tax authority balances these twin imperatives.[15]

More formally, given a vector $\mathbf{x} = (m, \rho, \epsilon, t)$ of the market value, under-reporting detection, location of jumps, and audit cost, the tax authority collects revenues of

$$\tau \int (1 - g(m)) \times r^*(c, \mathbf{x})\, f(\mathbf{x})\, d^4\mathbf{x} \tag{7}$$

The total revenue raised is the tax rate τ multiplied by the sum of all transaction values r^* integrated over both the m distribution and the distributions of ρ, ϵ and t. We allow each dollar of revenue to have an associated welfare-weight $(1 - g(m))$.[16]

As discussed, we allow the tax authority to care about the overpayment of taxes. Suppose an individual taxpayer owes τm in taxes, but reports r^*, and pays a tax of τr^*. Over-payment is the case when $m < r^*$. In this context, the government also cares about appeals. Setting a guidance value too high would lead to an increase in costs.

We model this using a general weighting function $\zeta(\cdot)$, and allow the government to place weight ψ on this inaccuracy in its optimal tax policy determination. For example, the government could care only about over-payers and in that instance $\zeta(r^*(m, \mathbf{x}), \mathbf{x}) := \psi \max(c - m, 0)$. In aggregate, the government's inaccuracy penalty can be written as:

$$\tau \int \zeta(r^*(c, \mathbf{x}), \mathbf{x})\, f(\mathbf{x})\, d^4\mathbf{x} \tag{8}$$

Finally, putting it all together, the government's overall problem, factoring in its twin objectives, is:

$$\max_c \underbrace{\left(\tau \int (1 - g(m)) \times r^*(c, \mathbf{x}) f(\mathbf{x}) d^4\mathbf{x}\right)}_{\text{welfare-weighted tax revenue}} - \underbrace{\left(\tau \int \zeta(r^*(c, \mathbf{x}), \mathbf{x}) f(\mathbf{x}) d^4\mathbf{x}\right)}_{\text{inaccuracy penalty}} \tag{9}$$

[15] There is also an extensive margin elasticity for the tax authority to consider, i.e., the total number of asset transactions may fall if c is raised sufficiently high. In our empirical application, we find that this is essentially zero for the tax that we consider and so do not consider this (see Appendix G). For completeness, in future drafts, we plan to include this extensive margin elasticity in the government's optimization problem.

[16] To select the welfare weights $g(m)$, we draw motivation from the fact that under the common specification of logarithmic utility of consumption, the marginal utility of consumption is proportional to inverse consumption. Using market transaction value m as a proportional proxy for taxpayers' level of consumption, we therefore select $g(m)$ to be proportional to $1/m$. Imposing the constraint that these welfare weights average to one, as is the case when there are no income effects on labor supply, this implies $\frac{\frac{1}{m}}{\int_0^\infty 1/m\, f(m) dm}$.

The general framework laid out in this section enables us to study optimal tax policy in the presence of misreporting, with individual taxpayers rationally optimizing their reporting behavior, and the government balancing its twin objectives of maximizing revenue and minimizing overpayment of taxes. We now turn to describing the institutional context, and empirically applying this framework to a rich dataset tracking the real estate market of Mumbai, India.

3 Institutional Background and Data

3.1 Institutional Background

Systems of property valuation for taxation purposes vary around the world, and can be broadly classified into two types. In the first type, taxation authorities generate "decentralized" (i.e., property-specific) assessed values, using a number of different approaches. In some jurisdictions, assessors determine property valuations using a combination of site visits and comparable analysis. In others, assessors determine the valuation by inputting features of the property (e.g., residential/commercial, square footage etc.) into a hedonic model.[17]

Physical assessment of individual properties can produce greater accuracy in determining market value, which is helpful given substantial unobserved heterogeneity in property quality even within small regions. Moreover, in decentralized assessment systems the guidance value is the statutory tax base; meaning that the owner has no reporting obligations. However, such individual assessments can be costly to implement and/or inaccurate given the scarcity of qualified assessors, and they can also be subject to manipulation, for example when assessors are bribed to lower assessments (Khan, Khwaja and Olken 2016). An alternative method of decentralized valuation is the so-called "self-assessment" method proposed by Harberger (1965), which encourages truthful self-reporting by giving the state the right to purchase the property at the property owner's self-assessed value.[18]

The Indian tax authorities (and many other jurisdictions, including, among others, Brazil, Colombia, Mexico, Thailand, Indonesia, Philippines, and New Zealand, see Appendix Table A1) utilize an alternative "centralized" system of property valuation for

[17] For example, the Danish system of tax assessor valuation at different points in time adopted both property-specific and model-based property assessments (Andersen, Badarinza, Liu, Marx and Ramadorai 2022). In the United States, depending on the location and property type, local governments often conduct comparable-sale based assessments to determine the tax base. See, for example, this description of New York state property taxes, accessed February 2023.

[18] Chang (2012) argues that even with these incentives property owners in Taiwan greatly underreported property values, because the probability of the state actually exercising the right was too low.

taxation purposes. In such systems, the authorities assign location-specific (usually per-square-foot) valuations as a lower bound tax base for all properties in the physical location, and periodically update these valuations as market prices evolve. The statutory requirement is that owners report the true market value of their property, with the tax base set as the maximum of the lower bound tax base (henceforth "guidance value") and the owner-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 (note that the true market value can be different from either or both of the guidance and reported values). Owners also generally have some form of (costly) recourse available to prove that lower valuations than the guidance value can be justified. Such centralization is more cost-effective than property-specific assessments, and can lower the probability of captured assessors. However, guidance values (especially when infrequently updated) can be inaccurate measures of market value.[19]

Conversations with market participants, and reports from regulators in the Indian centralized tax assessment system suggest that under-reporting follows a typical pattern.[20] The buyer (usually an individual) and seller (either an individual 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 (a frequent choice is to report exactly the guidance value; interviews with market participants suggest auditing is virtually non-existent as long as the reported value equals or exceeds the guidance value.).[21] To prevent detection of the under-reporting by paper/digital trail, the gap between the transaction price and the reported price is paid physically in cash; an alternative is that part of the true transac-

[19] In parts of the U.S., tax authorities apply formulaic growth rates to historical assessed values, which is a similar approach (e.g., California's Proposition 13, see http://Santa-Clara-property-assessments accessed February 2023.)

[20] See, for example, `https://timesofindia.indiatimes.com/blogs/law-street/black-money-does-the-devil-lie-in-real-estate/`. A similar discussion of under-reporting real estate transactions in rural France is recounted in Mayle (2000): "This is the two-price purchase, and a typical example would work as follows: Monsieur Rivarel, a businessman in Aix, wishes to sell an old country house that he inherited. He wants a million francs. As it is not his principal residence, he will be liable for tax on the proceeds of the sale, a thought that causes him great distress. He therefore decides that the official, recorded price—the prix déclaré—will be 600,000 francs, and he will grit his teeth and pay tax on that. His consolation is that the balance of 400,000 francs will be paid in cash, under the table. This, as he will point out, is an affaire intéressante not only for him, but for the buyer, because the official fees and charges will be based on the lower, declared price. Voila! Everyone is happy."

[21] While the law states that guidance values should be formulaic, following centrally assigned guidance values, 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 also detects under-reporting arising from such inspector discretion.

tion price can be misinvoiced as a higher "service charge" (this is especially prevalent in individual-developer contracts). This ensures that the bank records associated with the transaction are in agreement with the reported value on the registration, which is important, since buyers and sellers are required to report their tax-identification numbers to enable cross-verification with transactors' bank accounts. Typical methods of obtaining large sums of cash include accrued currency from business operations, a sequence of smaller withdrawals from bank accounts over time, or writing a check as a "gift" to a friend or relative in exchange for the cash. In some cases, cash funds can also be sourced or earned completely outside the tax net; this is referred to as "black money" in the Indian context.

3.2 Data

While market participants state 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 the appetite for under-reporting (i.e., a lower overall price for a higher cash share). However, it is generally not an easy task (and one to our knowledge has not been undertaken) for the tax authority to acquire systematic data on market prices (by interrogating brokers and/or developers, or scraping and painstakingly matching data from online housing listings) and conduct the kinds of comparisons that we undertake in this paper.[22] We now describe the specific data sources that we compile and merge for this purpose.

3.2.1 Registrar Data

The first dataset that we employ 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 guidance value; 3) the transaction tax paid; 4) the property's floor space area; 5) information about buyer and seller type (corporate or individual); 6) the transaction date; 7) the registration date. These data are publicly available from the registrar's website, and cover all available transactions between 2013 and 2018. See Appendix Section I.i for a complete description of this data and example documents.

We augment the IGR administrative data with data provided by Propstack Analytics which cover the period from 2013 to 2022. Propstack is a for-profit real estate firm that uses transactions to provide pricing and ownership information via its Zapkey

[22] Major listing websites include `www.housing.com`, `www.99acres.com` and `www.proptiger.com`.

data platform.[23] In addition to the IGR variables described above, Propstack Analytics also provides us: 1) an indicator for whether a property was 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 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 A2 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.[24]

Between IGR and Propstack Analytics, the data cover the universe of registered residential real estate transactions in 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.[25] There is also comprehensive spatial information available for this region, which we use in our empirical analysis.

3.2.2 Propequity Data

Our independent source of market price data (i.e., the proxy p measuring m) comes from Propequity, a real estate analytics firm that maintains a subscription real estate information portal for the Indian real estate market.[26] 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

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

[24] The government guidance value is determined by multiplying the guidance value (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 guidance value 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.

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

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

over locations. Over the time period 2013 to 2022 this dataset includes information 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 and 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.[27]

These price data serve as our proxy p for m, the true market value of transactions in the administrative data. We match Propstack transaction level data to Propequity price data, using the name of the project in which the transaction occurs, 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 Propstack transactions are exact project-level matches to Propequity. We match the remaining 40% of transactions to the most geographically proximate Propequity project. Figure A3 documents the match quality in the data. Overall, 95% of all registered transactions are matched to Propequity transactions within 500 meters of the latitude-longitude of the purchase transaction, and we eliminate from our analysis successful matches that are more than 1 kilometer away from the location of the associated Propstack transaction. Figure A4 shows the spatial distribution of the transactions

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

in our final sample for study.[28] While the overall match quality is high, we carefully investigate effects of measurement error in p in greater detail later in the paper.

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 made by individuals. The average property is 76 square meters (818 square feet) in size. The average reported value r over the sample is US$ 321,770. The average government guidance value c 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%, though we approach this more rigorously below.

Appendix Figure A5 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 c values 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. As mentioned earlier, the positive relationship between p and c lends credence to the idea that guidance 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 and Wong 2018). We next turn to documenting bunching and computing a measure of property value under-reporting using the data.

4 Empirical Results

Figure 5a plots the empirical distribution of $\frac{r-c}{c}$ around zero; to our knowledge, the first analysis of reporting behavior around guidance 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%.

The plot reveals that 13.9% of households report more than 1% below the government-assessed value; the average (median) r for this group is 24.9% (14.7%) lower than c. On the face of it, this is evidence that c imperfectly tracks m, since a non-negligible fraction of households is willing to pay verification costs to certify $m < c$. Second, there is a clear spike in r at government-assessed c.[29]

The green line with triangles in Figure 5a shows how p, the Propequity proxy for market values is distributed around c. The $\frac{p-c}{c}$ distribution shows no bunching at

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

[29] Appendix Figure A6 presents a version of Figure 5a with one-tenth of 1% bins, showing that the bunching of transactions at the guidance value is distinct and sharply identified.

zero, exhibiting a smooth distribution centered at roughly 30%. To the right of the bunching region, the market value distribution resembles a right-shifted version of the $\frac{r-c}{c}$ distribution - consistent with tax inaccuracy in the form of under-payment of taxes relative to market values. There are a substantial number of transactions where we estimate $\frac{p-c}{c} < 0$, which is consistent with the over-payment of taxes relative to market values. Relative to the distribution of r, the distribution of p has more mass between -0.25 and 0 away from c. One explanation for this, consistent with the model, is that there are incentives to bunch at zero even for those transactions whose $m < c$. A second explanation is that measurement error smooths the p distribution around c.

Overall, the raw data suggest both the under- and over-payment of taxes is common in this context, and motivates our exploration of tax accuracy as important driver of tax policy. If we had access to a transaction-level measure of m, we could directly estimate under and over-payments without inspecting bunching. That said, understanding how a bunching-based strategy to detect misreporting would work with perfectly observed m is a helpful benchmark for evaluating how the strategy would operate in the more realistic case of observing a proxy for m. Measurement error is a natural challenge with proxies for m. Section B in the Appendix evaluates the issue of how much an imperfect proxy p for m can influence our inferences of over- and under-payment, and tests different approaches to dealing with this issue using simulations. This helps us develop an aggregation procedure in the presence of measurement error from genuine under-reporting. The simulation shows that aggregating the transactions into bins associated with different noisy p generates the same approximate mass within these bins as aggregating m, if there is classical measurement error in p (i.e., p is centred around the true m).

We also adopt two alternative approaches to this issue. The first alternative approach directly uses a subset of the data that has plausibly less measurement error where we only use the p for any given transaction when there is an exact match of property locations across the two data sources (reported values to the registrar and our source for the proxy for market value). The second alternative approach rests on the observation that for measurement error to completely explain our results, c must perfectly track m across both time and space, yielding perfect bunching of r at c. All three of these approaches suggest that measurement error alone is unlikely to explain our main finding of under- and over-payments relative to the statutory requirements. Keeping these ideas in mind, we turn to describing the remainder of our baseline empirical results.

Figure 5b applies the aggregation procedure to the actual data, plotting aggregated reported and Propequity values in the data by $\frac{r-c}{c}$ bins. The green triangles indicate the

total amount of Propequity estimated value (p) transacted within 0.02 width $\frac{r-c}{c}$ bins, and the blue dots indicate the total value reported (r) for the same set of transactions within each $\frac{r-c}{c}$ bin. The figure is consistent with Figure B28f, i.e., we confirm that there is under-reporting using the aggregation approach, and the figure reveals that the largest amount of unreported value (the differential between aggregated r and p) comes from bunching transactions that report r exactly equal to c. It is also worth noting that in Figure 5b, the overall overlap of the p and r distributions is tighter than in the plot of transaction counts in Figure 5a, with the sharpest deviation evident between the distributions at exactly c.[30]

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 guidance value, there is no visible gap between aggregate reported value r and estimated market values p. Going back to our model, this finding suggests heterogeneity in $\hat{\pi}$ across the $\frac{r-c}{c}$ distribution that results in buyers reporting more truthfully as we move away from c. Overall, when we aggregate the rupee value of under-reporting across all bins in Figure 5b, we estimate ₹455.1 billion (US$9.1 billion) in under-reported real estate value in our sample over the period 2013-2022.

Figure 5c directly shows the fraction of estimated market value under-reported by $\frac{r-c}{c}$ bin. Transactions with r up to 50% lower than c have an estimated under-reporting rate of roughly 20% on average.[31] There is a sharp discontinuity at zero and a marked change in slope as we move to transactions that report $r > c$, with estimated under-reporting dropping to zero for transactions that with r 50% above c. For properties that report 60% above c, we estimate negative under-reporting (i.e., over-reporting) rates, although these estimates are based on a relatively small number of transactions (the counts per bin in this region are shown in Figure 5a), and the 95% confidence interval includes zero.[32]

[30] For bins with $\frac{r-c}{c} < 0$ (i.e. transactions with reported values less than guidance values) the blue circles in Figure 5b aggregate c. This is because if $r < c$ the tax base is effectively c, i.e., the guidance value, since the government assesses taxes at c pending a successful appeal. Appendix Figure A8 replaces the blue circles in Figure 5b with the total r within each bin, which means that to the right of zero, these figures are identical, but will differ to the left of zero. The figure shows that the total aggregated r values 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.

[31] Note that these transactions are measured as reporting the guidance value (because when $r < c$ the tax base is assumed to be c). A caveat is that it may be that for some of these transactions, $m < c$ and buyers disputed the accuracy of the government-issued guidance value c, so the under-reporting rate at zero (i.e., $r = c$) is potentially more accurate as it is less subject to this confound.

[32] Within each bin, we sample with replacement the same number of observations. This allows us to compute the aggregate values and under-reporting rates for different samples. We then compute the bootstrapped standard errors and then construct the 95% confidence intervals reported in Figure 5c. We use the same procedure for the full sample, and also in other cases where we construct standard errors in this paper.

4.1 External Constraints and Incentives to Misreport

Next, we briefly explore the effects of external constraints on misreporting behavior. Section F in the Appendix explores how under-reporting varies by whether the counter-party is a firm (developer) or another household in resale transactions, and by transaction size. To summarize, bunching in resale transactions is greater for resale than for developer driven transactions, consistent with Kleven, Kreiner and Saez (2016) who argue that under-reporting of wages is lower amongst large firms because of the greater risk of a disgruntled employee revealing the under-reporting collusion—transactions involving households on both sides are less subject to such concerns. We also find that smaller size transactions show a more substantial jump in the under-reporting rate when $r = c$, consistent with the idea that households with larger value transactions may have a higher belief detection ρ compared to those with smaller transaction size.

4.2 Financial Constraints, Mortgages, and Under-reporting

One particularly important non-penalty driven incentive to truthfully report in our extended model is the desire to alleviate financial constraints when borrowing to fund a purchase transaction. Mortgage lending policies in India (and indeed in many other jurisdictions) often take the official reported/registered 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. Thus, the desire to unlock greater mortgage financing can increase incentives to truthfully report. The resulting prediction is that 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. In support of this idea, Appendix Figure A27 presents the relationship between the percentage of mortgages in each $(r-c)/c$ bin in Figure 5 and the under-reporting rate within each of those bins. The figure shows a strong negative relationship between the incidence of mortgage-based transactions and the under-reporting rate, with a slope of $-0.018 (s.e. = 0.005)$ and an R^2 of 22%.

To pursue this more formally, we use the extended model framework in Section 2.1.1. The optimal reporting behavior for individual taxpayers governed by equation (6). Although optimal reporting depends on the standard parameters, with external financial constraints, the loan-to-value ratio $(1-\gamma)$ and the tightness of the financial

constraint μ become important drivers of this behavior. Importantly, both $(1-\gamma)$ and μ drive reported values closer to the true market values of the transaction, in line with our empirical presentation in Appendix Figure A27.

To further test the implications of this simple application of our extended 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 data, we undertake this matching exercise, and we are able to match roughly 31,000 reported transactions to a mortgage.[33] We assume that any transaction that is unmatched to a mortgage transaction has a loan-to-value (LTV) ratio of zero, since this biases us against finding any differences between zero LTV and positive LTV transaction samples (this mechanically makes the zero LTV sample more similar to the positive LTV sample, because some of the zero LTV sample likely have a positive but unmeasured LTV).

Given our constrained data environment, any test can only be conducted using a very small sample since we are only able to match 8,913 of the 31,000 mortgage-matched transactions to a Propequity value p. To more fully utilize the matched data, we take a different approach, plotting how bunching behavior varies with loan-to-value ratios for the full set of 31,000 matched mortgage transactions. Figure 7a 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 (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. As we do not utilize exogenous variation in credit constraints here, we note that this result could also be driven by a negative correlation between preferences for tax evasion and agents' credit constraints.

In Figure 7b, 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 co-operative banks demonstrate the greatest extent of bunching, followed by banks which we were unable to perfectly classify, then 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

[33] We describe this matching process in detail in Appendix Section II, but note here that there are 78,414 mortgages which are not matched to transactions, which means there is likely measurement error and potentially selection in the mortgage-matched sample.

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 along the dimension of property tax under-reporting. For example, Mishra, Prabhala and Rajan (2022) show that public-sector Indian banks appear to have laxer credit screening standards and slower technology adoption than private sector Indian banks, resulting in higher non-performing borrowers for the public-sector banks, and link this finding to differences in organization culture.

Finally, Figure 7c 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.

5 Structural Estimation

Our goal in this section is to provide the first estimate of a government's revealed preference for tax accuracy versus tax revenue (i.e. the parameter ψ in equation 9). We begin by estimating the parameters that govern households' optimal reporting and evasion behavior. With those parameters in hand, we can then estimate the ψ that gives us the closest match to the observed enforcement policy (guidance values in our context).

5.1 Parameters Governing Individual Reporting Behavior r^*

Five parameters govern optimal reporting/evasion behavior in our structural model of tax reporting. First, $\pi_{r \geq c}$ is the household's perceived audit probability when

they choose to report more than the guidance value $r > c$.[34] A higher audit probability will, all things equal, generate a higher reporting amount. The level of the audit probability will also determine the discrete change in the audit probability that the household faces when it reports r just below versus just above c. In particular, if $\pi_{r \geq c} < 1$ then households have a strong incentive to bunch at c, or to the right of c ("diffuse-bunching"), to avoid the sure audit coming from reporting $r < c$ and associated transaction cost of being audited.

The second parameter ρ governs the households' beliefs about how much evasion is detected, conditional on being audited. Higher ρ values also lead to greater reported values, as conditional on being audited the expected penalty is higher buyers with larger ρ.

The parameter t determines the households' transaction cost of being audited ($\frac{t}{\tau}$)) a greater t value will lead to greater reporting (and in particular an avoidance of reporting $r < c$ as this means the household surely pays the transaction cost of audit).

Finally, ϵ_i determines where household i believes the government decides to audit all transactions (i.e. $\pi_{r<c} = 1$) versus only auditing a subset of transactions ($\pi_{r \geq c} < 1$) – positive values of ϵ_i mean households will bunch to the right of the guidance value c to avoid what they believe is a 100% chance of being audited and is included to generate the observed diffused bunching in the data.

We draw $\rho \sim \chi^2(K)$, where we search for K over a grid spanning $[0, 4]$. Given that $\pi_{r \geq c} \in (0.25, 1]$, we use this as the support to structurally estimate the value for $\pi_{r \geq c}$.[35] $\epsilon_i \sim \chi^2(E)$, where we search a grid from $[0, 2]$ for E. This value for ϵ_i is expressed in terms of c, i.e., if $\epsilon_i = 0.5$, then the individual believes the jump from $\pi_{r \geq c}$ to $\pi_{r<c}$ is at $1.5 \times c$. t is drawn from a $\chi^2(T)$ searched for over the grid $[0, 1]$, as it is expressed as a rate relative to c (i.e. a $t = .2$ implies a transaction cost of audit equal to $.2c$).

We observe data on reported (r_i), guidance (c_i) and (estimated) market values (m_i). We have a transaction tax $\tau = .05$ and a penalty multiplier from our context of 4. For each round of the optimization we choose one value for each of the five parameters discussed above. We then randomly draw values from the distributions according to these parameters, which in turn determine the optimal reporting amount r_i^* for each transaction based on our model. We iterate over all possible values of the parameters to find the parameter set that minimizes the following loss function:

[34] Recall, we fix the household's audit probability at 1 if they choose to report $r < c$ as this automatically triggers an audit.

[35] The lower bound for $\pi_{r \geq c}$ is $\frac{1}{n}$, because if it were lower the household would always report c as it would be worth the penalty in expected value terms.

$$\sqrt{\left(\frac{1}{K}\sum_{k=1}^{K}\frac{(r_k - r_k^*)^2}{\bar{r_k}}\right) + \left(\frac{1}{K}\sum_{k=1}^{K}\frac{(v_k - v_k^*)^2}{\bar{v_k}}\right)}, \tag{10}$$

where k are the 2% $\frac{(r-c)}{c}$ bins from -1 to 1 (a total of 99 bins), r_k is the total observed reported counts and r_k^* the total model-implied reported counts within each bin, and v_k the total observed reported value, and v_k^* the model-implied reported values within each bin, for a given combination of the four parameters. Since we equally-weight two different sets of moments with count and dollar-value as units, we normalize each error by the average bin count $\bar{r_k}$, and average bin value $\bar{v_k}$, observed in the data.

We find that the optimal values for $\{\pi_{r\geq c}, K, T, E\}$ to draw $\{\pi_{r\geq c}, \rho, t, \epsilon_i\}$ to be $\{0.289, 1.33, 0.44, 0.88\}$. Our estimate of the perceived probability of audit is $\pi_{r\geq c}$ is 28.9%. The fact that this is substantially lower than one results because we observe fewer transactions to the immediate left of c relative to the right in the data; because the probability of audit, and resulting payment of the transaction cost t changes discretely at c, agents have an incentive to report just at or above c versus below. $\rho = 1.33$ implies a generally low belief of detection probabilities conditional on audit; with $\rho = 1.33$ the bunchers in our analysis, who under-report by approximately 25% (see Figure A2) believe that there is a 16% detection probability conditional on audit.[36]. This parameter results primarily from the substantial amount of under-reporting that we observe in the data both at and away from the cut-off.

We estimate E, the mean of the χ^2 distribution as .88. Note that this implies approximately half of buyers believe the probability of audit is 1 even if they report 35% above c; however, given the low estimated values of ρ, these households essentially believe that while they are sure to be audited, the audit is unlikely to discover the full amount of under-reporting. This high-level of E matches the data pattern of substantial diffuse bunching around the guidance value c.

Figure 8a presents the best model fit. The blue dots are the observed counts from the raw data, the red triangles are the estimates of counts within each 2% bin implied by the model, and the purple left-triangles plot the market value distribution from the data. Figures 8b and 8c present the aggregate tax base values (max of r and c for each transaction aggregated up) and the under-reporting rate. Overall the model fit to the data matches the main data patterns we wish to include in our later estimation of the government's preference for tax accuracy: 1) a sizeable fraction of households report $r < c$ and are consistent with the "over-payers" we observe in the data 2) there is substantial bunching of reporting behavior at the guidance value c 3) under-reporting is

[36] (1-.75)$^{1.33}$ = .16.

highest at the bunch point and declines as we observe households that report more than c. Given these results, we expect the model to perform well in producing reporting behavior under the different possible guidance values we must consider to estimate the government's revealed preference for tax accuracy.

5.2 Estimating the Government's Preference for Tax Accuracy

Comparative Statics: To help develop intuition, Figure 4 presents some comparative statics for optimal government policy in this framework. We construct these figures by using the distributions of market values in our dataset, divided by the guidance values (we normalize the guidance value to 1). We set the parameters $\{\pi_{r\geq c}, K, T, E\}$ to draw $\{\pi_{r\geq c}, \rho, t, \epsilon\}$ to be $\{0.289, 1.33, 0.44, 0.88\}$, as were estimated in our structural model of reporting above. Once we calculate reporting behavior for each transaction we can sum them according to the government's objective function (equation 8) to get the government's welfare.

We first set $\zeta(.)$ to the baseline case of $\psi \mid \max(r^*(m,c),c) - m \mid$, capturing the symmetric concern of the tax authority for over- and under-payment. Figure 4a shows the optimal guidance value c^* for various levels of the inaccuracy weight parameter ψ in this case. When ψ increases, authorities lower c^*. However, the c^* converges to a constant with increases in ψ, suggesting that under a symmetric $\zeta(.)$, there is little room to lower c^* without hurting the revenue maximization objective.

Figures 4b show the optimal values when the penalty multiplier n set by the government varies, thus affecting the optimal individual reporting behavior. Increases in n for a given ψ generate lower levels of c^*, since individuals are more likely to be truthful in their reporting behavior when penalties increase.

The next two panels, Figures 4c and 4d present the case where external financial constraints affect optimal tax policy. We find here that decreases in γ that are influenced by loan-to-value ratios or lender-demanded truthful reporting policies can lower the optimal guidance value set by the tax authorities. Put differently, such external financial constraints force greater truth-telling in the market, thus enabling authorities to focus on reducing the burden on over-payers. Policies that vary γ are less useful unless the strength of the down-payment constraint is binding on potential buyers. μ, the parameter that determines the tightness of the financial constraint, can also affect optimal taxation, both through its relationship with the transaction tax rate τ, and in terms of how truthful the individual is likely to be. Figure 4d suggests that this constraint is only meaningful when $\mu < \tau$, i.e., any value of $\mu > \tau$ will already induce truthful reporting, thus inducing a corner solution of zero for the guidance value. For values of $\mu < \tau$, the figure suggests that smaller values of μ, i.e., slack in the financial

constraint, is only likely to increase the guidance value required for balancing the twin-objectives of maximizing revenue and maintaining fairness. Another way to interpret this result in the context of the Indian setting is that as the economy moves towards more financial formality in mortgage/borrowing markets, tax authorities will find it easier to satisfy their desire for tax accuracy, given the complementary role of financial markets in inducing truthful reporting.

Figures 4e and 4f relax the symmetric $\zeta(.)$ assumption. We assume in turn that the tax authorities care only about under-payers, i.e., set $\mid \max(r^*(m,c),c) - m \mid = 0$ when $m < c$ in panel (e); and that they only care about over-payers, i.e., set $\mid \max(r^*(m,c),c) - m \mid = 0$ if $m \geq c$ in panel (f). We vary ψ, the weight on the inaccuracy component of the government's objective function to understand its effects on c^*. When the authorities only care about under-payers, they are no longer constrained and can set c^* very high, simultaneously fulfilling their objectives of 1) maximizing revenues and 2) driving out under-payers. However, when they care only about over-payers, c^* falls rapidly to values much lower than under the symmetric $\zeta(.)$ case.

Structurally estimating the government's weight on inaccuracy ψ: We structurally estimate ψ by fixing the inaccuracy penalty $\zeta(.)$ function to only care about over-payers and not about the under-payers as in equation (11) below. We motivate this assumption with the observation that there are generally more complaints from households when they face guidance values c greater than their m values, and so it seems reasonable that a government would be very sensitive to this particular form of accuracy.

$$\zeta(r^*(m,c)) := \begin{cases} 0 & \text{if } m \geq c \\ \psi \mid \max(r^*(m,c),c) - m \mid & \text{if } m < c \end{cases} \tag{11}$$

We further assume that the government's weighting ψ does not differ across regions, i.e., it is the same for the entire whole city of Mumbai. We set as our target moments the observed guidance values c_{st} for all sub-zones s in all years t in our data. The algorithm proceeds heuristically as follows. We first pick a ψ value. Given this ψ we then proceed to estimate c^* for that given ψ that maximizes the government's objective function (which includes both revenue maximization and tax accuracy maximization specified in equation 9). We then choose the optimal ψ that produces a c^* that is closest to the observed guidance values c in the data according to this loss-function:

$$\min_{\psi} \sqrt{\sum (C^*_{st}(\psi) - C_{st})^2 \omega_{st}} \tag{12}$$

ω_{st} represents weights for each subzone s and time t. For simplicity, we place

greater weight on sub-zone $\times$ year with more transactions than those with fewer ones, a simple form of inverse-variance weighting.

We assume that the extensive margin elasticity to c is zero—in support of this assumption, Appendix Section G presents evidence that in the case of Mumbai, revisions to c do not result in any meaningful changes to the total volume of transactions. We also assume when structurally estimating ψ that while the government does not observe the true values m for every transaction, they do observe the m distribution for the newly-launched projects in Mumbai, and the reported value distribution r for other transactions. To operationalize this assumption, we use a mixture of these two distributions, m for the transactions with an exact match in the data, and r for those with non-exact matches in the data.

We structurally estimate $\psi = ₹1.05$; this implies the government is willing to give up ₹1.05 per ₹1 of over-payment. We interpret this as a substantial magnitude; a government that only focuses on maximizing tax revenue would have a ψ of zero. Intuitively, guidance values that aimed to just maximize tax revenue would be substantially higher, but the constraint of having a preference for not over-taxing some agents too much leads the government to essentially "under-enforce" tax policy.

Optimal vs. Observed Guidance Values How do optimal guidance values, calculated based on our structural model of reporting behavior, compare to observed guidance values in Mumbai? To address this we calculate optimal guidance values for each year based on our structural model of reporting, but estimate the parameters of the structural model of reporting using data from all other years. For instance, for the year 2016, for each sub-zone in the data, we structurally estimate ψ and an associated c^* using all data other than observations from 2016. We then compare the optimal c^* we calculate for 2016 to the observed guidance values c in 2016.

Figure 9a plots a histogram of the difference between the model-implied c^* and the actual c in the data across all years in our data. The modal observation is centered around 1.3, meaning that our estimates of c^* suggest that policymakers can raise guidance values for some sub-zones given their objectives as discussed above.

We then use the distance between this model-implied c^* and observed c for each sub-zone to forecast revisions to c over the next year (estimated 2016 differences between model-implied c^* and observed c to forecast 2017 revisions to c in the example above). This checks whether regions where the optimal guidance value is lower or higher than the prevailing guidance value are likely to witness an upward or downward revision in the future. Figure 9b presents the average reduction in guidance value predicted by the model. The top panel presents the in-sample prediction, and the

bottom-panel presents the out-of-sample prediction. In both cases, the figure shows that both upward and downward revisions are well-predicted by deviations between model-implied c^* and actual c. While our directional forecasts are broadly in line with the data, the magnitude of observed revisions are small relative to our model—suggesting that there are other factors at play to explain quantitative responses than in our simple structural model.[37] Nevertheless, this serves as useful out-of-sample validation of our model.

6 Conclusion

We develop a new framework for optimal taxation and enforcement when policymakers care about both welfare maximization and tax accuracy. This framework accommodates the widespread sentiment that there is value in reducing tax noncompliance separate from the revenues raised, and correspondingly, there is a cost to overcollecting revenues beyond what is legally owed. Our model formalizes these sentiments by adding a priority for tax accuracy in the policymaker's objective function.

We apply this framework to the empirical setting of property transaction taxes in Mumbai with self-reported transaction values. We recover the degree of misreporting, and the elasticity of misreporting and transaction volume to enforcement, based on the degree of bunching in reported valuations around government-assessed guidance values which serve as a floor on the tax base. Existing policy suggests a strong preference against tax inaccuracy—and especially against tax overcollection—on the part of policymakers. We also find evidence of a strong correlation between the degree of bunching of reported property values at guidance values and features of mortgage contracts–such as the LTV ratio on the loan, and the identity and financial health of the bank issuing the loan–on these properties. 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.".

[37] Appendix Table 3 confirms these visual results in a regression.

Allingham, Michael G. and Agnar Sandmo. 1972. "Income tax evasion: A theoretical analysis." *Journal of Public Economics* 1(3):323–338.

Alvarez, Fernando and David Argente. 2022. "On the effects of the availability of means of payments: The case of uber." *Quarterly Journal of Economics* 137(3):1737–1789.

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. 2022. "Reference dependence in the housing market." *American Economic Review* 112(10):3398–3440.

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

Avenancio-León, Carlos F and Troup Howard. 2022. "The assessment gap: Racial inequalities in property taxation." *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.

Bergolo, Marcelo, Rodrigo Ceni, Guillermo Cruces, Matias Giaccobasso and Ricardo Perez-Truglia. 2023. "Tax Audits as Scarecrows: Evidence from a Large-Scale Field Experiment." *American Economic Journal: Economic Policy* 15(1):110–153.

Best, Michael Carlos and Henrik Jacobsen Kleven. 2018. "Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK." *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.

Casaburi, Lorenzo and Ugo Troiano. 2016. "Ghost-house busters: The electoral response to a large anti–tax evasion program." *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." *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." *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.

Das, Shaswati. 2020. "Why Arun Jaitley Gave University Basic Income a Miss in Union Budget 2017." *Livemint* .
URL: *https://www.livemint.com/Politics/ns8z7JTxYLQc5AdGmlvtiN/Why-Jaitleys-Union-budget-didnt-provide-for-universal-basi.html*

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

Feldstein, Martin. 1995. "The effect of marginal tax rates on taxable income: a panel study of the 1986 Tax Reform Act." *Journal of Political Economy* 103(3):551–572.

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.

Garmaise, Mark J and Tobias J Moskowitz. 2004. "Confronting information asymmetries: Evidence from real estate markets." *Review of Financial Studies* 17(2):405–437.

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." *Quarterly journal of economics* 116(4):1233–1260.

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

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

Harari, Mariaflavia and Maisy Wong. 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.

Hendren, Nathaniel, Ben Sprung-Keyser and Ellen Stuart. 2023. A Welfare Analysis of Tax Audits Across the Income Distribution. Technical report National Bureau of Economic Research.

Karmakar, Sudipto and Abhinav Narayanan. 2020. "Do households care about cash? Exploring the heterogeneous effects of India's demonetization." *Journal of Asian Economics* 69:101203.

Keen, Michael and Joel Slemrod. 2017. "Optimal tax administration." *Journal of Public Economics* 152:133–142.

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

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.

Mayle, Peter. 2000. *A year in Provence*. Penguin UK.

Mayshar, Joram. 1991. "Taxation with costly administration." *The Scandinavian Journal of Economics* pp. 75–88.

Mirrlees, James A. 1971. "An exploration in the theory of optimum income taxation." *The review of economic studies* 38(2):175–208.

Mirrlees, James A. 1986. "The theory of optimal taxation." *Handbook of mathematical economics* 3:1197–1249.

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?" *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.".

Okunogbe, Oyebola and Fabrizio Santoro. 2022. "The promise and limitations of information technology for tax mobilisation.".

Panchapagesan, V. and N. Karthik. 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* .

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.

Ramsey, Frank P. 1927. "A Contribution to the Theory of Taxation." *The economic journal* 37(145):47–61.

Regan, Tanner. 2023. "The Economic Consequences of Statistical Property Valuation: Evidence from Kampala, Uganda." Abstract.

Riots erupt in Nigerian cities as bank policy leads to scarcity of cash. 2023. *The Guardian* . **URL:** *https://www.theguardian.com/world/2023/feb/15/angry-protests-erupt-across-nigeria-against-scarcity-of-cash*

Rocha, Roberto, Thiago Scot and Javier Feinmann. 2023. "Property Taxes in Sao Paulo.".

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

Saez, Emmanuel, Joel Slemrod and Seth H Giertz. 2012. "The elasticity of taxable income with respect to marginal tax rates: A critical review." *Journal of economic literature* 50(1):3–50.

Sandmo, Agnar. 1981. "Income tax evasion, labour supply, and the equity—efficiency tradeoff." *Journal of public economics* 16(3):265–288.

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

Slemrod, Joel, Caroline Weber and Hui Shan. 2017. "The behavioral response to housing transfer taxes: Evidence from a notched change in DC policy." *Journal of Urban Economics* 100:137–153.

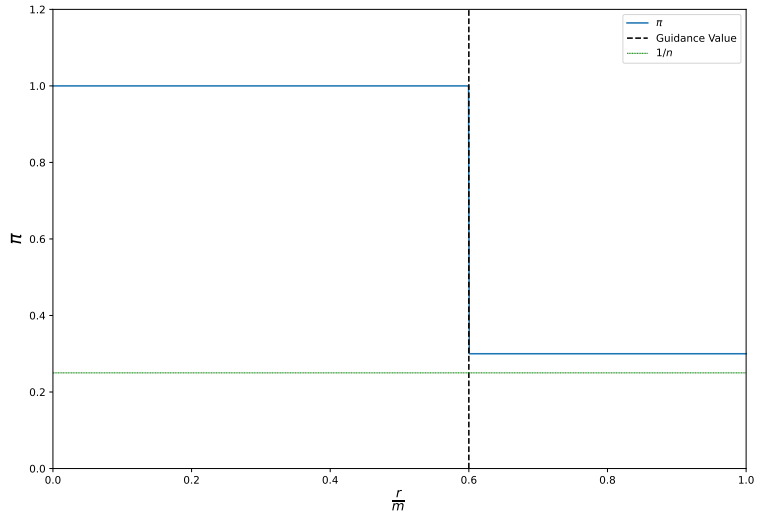
Sood, Aradhya. 2020. Land Market Frictions in Developing Countries: Evidence from Manufacturing Firms in India. Technical report Mimeo. https://aradhyasood. github. io/Sood_Land_Frictions_India1. pdf.

Srinivasan, T.N. 1973. "Tax evasion: A model." *Journal of Public Economics* 2(4):339–346.

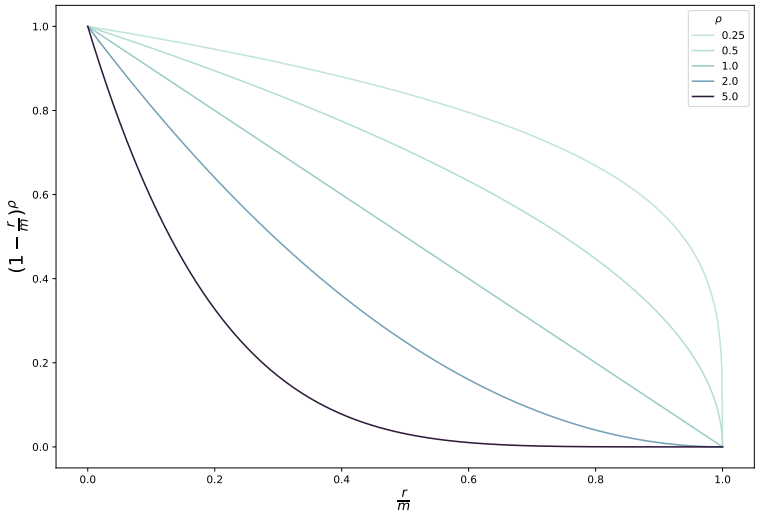
Tandon, N.K. 1987. "Under-reporting of Real Estate Transactions in India." Unpublished Manuscript.

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

Yitzhaki, Shlomo. 1979. "A note on optimal taxation and administrative costs." *The American Economic Review* 69(3):475–480.



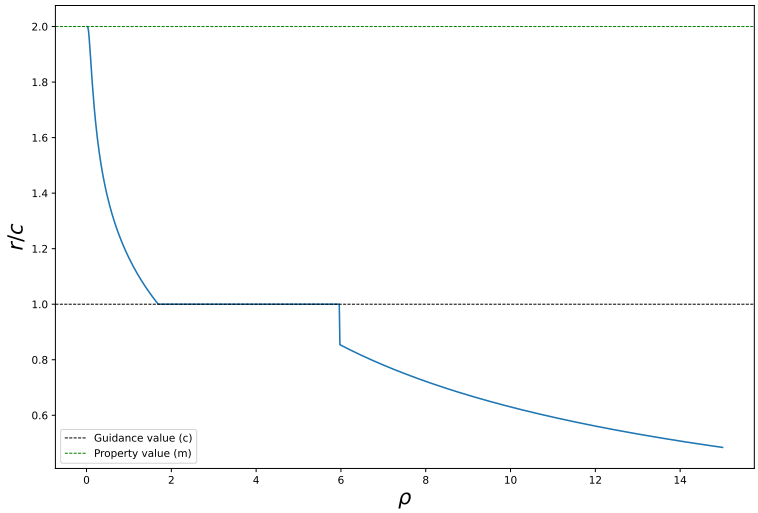
(a) Audit Probability π



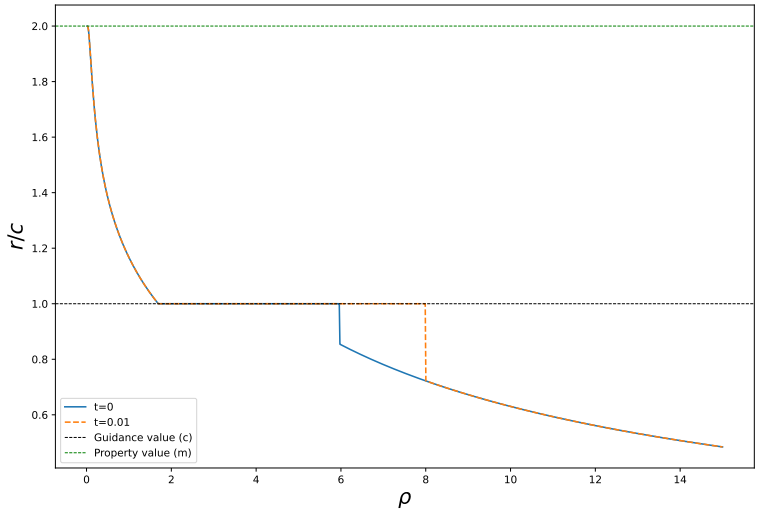
(b) Detection Belief ρ

Figure 1
Belief Parameters in the Model

Panel(a) presents the audit probability π, and Panel (b) presents the detection belief parameter ρ.



(a) Reported Values for different ρ



(b) Increasing t

Figure 2
The Role of ρ, π and t

Panel(a) presents reported values, relative to c, for various values of ρ. Panel (b) presents reported values, relative to c, for various values of ρ, and in the present of transaction costs, t.

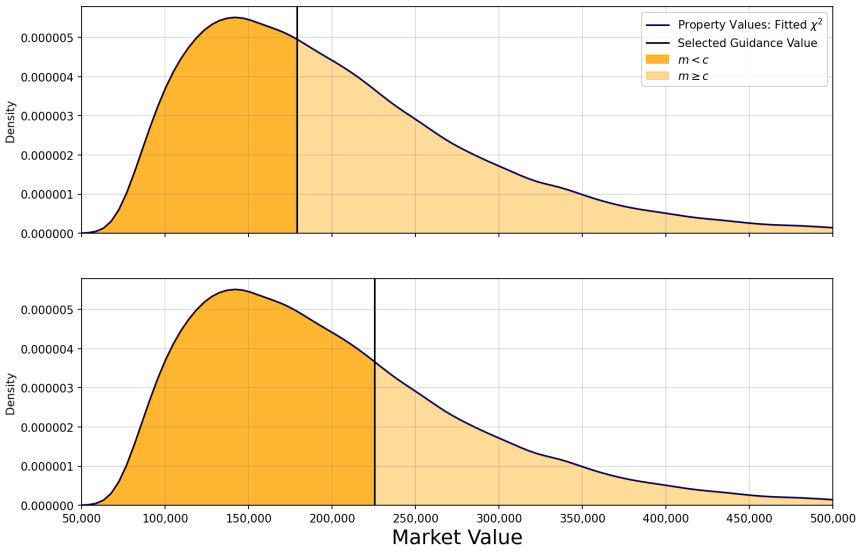
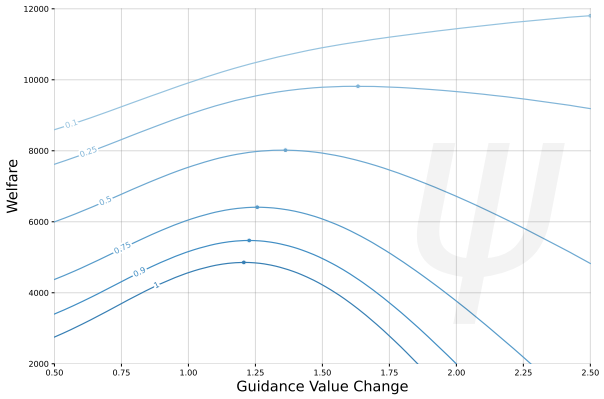
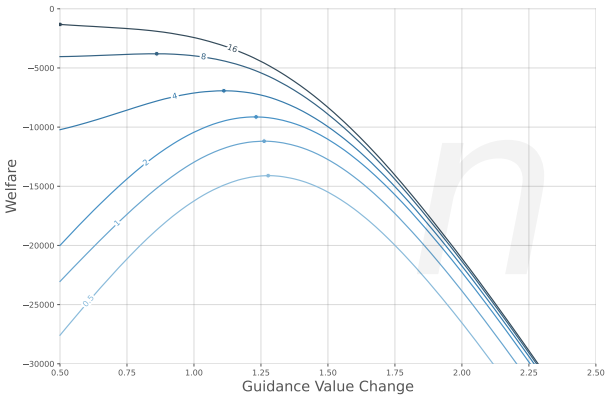


Figure 3
Optimal Government Policy: An Illustration

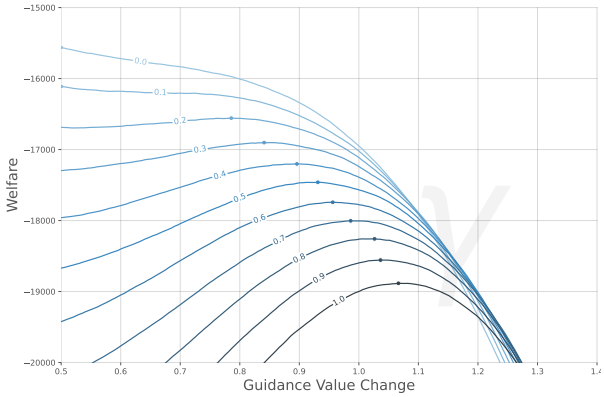
The two panels present distribution of market values in the asset market with two different guidance value cut-offs for c.



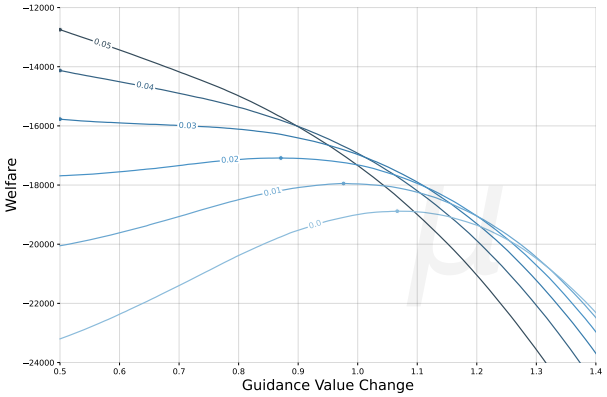
(a) Optimal guidance value for different ψ



(b) Optimal guidance value by n

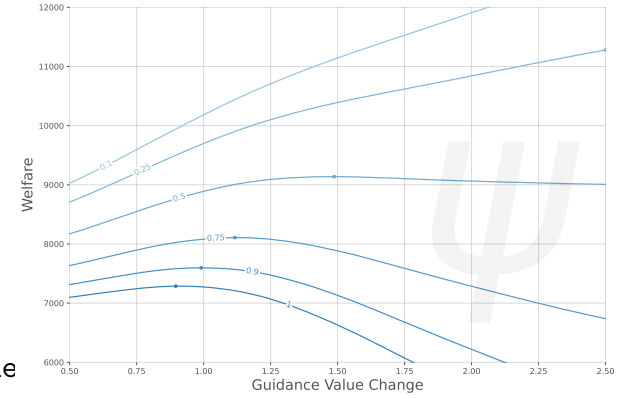


(c) Optimal guidance value for different γ



(d) Optimal guidance value by μ

1_DOC/1_FIGURES/Mode



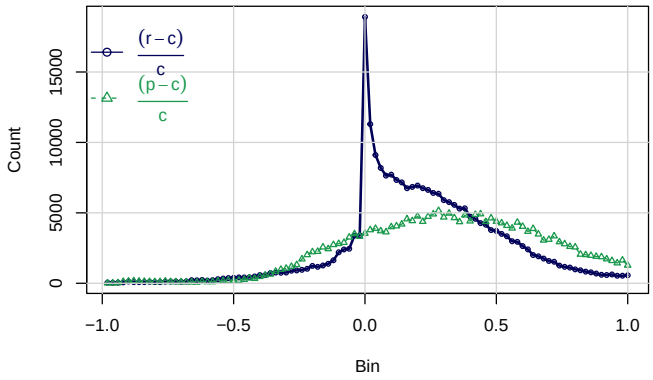
.WelfareUnc

(e) Optimal guidance value when only under-payers matter

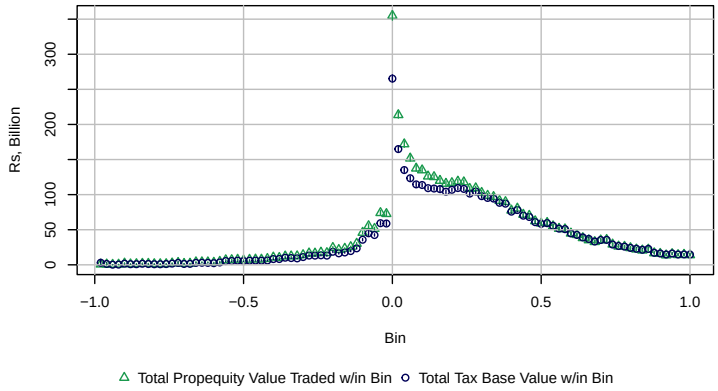
(f) Optimal guidance value when only over-payers matter

Figure 4
Comparative Statics: Optimal Government Policy

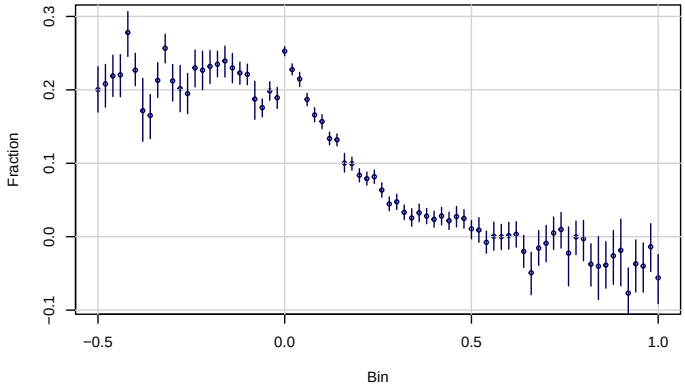
Panel(a) presents the comparative statics for optimal guidance value varying ψ, (b) for various values of the penalty parameter n, panel (c) for various values of γ, the LTV set by the regulator, panel (d) for various values of μ, the down-payment constraint, panel (e) when only under-payers matter (therefore finding the guidance value should be as high as possible, so as to avoid under-payers) and panel (f) when only over-payers matter.



(a) Bunching of Reported and Propequity Values Around Circle Values



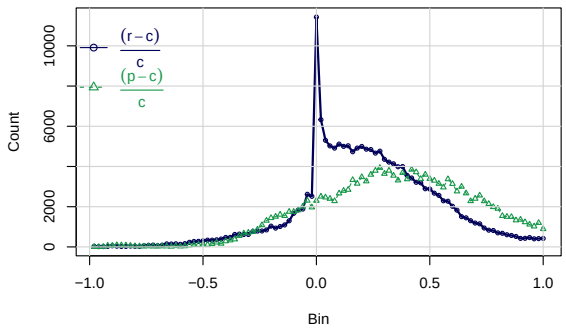
(b) Aggregate Taxbase and Propequity Values by Reporting Behavior Bins



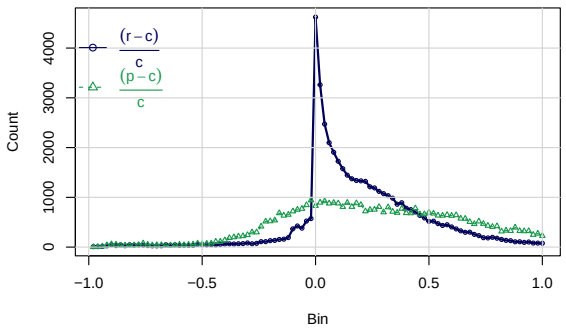
(c) Under-Reporting Rate by Reporting Behavior Bins

Figure 5
Main Results

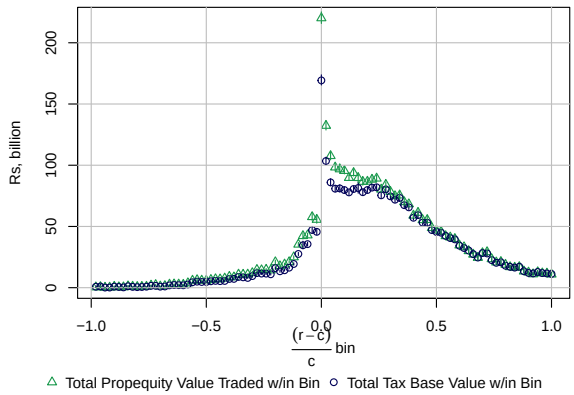
The blue line/circle shows the distribution of reported values across 2% reported value bins, where a reported value bin is measured as a deviation from the guidance value. The green line/triangle shows the distribution of our noisily measured estimate of the market price (the Propequity values) for the same underlying set of transactions reported in the blue line/circle. The blue circles in (c) present the under-reporting rate within 2% reported value bins. Vertical bars represent 95% bootstrapped confidence intervals.



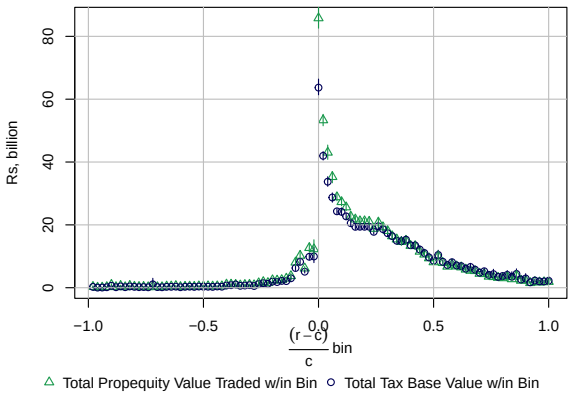
(a) Developer Sales



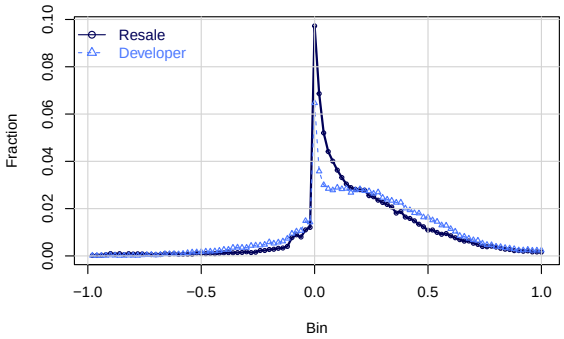
(b) Resale



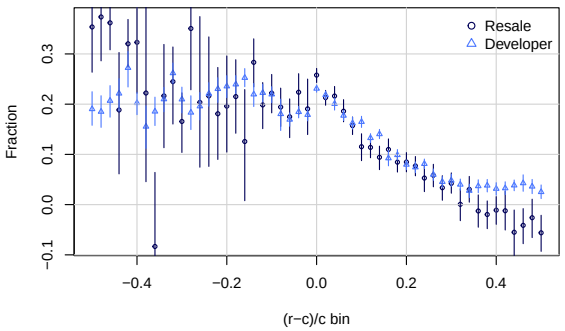
(c) Developer Sales



(d) Resale



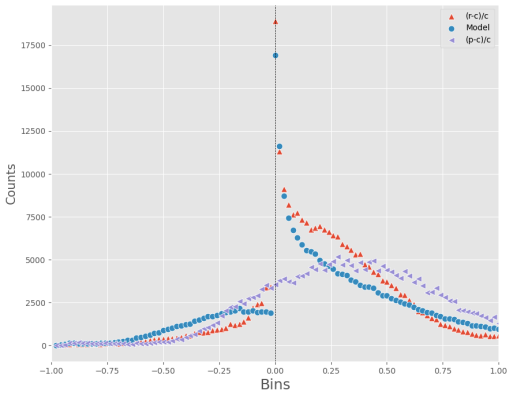
(e) Reported Counts (Density) Comparison



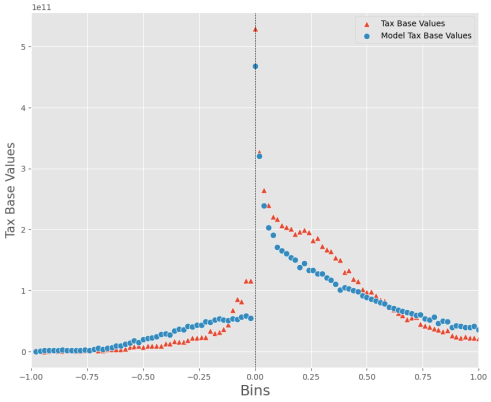
(f) Under-Reporting Rates

Figure 6
Developer vs. Resale Heterogeneity

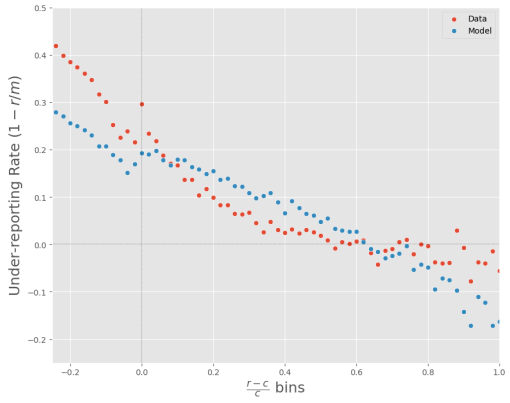
See Figures 5a, 5b and 5c for detailed descriptions.



(a) Count Distribution



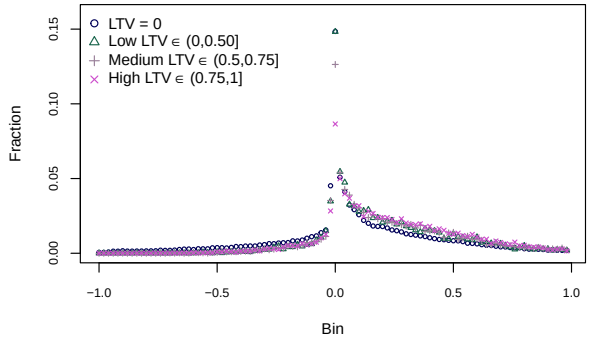
(b) Value Distribution



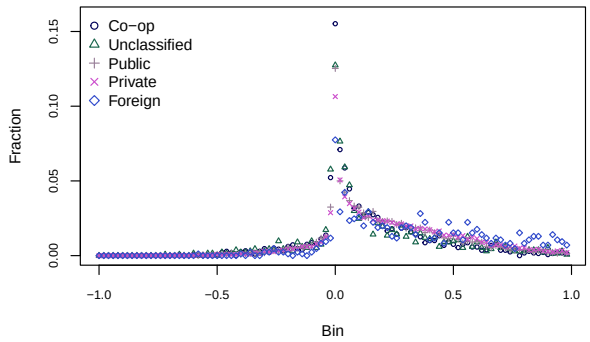
(c) Under-reporting Rate

Figure 8
Structural Estimation and Model Fit

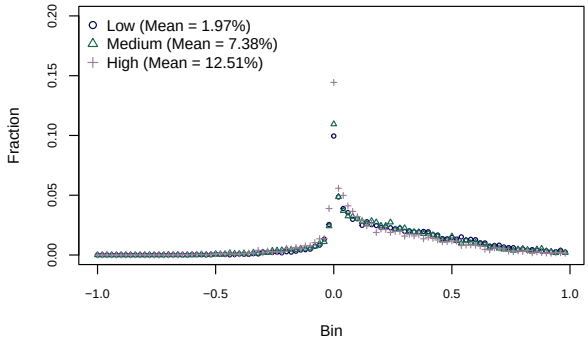
Panel(a) and (b) presents the model fit to the targeted moments, the distribution of transaction counts in 2% bins and the value distribution. Panel (c) presents the model fit to the untargeted under-reporting rate.



(a) Loan-to-Value Bunching Heterogeneity



(b) Bank Type Bunching Heterogeneity



(c) Non-Performing Loan Heterogeneity

Figure 7
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 from the IGR data until 2019. 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.

../model/Figures/OOS_Subzones_Histogram.pdf

(a) c^*/c: Histogram

../model/Figures/OOS_Subzones_increase.pdf

(b) Predicting Revisions in c by c^*/c

Figure 9
Predicting Revisions in c

Panel(a) presents the histogram of distance between c^* and c, while Panel (b) presents the average distance within each c^* distance bin (top sub-panel) and the revision in the bottom sub-panel.

	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

Table 1
Summary Statistics on Transactions

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. A primary transaction is one where the housing unit is sold by a real estate developer.

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	58	75	88	2.10	0.124
15–25	88	99	111	1.20	0.034
25–35	111	123	135	1.53	0.060
35–45	135	147	160	1.53	0.088
45–55	160	174	190	1.23	0.043
55–65	190	208	230	1.24	0.041
65–75	230	254	283	0.94	0.033
75–85	283	324	387	1.05	0.032
85–95	387	494	715	0.75	0.028

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
Percent increase in c and distance to c^*

	Pct. increase in c	Pct. increase in c	Pct. increase in c	Pct. increase in c
C^*/C	0.038***	0.060***	0.030***	0.053***
	(0.002)	(0.009)	(0.007)	(0.014)
$(C^*/C)^2$		-0.015*		-0.012
		(0.006)		(0.007)
Year Fixed Effects	No	No	Yes	Yes
Subzone Fixed Effects	No	No	Yes	Yes
Number of Years	9	9	9	9
Number of Subzones	197	197	197	197
Obs	862	862	862	862

Standard errors in parentheses. * $p < .05$, ** $p < .01$, *** $p < .001$

Optimal Tax Policy with Misreporting: Theory, and Evidence from Real Estate

Santosh Anagol Vimal Balasubramaniam
Benjamin B. Lockwood Tarun Ramadorai Antoine Uettwiller

March 20, 2025

A Tables & Figures

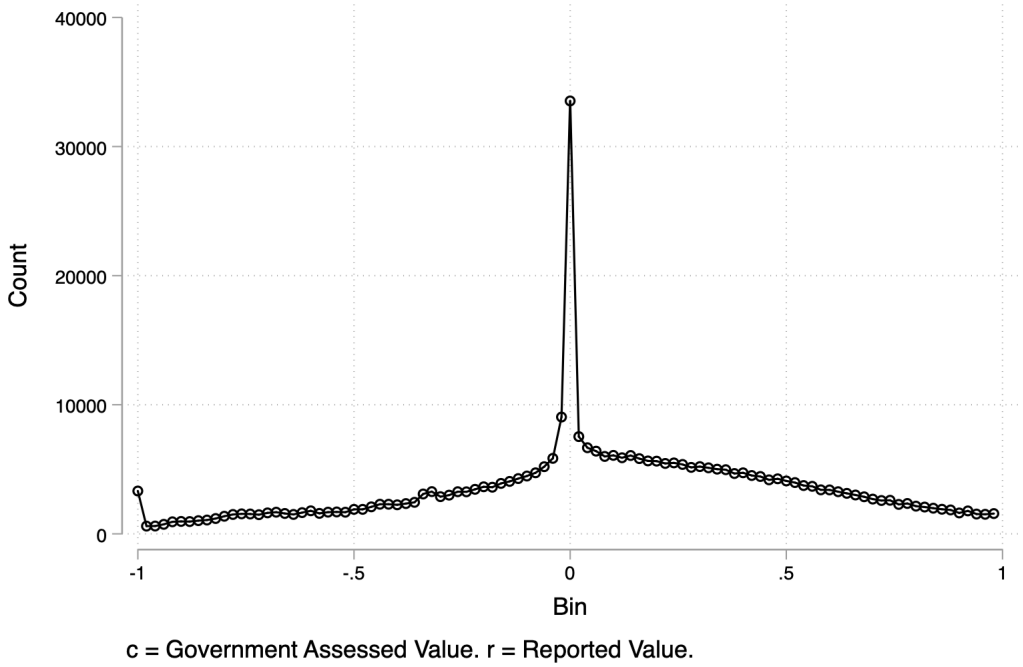
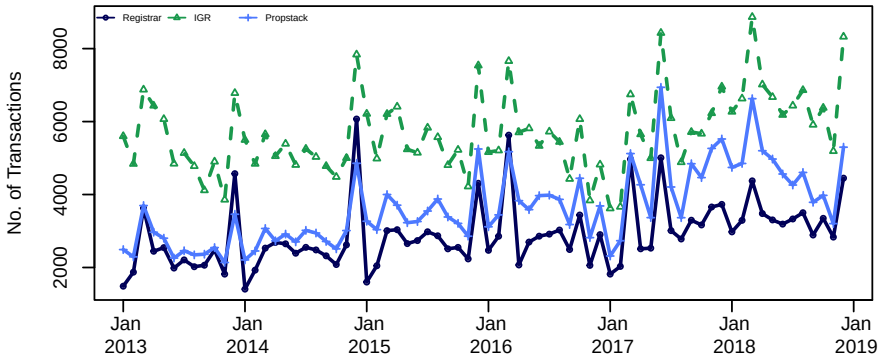
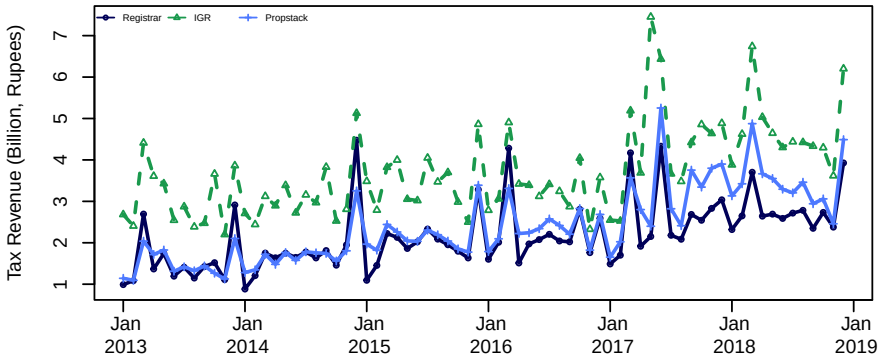


Figure A1
Bunching Estimates for Sao Paulo, Brazil

This figure replicates Rocha, Scot and Feinmann (2023) who show similar bunching patterns using the ITBI municipal transaction tax on properties in Sao Paulo, Brazil. The transaction tax rate is 3%, and it is charged on the higher of the buyer's reported value or a guidance guidance value. The line shows the distribution of reported values across 2% reported value bins, where a reported value bin is measured as a deviation from the guidance guidance value.



(a) Transactions



(b) Tax Revenue

Figure A2
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.

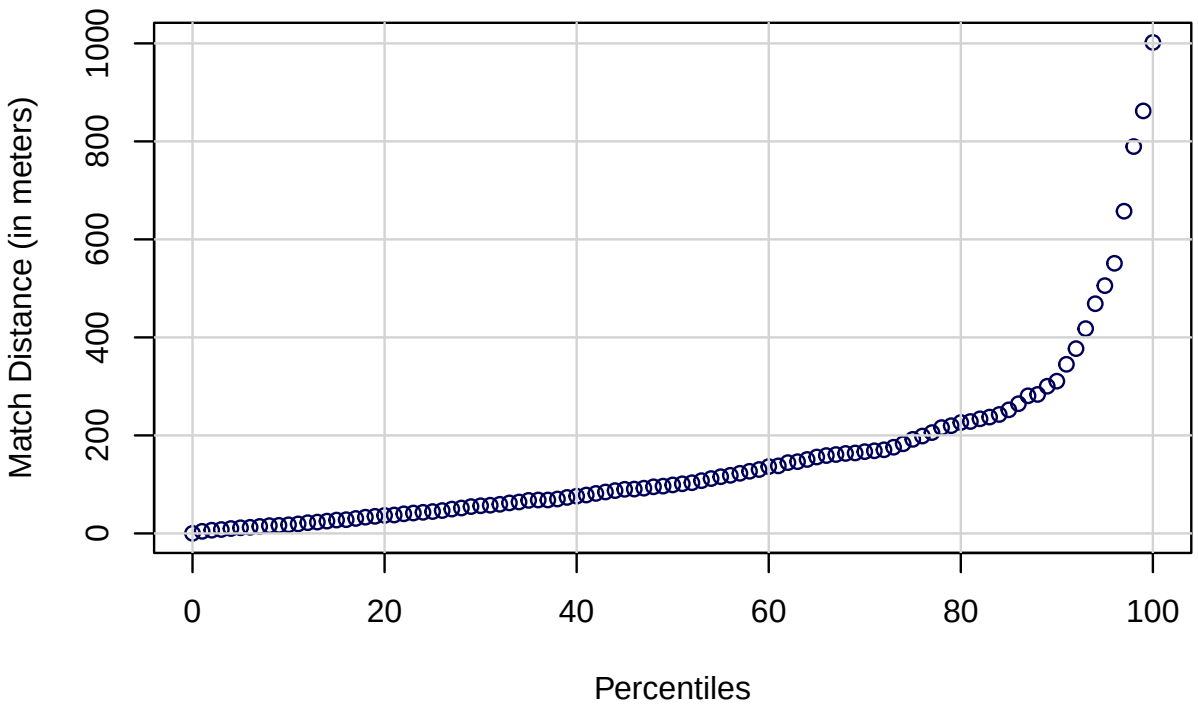


Figure A3
Propstack-Propequity Match Quality

This figure plots the empirical distribution of the match distance (in meters) between a transaction in our Propstack data (source of administrative data on reported and guidance values) and the match from our Propequity data (source of estimated price data). 80% of the transactions matched to Propequity price data are within 200 meters, and 95% of the transactions are matched to Propequity transactions within 500 meters.

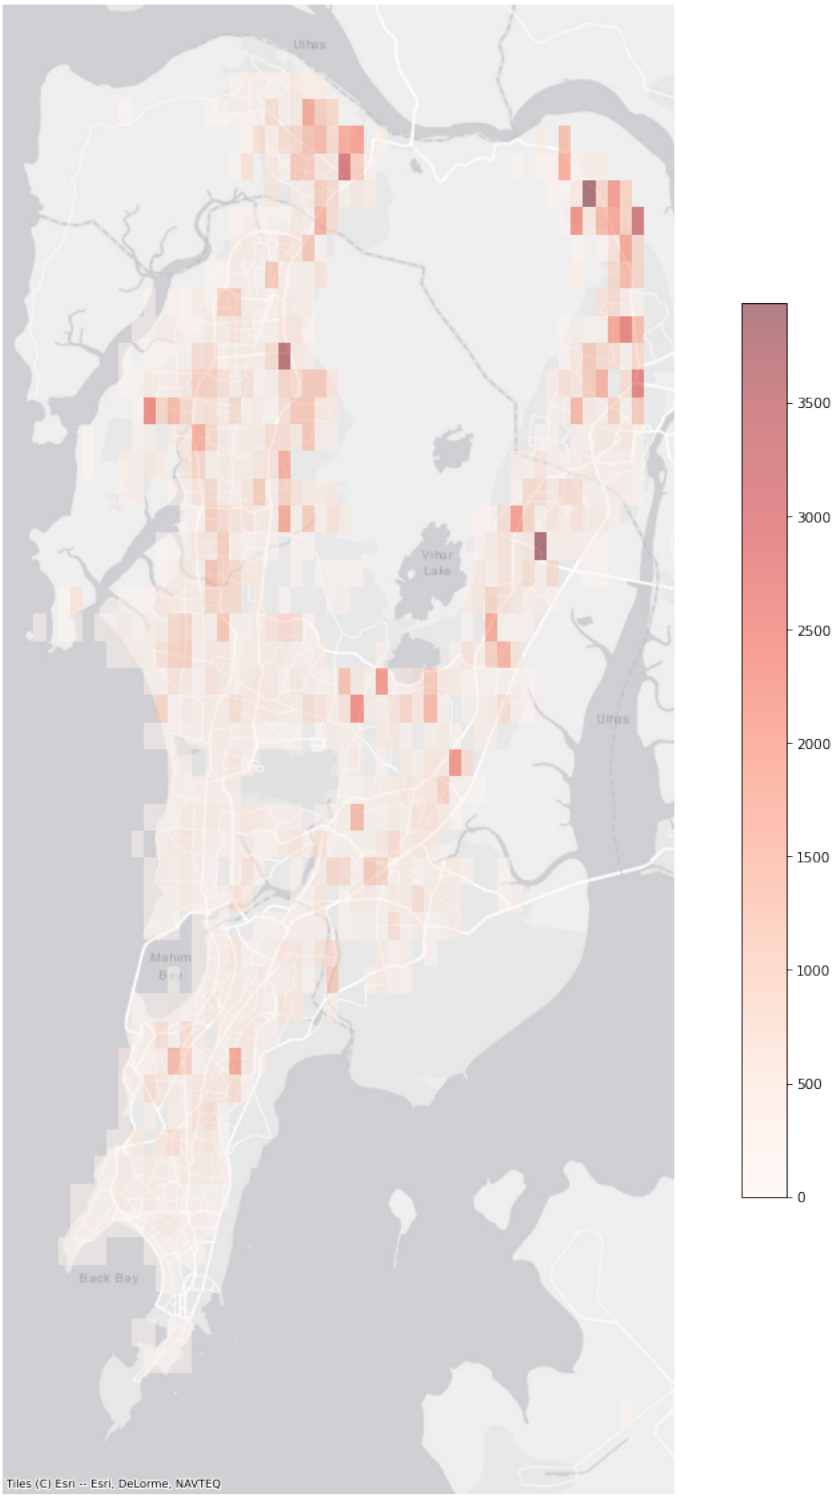


Figure A4
Heatmap of Transactions in our Final Sample

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

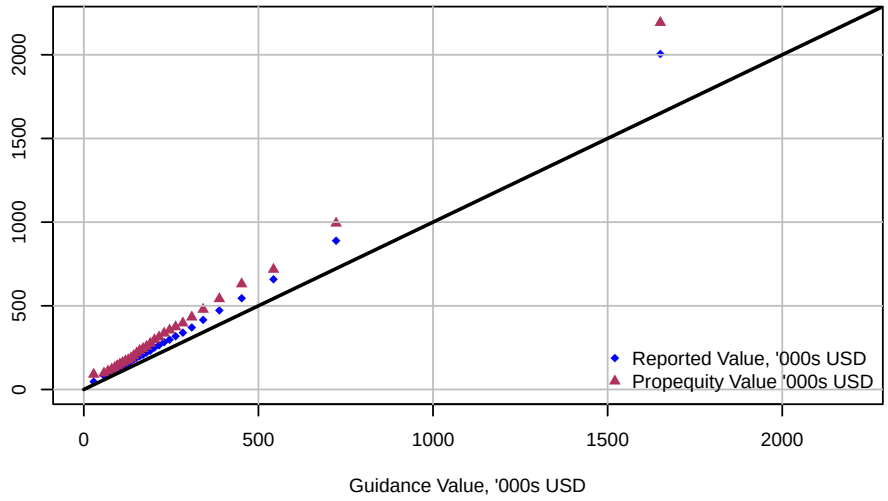


Figure A5
Correlation between Reported Value, Propequity (Estimated Market) Values, and Guidance Values

This figure presents a binned scatterplot of the average reported values (blue diamonds) and estimated market values from the Propequity data (maroon triangles) within guidance value bins. The black line is the 45-degree line.

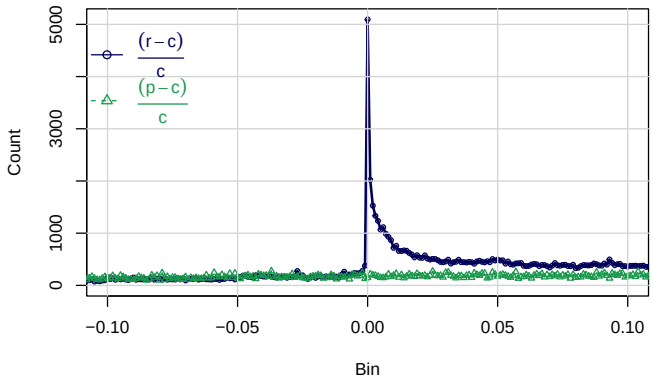
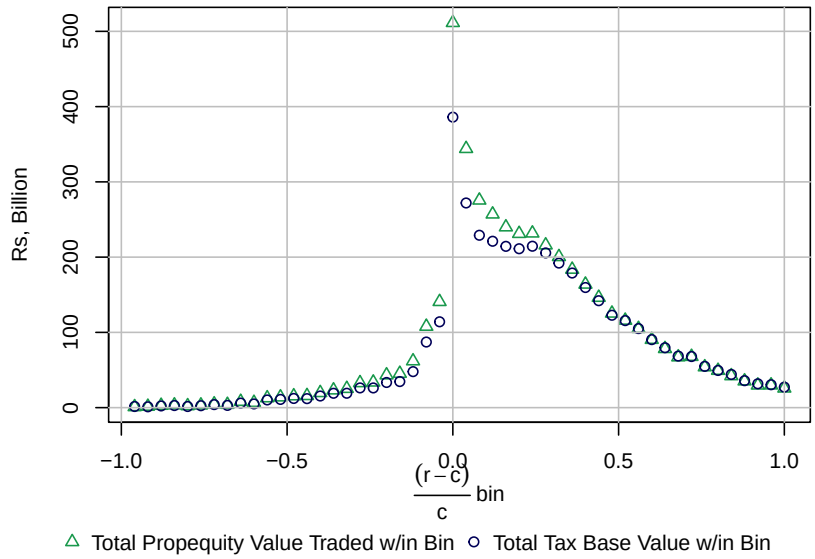
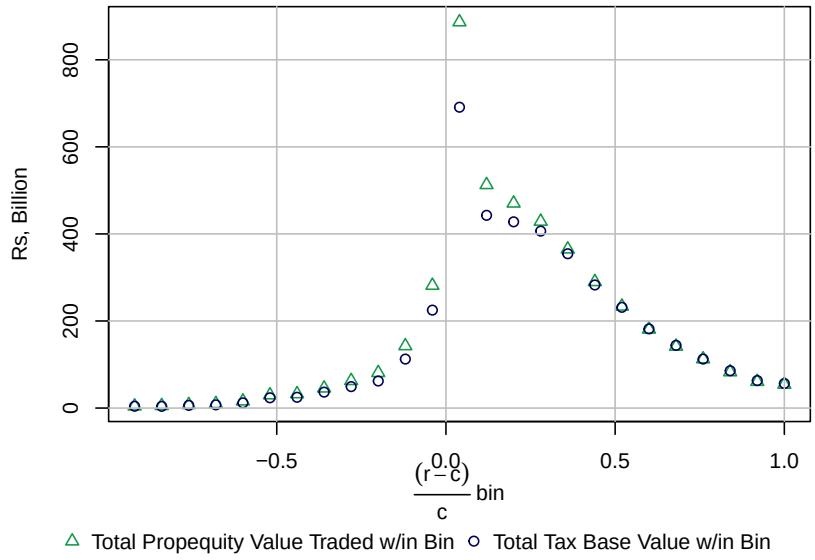


Figure A6
Bunching of Reported and Propequity Values Around Circle Values in 0.1% Bins

The blue line/circle shows the distribution of reported values across 0.1% reported value bins, where a reported value bin is measured as a deviation from the guidance value. The green line/triangle shows the distribution of our noisily measured estimate of the market price (the Propequity values) for the same underlying set of transactions reported in the blue line/circle.



(a) Bin width of 4%



(b) Bin width of 8%

Figure A7
Robustness: Aggregate Tax Base 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 2 for details.

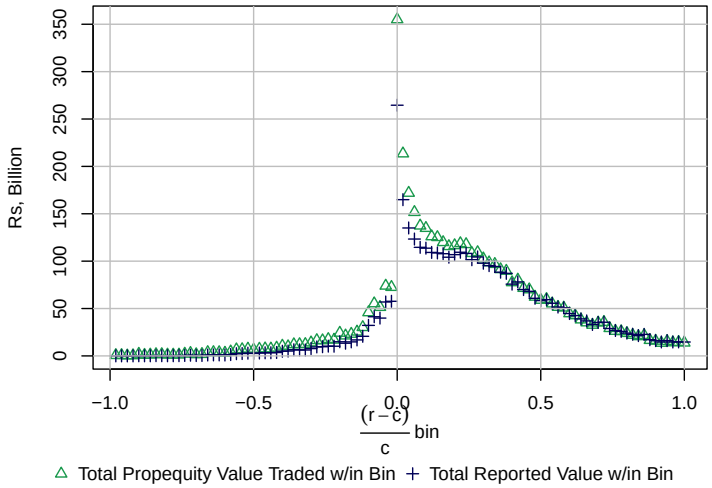


Figure A8
Aggregate Reported and Propequity Values by Reporting Behavior Bins

This figure reports the aggregated reported and Propequity (estimated market) values by reporting behavior bins. It differs from Figure 5 in that this figure aggregates reported values even if they are lower than a transaction's guidance value.

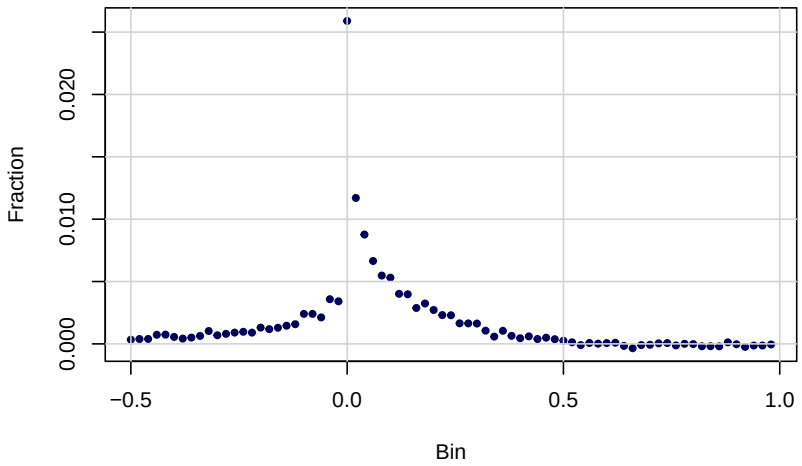


Figure A9
Truthful Reporting Hypothesis:
Measurement Error Contribution by Reporting Behavior Bin

Each point is equal to aggregate difference between our estimated market values and the tax-base value within a $\frac{r-c}{c}$ bin divided by the total aggregated tax-base value across the whole sample. Assuming that the tax-base values are the true market values, the figure then shows the pattern of measurement error required to rationalize our data under the truthful reporting hypothesis. The points sum to the aggregate measurement error amount as discussed in the decomposition in the text. The figure shows that measurement error would have to be strongly concentrated near bunching transactions to explain the relationship between reported and estimated market values in our data.

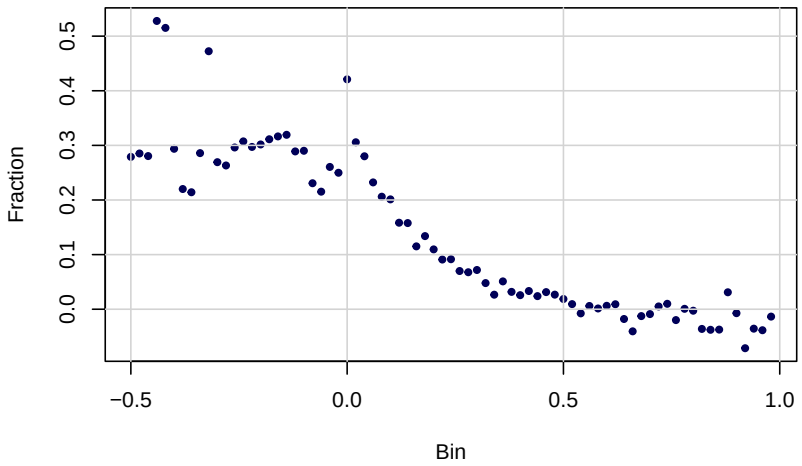


Figure A10
Estimated Measurement Error Assuming Truthful Reporting: Full Sample

The blue circles show our estimate of measurement error in our estimated market price data under the assumption of completely truthful reporting (i.e. that the tax base value for each transaction is actually the true market value). The estimated measurement error is the sum of the differences between our estimated market values and the tax base value within each bin, divided by the total tax base value in each bin. This version of the figure does not trim large estimated market values.

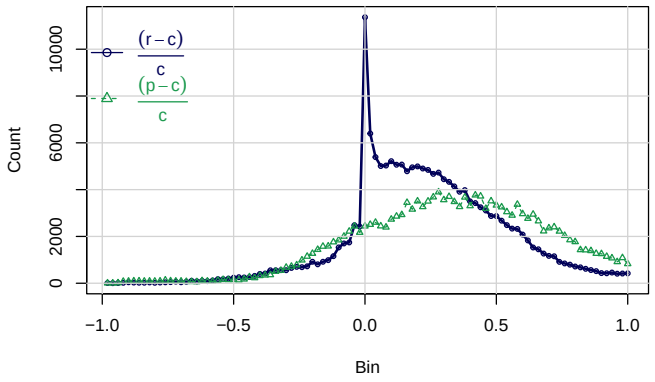


Figure A11
Exact Match: Bunching of Reported and Propequity Values Around Circle Values

This figure replicates our main findings using 60% of the all transactions in our sample where we have an exact project match for our estimate of the market price (the Propequity 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 guidance value. The green line shows the distribution our noisily measured estimate of the market price (the Propequity values) for the same underlying set of transactions reported in the blue line.

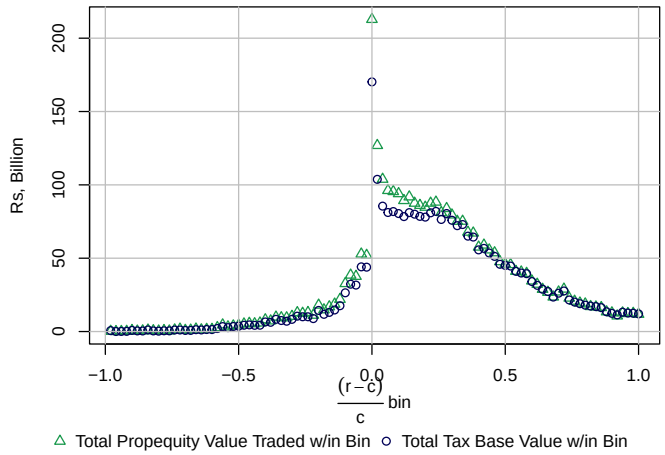


Figure A12
Exact Match: Aggregate Reported and Propequity Values by Reporting Behavior Bins

This figure replicates our main findings using 60% of the all transactions in our sample where we have an exact project match for our estimate of the market price (the Propequity values). 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 guidance value. The blue circles show the aggregate noisily measured estimate of the market value (the Propequity values) for the same underlying set of transactions reported in the green triangles.

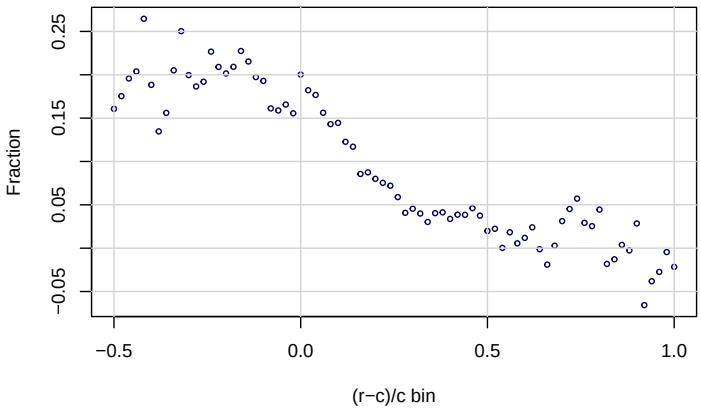
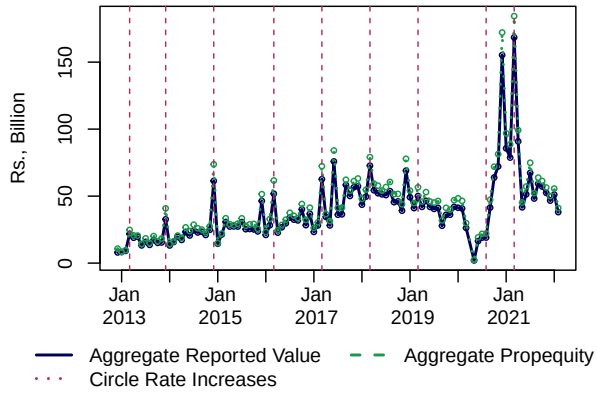
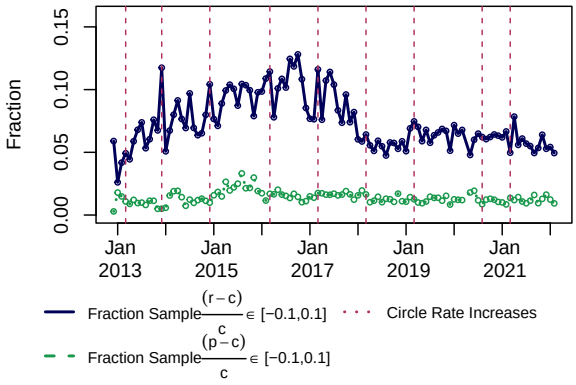


Figure A13
Exact Match: Under-Reporting Rate by Reporting Behavior Bins

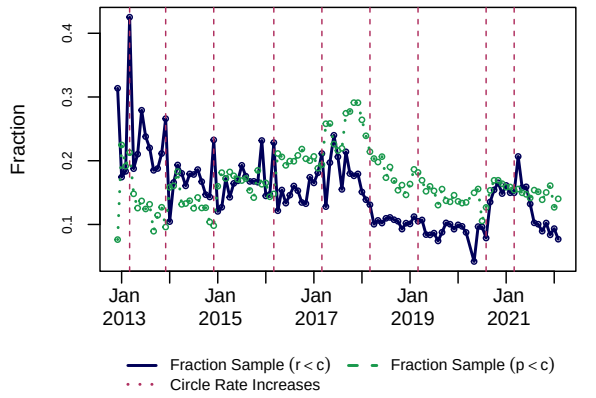
This figure replicates our main findings using 60% of the all transactions in our sample where we have an exact project match for our estimate of the market price (the Propequity values). The blue circles show the estimated under-reporting rate within 2% reported value bins, where a reported value bin is measured as a deviation from the guidance value.



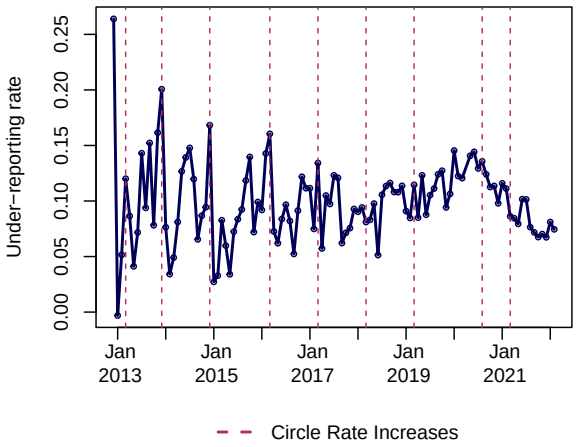
(a) Monthly Aggregates



(b) Monthly Bunching



(c) Reporting/Market Values $<$ Guidance Values



(d) Monthly Under-Reporting Rate

Figure A14
Heterogeneity Over Time

Red vertical dashed lines refer to scheduled circle rate (guidance value) increases. The circle rates were increased in 2013, 2014 and 2015; The government kept circle rates the same in 2016, 2017 and 2018, but the market likely still expected a possible increase in circle rates in those years. Circle rates were reduced in 2020 in response to the Covid-19 pandemic, but they increased again in 2021. The large increase in transactions in 2021 is likely due to a transaction tax rate reduction during that time period.

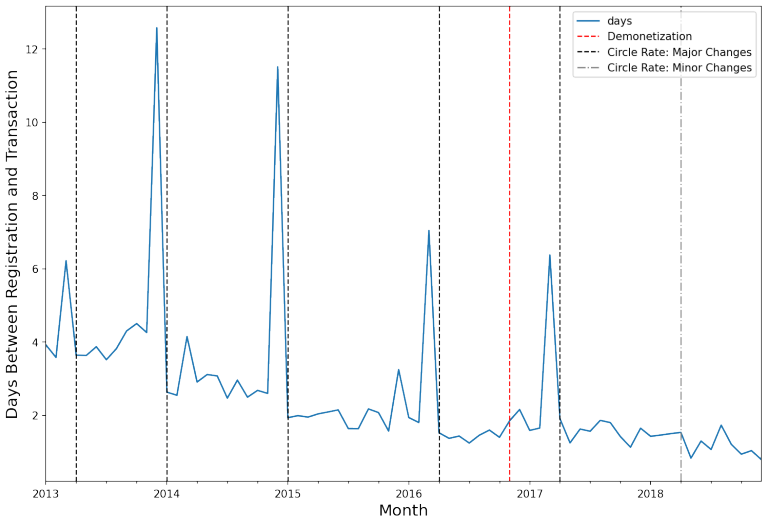
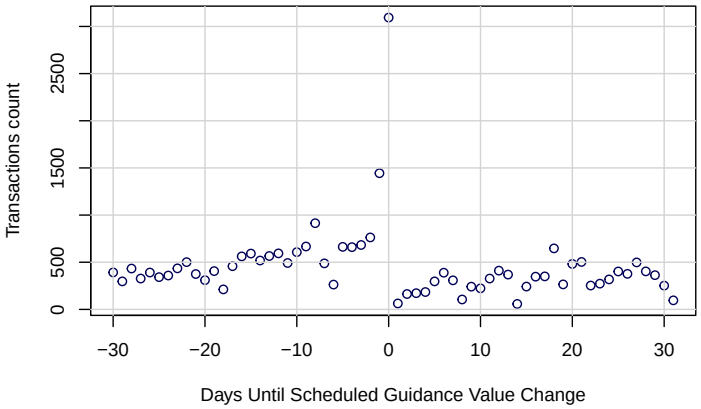
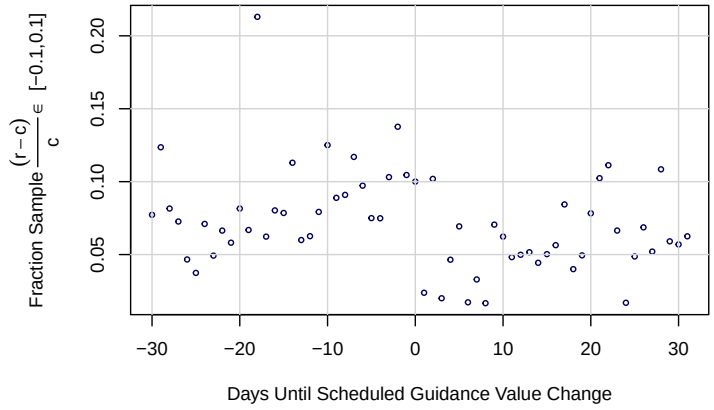


Figure A15
Agreement Date Backdating

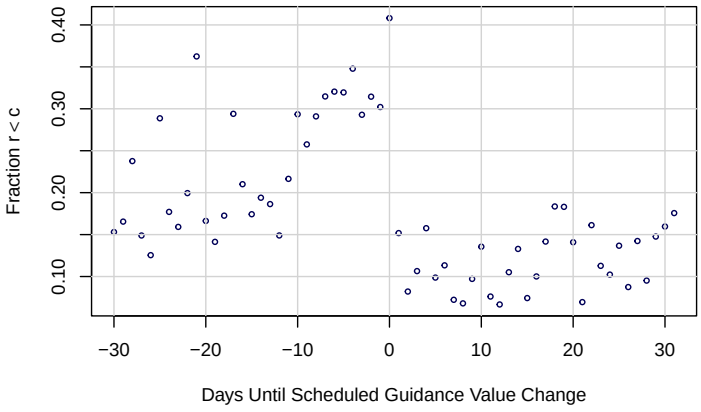
This figure plots the agreement month on the x-axis and the average number of days between the registration and transaction dates on the y-axis.



(a) Transaction Counts



(b) Fraction Bunching



(c) Fraction $r < c$

Figure A16
Reporting Behavior in Days Around Scheduled Guidance Value Increases

Panel (a) shows counts of transactions made per day within a 60 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.

				2BHK COST SHEET-
Saleable Area		1237	1242	12
Basic Rate		6099	6099	
Basic Cost		7544463	7574958	
Car Parking Charges (1 car park slot)		400000	400000	
Preferential Location Charges (Rate per sqft)		50	50	
Preferential Location Charges Cost		61850	62100	
Basic Unit Value		**8006313**	**8037058**	**8061**
				Additi
Corpus Fund Rs 40/- per sq.ft		49480	49680	
Infrastructure charges Rs 75/- per sq.ft		92775	93150	
Maintenance Charges for 12 months @ Rs. 3.50/sq.ft		51954	52164	
CMWSSB & TNEB Charges		100000	100000	
Documentation Charges		25000	25000	
Club Amenities charges		50000	50000	
GST@18%		57551	57657	
Total Additional Charges Including GST@18%		**426760**	**427651**	**428**
				PAYMEN
Application Amount		300000	300000	
Allotment stage Payment (To be paid on or before15th day from Booking Date)	**10%**	500631	503706	
1st Quarterly Instalment - 11th Quarterly Instalment (To be paid on or before 45th day from Booking Date)	**85%**	6805366	6831499	
12 Quarterly Installment upon Intimation on Possession (Inclusive of Additional charges)	**5%**	827076	829503	
Total Payable		**8,433,073**	**8,464,709**	**8,**

Figure A17
Example Price Sheet from Propequity

This extract presents the detailed breakdown of the costs covered by our data provider. Our estimate of the market value includes both the property purchase cost and other ancillary services.

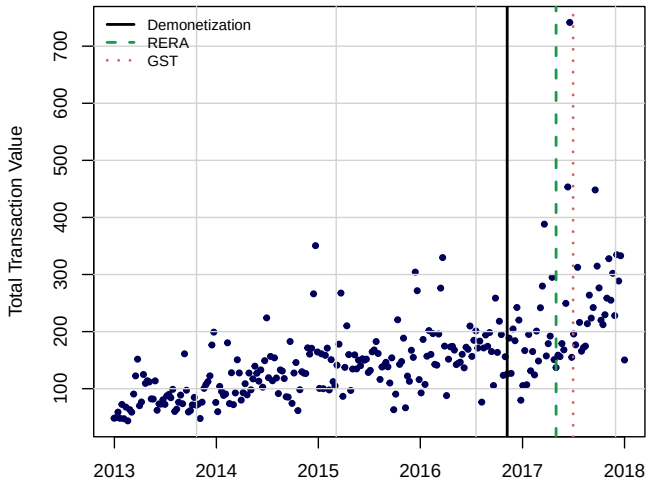


Figure A18
Weekly Total Estimated Transaction Value

See text for description of demonetization, and Appendix E for description of RERA and GST policies.

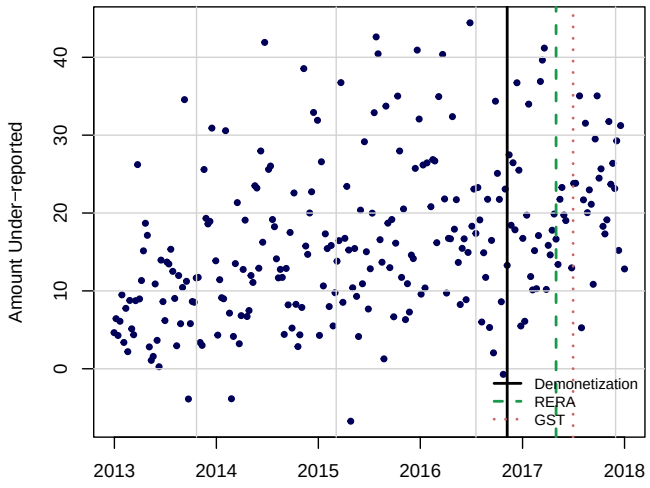


Figure A19
Weekly Estimated Amount Under-Reported (US$ Millions)

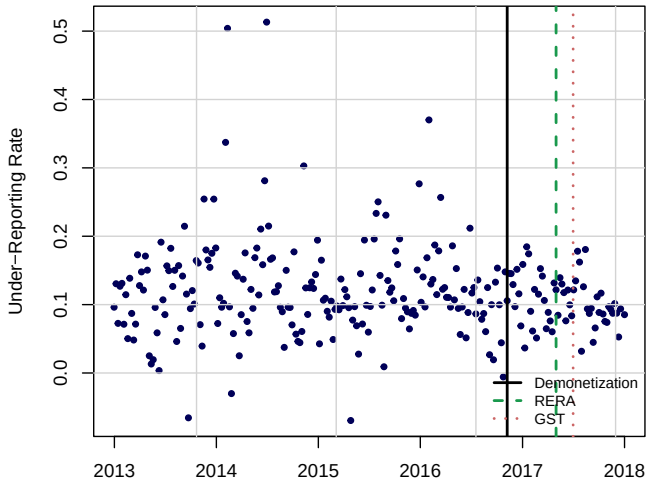


Figure A20
Weekly Estimated Under-Reporting Rate

(a) Fraction of Sample

(b) Under-Reporting Rate

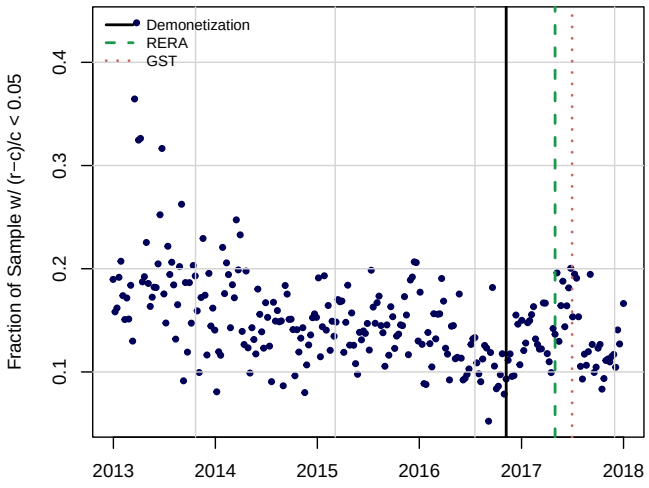


Figure A22
Fraction of Sample with $(r-c)/c < -.05$

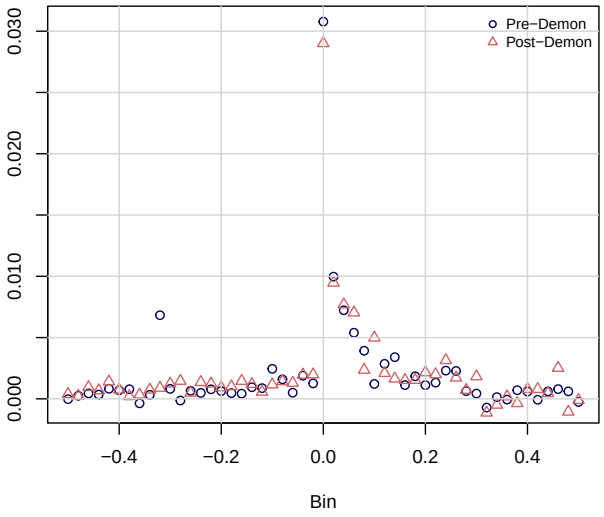
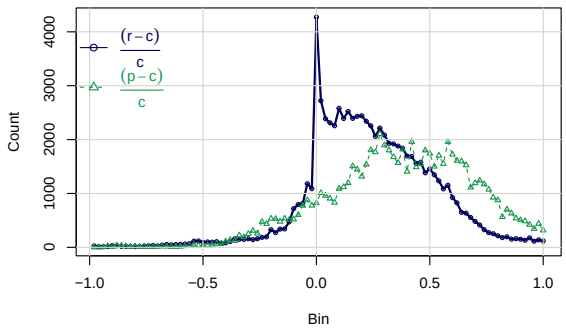
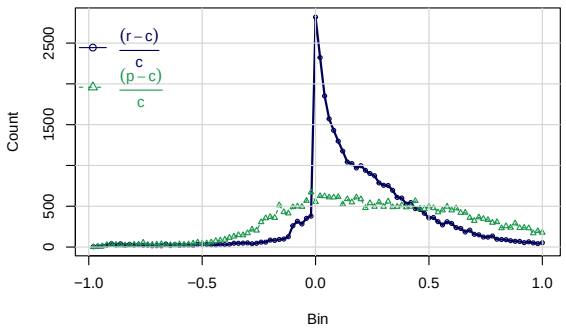


Figure A23
Contribution of Bin to under-reporting in (r-c)/c Bins 180 Pre- and Post-Demonetization

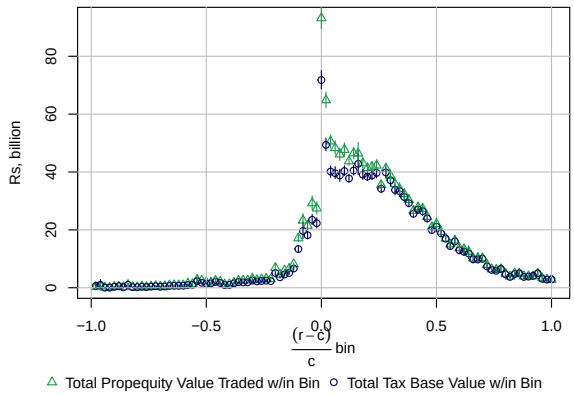
Each point is equal to the total amount of under-reporting within a $\frac{r-c}{c}$ bin divided by the aggregated estimated amount of transactions in the given time-period (pre-demon or post-demon). The sum of all points in the Pre-Demon (Post-Demon) series adds up to the aggregate under-reporting rate in the Pre-Demon (Post-Demon) Period.



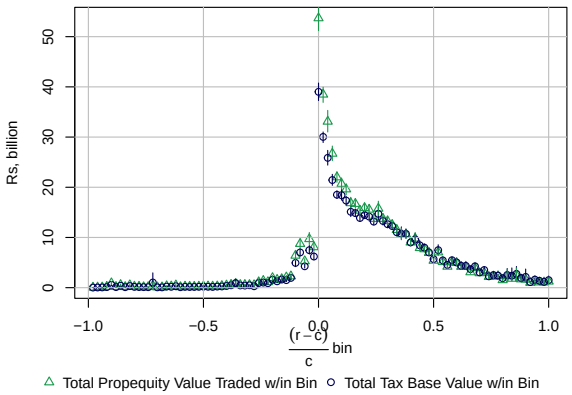
(a) Developer Sales



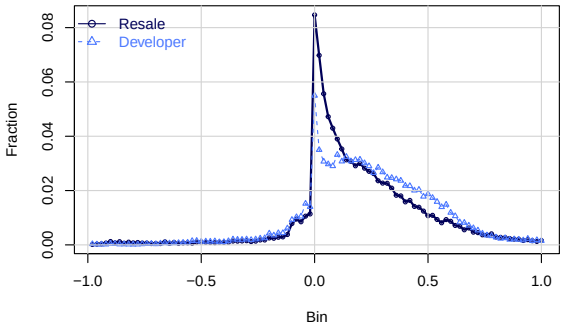
(b) Resale



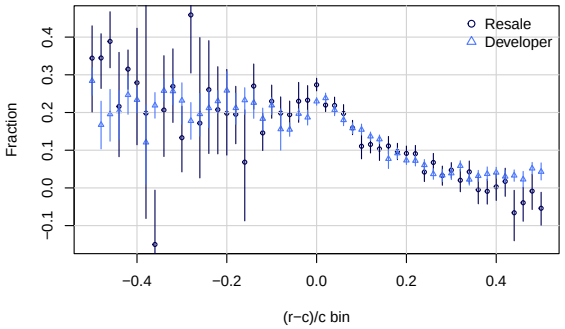
(c) Developer Sales



(d) Resale



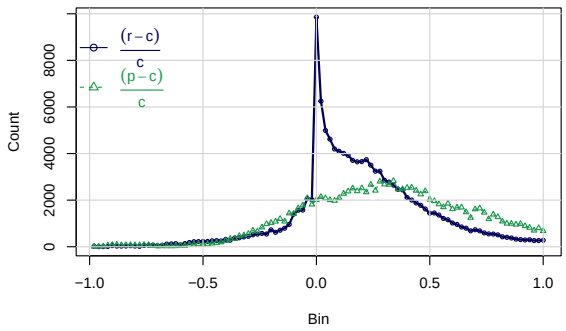
(e) Reported Counts (Density) Comparison



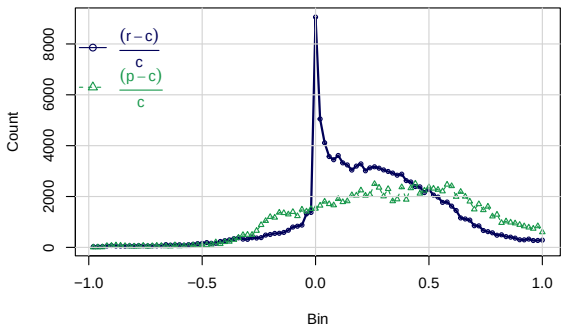
(f) Under-Reporting Rates

Figure A24
Developer vs. Resale Heterogeneity (Since Jan 2019)

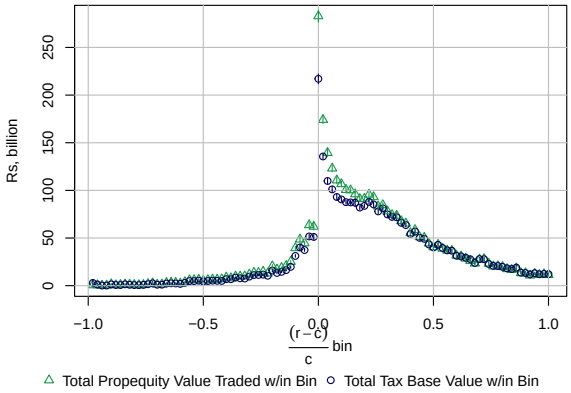
See Figures 5a, 5b and 5c for detailed descriptions. This restricts the sample to begin from January 2019 when all transactions are flagged as developer or resale transactions.



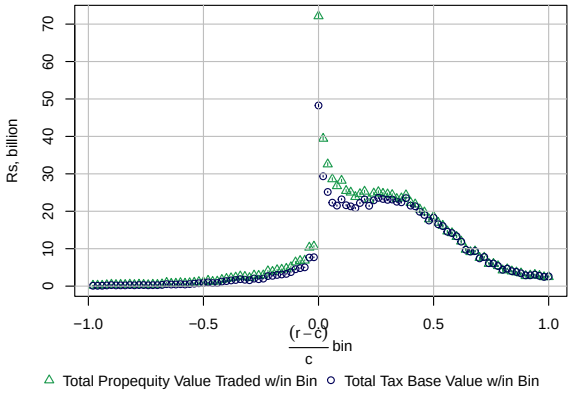
(a) Above Median Circle Value



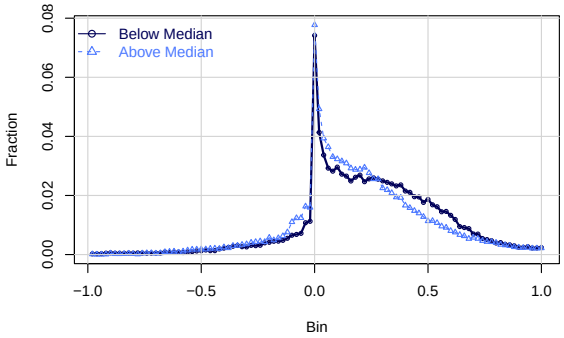
(b) Below Median Circle Value



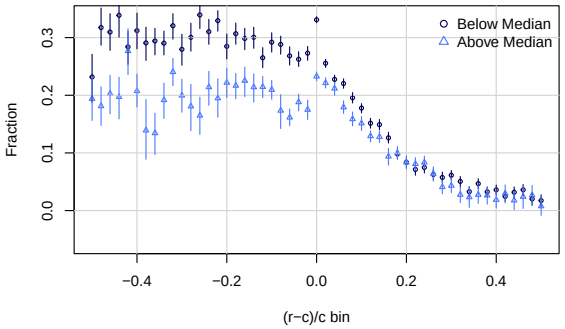
(c) Above Median Circle Value



(d) Below Median Circle Value



(e) Reported Counts (Density) Comparison



(f) Under-Reporting Rates

Figure A25
Above vs. Below Median Circle Value Heterogeneity

See Figures 5a, 5b and 5c for detailed descriptions.

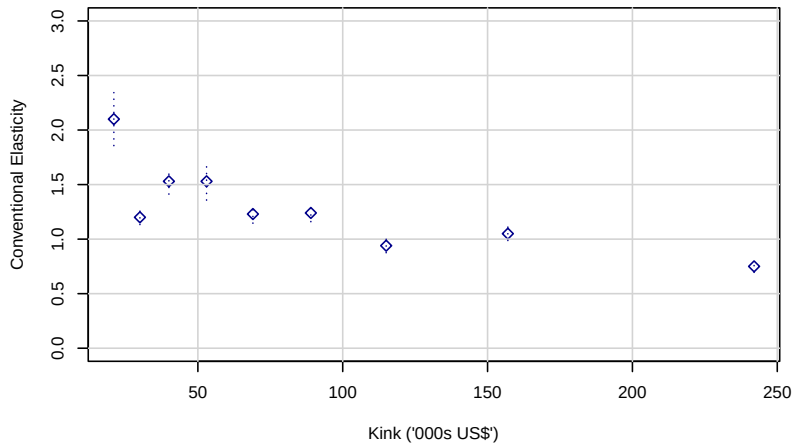


Figure A26
Reported Value Elasticity to Transaction Tax Rate

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.

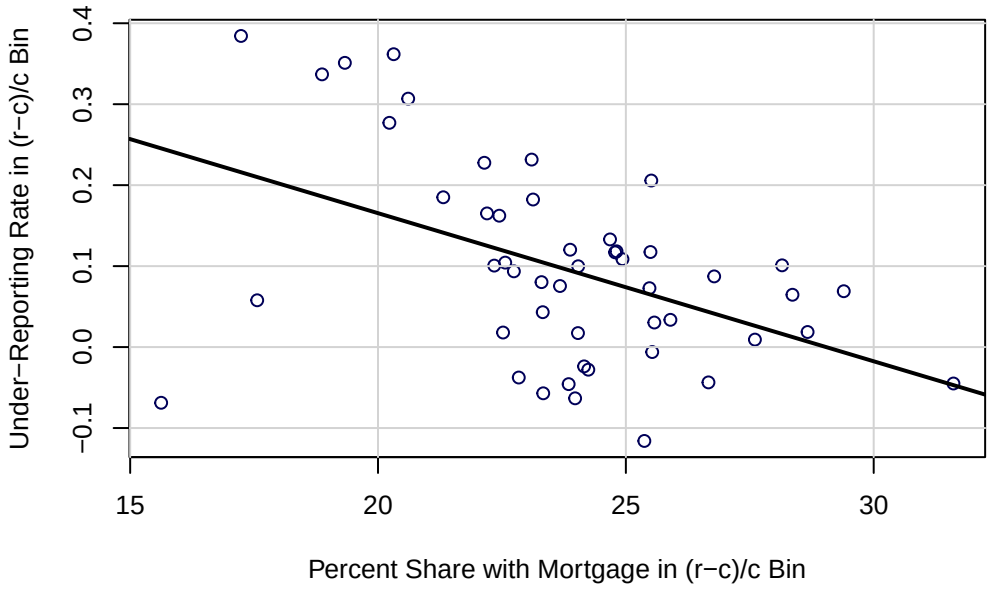


Figure A27
Incidence of Mortgages and Under-reporting

This figure plots the percentage of mortgage-based transactions in each $(r-c)/c$ bin (x-axis) and the under-reporting rate for each $(r-c)/c$ bin (y-axis). The black line presents the fitted line from a linear regression.

Table A1
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 A1 – 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 A1 – 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%
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

Continued on next page

Table A1 – continued from previous page

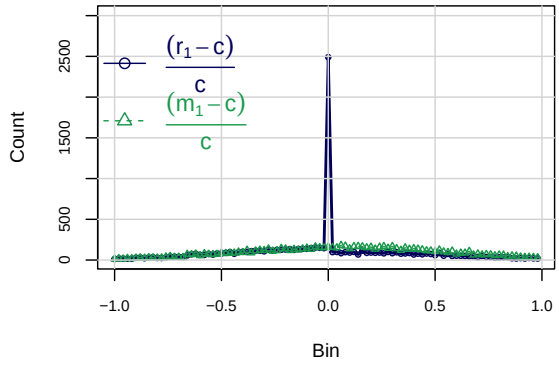
City	Country	Transaction Tax Rate	Additional tax	Assessment Method	Registration Fee
Toronto	Canada	0.5-2.5%	Ontario land transfer tax	Decentralized	
Belo Horizonte	Brazil	3%		Centralized	
Singapore	Singapore	1-3%	5-15% additional buyer stamp duty dependent on citizenship	Decentralized	SGD70 (US$52)
Philadelphia	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%
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	

Table A2
Under-reporting Before Circle Rate Changes

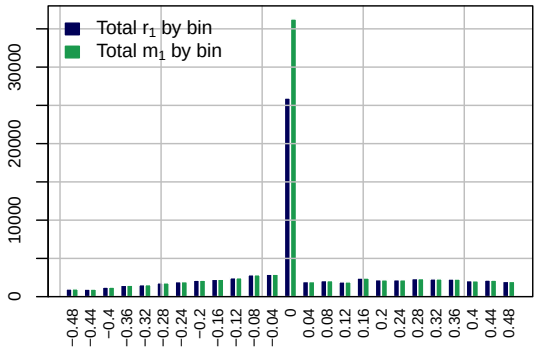
Dep Var: Under-reporting Rate	(1)	(2)	(3)	(4)
Month Before Policy Change	0.061***	0.054***	0.064***	0.068***
	(0.000)	(0.000)	(0.000)	(0.000)
Mean Dep Var.	0.06	0.06	0.06	0.06
Intercept	Yes	Yes	Yes	Yes
Time-trend	No	Yes	Yes	Yes
Month of year FE	No	No	Yes	Yes
Year FE	No	No	No	Yes
No. Obs.	260,614	260,614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

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

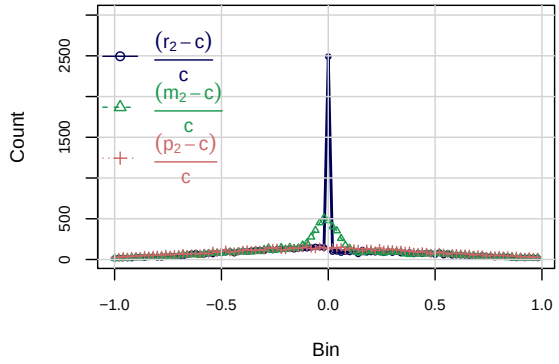
B Measurement Error



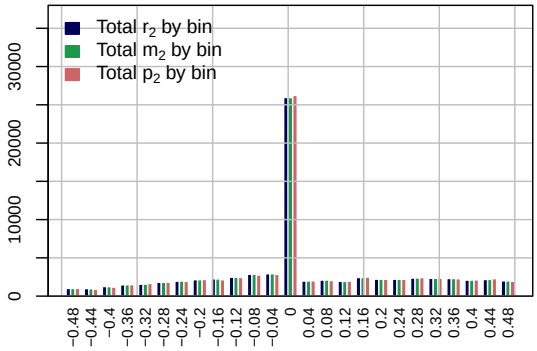
(a) High under-reporting, without measurement error



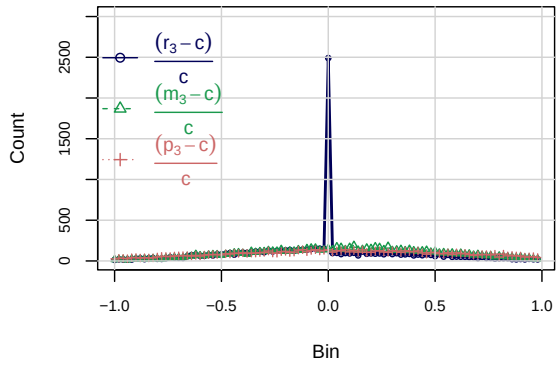
(b) High under-reporting, without measurement error



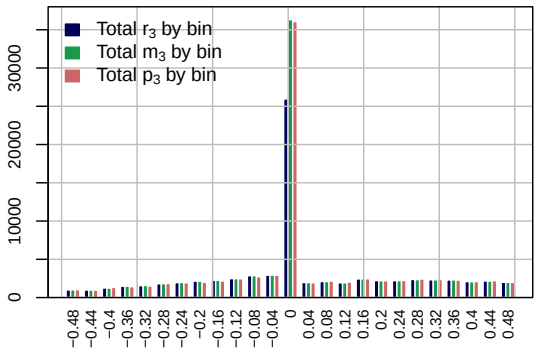
(c) No under-reporting, with measurement error



(d) No under-reporting, with measurement error



(e) High under-reporting, with measurement error



(f) High under-reporting, with measurement error

Figure B28
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 guidance value. Measurement error refers to noise in our estimates of market prices relative to the true unobserved market price.

Let $f^j(m)$ be the probability density function of market values for all households of type j, where $j \in \{\text{under-reporter,truthful-reporter}\}$. From the simplified model, observing household $i's$ triple (m, r, c) reveals its (unobservable) type perfectly. If $r = c$ and $m > r$, then $j =$ under-reporter. If $r = m$, then $j =$ truthful reporter. θ (the share of under-reporters) is then simply the fraction of all households with $r = c$ and $m > r$. The *aggregate* amount of under-reporting is the difference between aggregate m and aggregate r for households with $r = c$, because aggregate under-reporting for truthful reporters equals zero, and all under-reporters bunch at $r = c$. At the bunch point c, under-reporters are well identified.

Figure B28a simulates the distributions of r and m around c assuming that $\theta = 40\%$, $(m - c) \sim \mathcal{N}(\mu = 1, \sigma^2 = 10)$, and $c \sim \mathcal{N}(\mu = 10, \sigma^2 = 1)$. The x-axis is $\frac{r-c}{c}$ for the r distribution, and $\frac{m-c}{c}$ for the m distribution. The figure reveals substantial bunching of r around c, with 40% of households with $m \geq c$ choosing to report c. The underlying θ parameter can be backed out by inspecting how the (bunched) distribution of r around c differs from the (smoother) distribution of m around c.

As mentioned, the extent to which the government-assessed values c track market values m is a key confound. To evaluate this issue, we need an independent measure of m which is neither r nor c. For now, Figure B28a is drawn under the assumption that we have access to a measure p which is a perfect estimate of m.

B.1 Estimating Under-reporting with Measurement Error in the Market Value Proxy

As discussed in the paper, the patterns documented in Figure 5a are consistent with under-reporting, as shown in Figure B28a. These patterns are also potentially consistent with *no* under-reporting in the special case when c perfectly tracks m and p, the proxy for m, is measured with substantial error. To develop intuition, Figure B28 shows two (extreme) simulations from the simplified model with $\hat{\pi}$ set to zero. Both exhibit bunching of r at c, but generated from very different levels of under-reporting. In the "high under-reporting, without measurement error" case, $\theta = 0.4$ (i.e., 40% of households under-report) under the assumption that p is an error-free proxy (i.e., $p = m$). The excess bunching mass in this case is only due to under-reporting.

The figure shows that this bunching pattern can also be generated in a very different "no under-reporting, with measurement error" case (i.e., $\theta = 0$). This is possible if government guidance values c *exactly equal* market values m. This could happen for two possible reasons: 1) if the tax authority sets c values to perfectly match m; and 2) if buyers and sellers perfectly anchor on c when negotiating m.

Assessing which of these scenarios is more consistent with the data is difficult because we do not directly observe m; we only observe p, the proxy of transaction prices, which can be noisy. This leads to a confound since the excess bunching mass could either be generated by under-reporting, or by perfect coincidence of c and m plus an incorrectly measured counterfactual p which differs from m. The simulation is a stark illustration of this confound to emphasize the point that bunching at c alone does not perfectly identify under-reporting behavior. (Stark because we illustrate this using the identical distribution of reported values around guidance values ($\frac{r-c}{c}$) and show it can be consistent with $\theta = 0$ *or* $\theta = 0.4$).

B.1.1 Simulation to Illustrate Effects of Measurement Error in p

Figure B28c 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 B28c we set $p = m + \epsilon$, where $\epsilon \sim \mathcal{N}(\mu = 0, \sigma^2 = 16)$. 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 a smoothed version of the bunching in both m and r around c.

This confound is difficult to surmount since true transaction prices are generally unobservable given incentives to under-report. Reliance on proxies such as the Propequity measure that we use generate concerns of potential measurement error in such proxies. Hedonic models unfortunately are generated using observed reported values and therefore do not offer a direct solution to this challenge. The next section discusses our proposed solutions to derive sensible inferences in the face of measurement error.

B.1.2 Dealing with Classical Measurement Error: Aggregation

The first approach to distinguish the case with measurement error from genuine under-reporting is to compare *aggregated* reported and market values by $\frac{r-c}{c}$ bins. Figures B28b and B28d plot simulated aggregated reported and market values for the two cases considered above (in Figures B28a and B28c respectively). In these figures, we sum all measures of value (p, r, c, m) for the simulated transactions within $\frac{r-c}{c}$ bins of 0.02 width. In both figures, the x-axis indicates the total reported value of transactions; for example, in Figure B28b at zero, the blue bar is the sum of all r for transactions with $-0.02 \leq \frac{r-c}{c} < 0.02$, and the green bar is the total market value of transactions in each $\frac{r-c}{c}$ bin. In Figure B28d, we separately plot red bars, which show the sum of all p within each bin.

In Figure B28b, with simulated high under-reporting, but no measurement error in p, we can clearly see under-reporting in the zero bin, where aggregated r is substantially lower than aggregated m (since there is no measurement error assumed, the sum of all m also equals aggregate p). In contrast, in the simulation in which there is no under-reporting, but substantial measurement error in p, Figure B28d shows that aggregated m and aggregated noisy p are (approximately) the same within each bin. The insight here is that aggregation within bins smooths out symmetric measurement error in p, and allows us to observe differences between r and m. This allows us to distinguish between the two cases, which are identical in the simulated count/bunching distributions plotted in Figures B28a and B28c. Put differently, truthful reporting with measurement error will exhibit no mass difference between aggregated p and r, and under-reporting will result in a mass difference between aggregated p and r in the central bin.[38] Of course, the data is likely to feature a mix of under-reporting and measurement error in the proxy p for m, and Figure B28f shows that in the presence of both under-reporting and measurement error, aggregation within bins smooths out measurement error and still facilitates 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 5a. To do so, we assume that $r_i = m_i$ (truthful reporting, no under-reporting), that c perfectly tracks m, and we simulate Propequity prices $p_i = m_i + \epsilon_i$. We find that with these assumptions, p that are on average 15% above m, with a standard deviation of 15% can replicate Figure 5a. This result illustrates the challenge in estimating under-reporting based purely on inspecting bunching in r

[38] 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 A7 present aggregated estimates with bin-width set to 4% and 8% respectively, confirming this in the data.

around c versus that in p, though it relies on the strong assumption that $c = m$.

Figure 5b applies the aggregation procedure to the actual data, plotting aggregated reported and Propequity values in the data by $\frac{r-c}{c}$ bins. The green triangles indicate the total amount of Propequity estimated value (p) transacted amongst transactions within a 0.02 width $\frac{r-c}{c}$ bin, and the blue dots indicate the total value reported (r) for the same set of transactions within the $\frac{r-c}{c}$ bin. The figure is consistent with Figure B28f, i.e., we confirm that there is under-reporting using the aggregation approach, and the figure reveals that the largest amount of unreported value (the differential between aggregated r and p) comes from bunching transactions that report r exactly equal to c. It is also worth noting that in Figure 5b, the overall overlap of the p and r distributions is tighter than in the plot of transaction counts in Figure 5a, with the sharpest deviation evident between the distributions at exactly c. This pattern is more closely consistent with the "two-type" simplified model (i.e., assuming $\hat{\pi} = 0$) presented in the beginning of this section, which predicts an atom of mass in the distribution of r at c and a smooth distribution that is closer to p otherwise.[39]

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 guidance value, there is no visible gap between aggregate reported value r and estimated market values p. Going back to our model, this finding suggests heterogeneity in $\hat{\pi}$ across the $\frac{r-c}{c}$ distribution that results in buyers reporting more truthfully as we move away from c. Overall, when we aggregate the rupee value of under-reporting across all bins in Figure 5b, we estimate ₹455.1 billion (US$9.1 billion) in under-reported real estate value in our sample over the period 2013-2022.

Figure 5c directly shows the fraction of estimated market value under-reported by $\frac{r-c}{c}$ bin. Transactions with r up to 50% lower than c have an estimated under-reporting rate of roughly 20% on average.[40] There is a sharp discontinuity at zero and a marked change in slope as we move to transactions that report $r > c$, with estimated under-reporting dropping to zero for transactions that with r 50% above c. For properties that report 60% above c, we estimate negative under-reporting (i.e., over-reporting) rates,

[39] For bins with $\frac{r-c}{c} < 0$ (i.e. transactions with reported values less than guidance values) the blue circles in Figure 5b aggregate c. This is because if $r < c$ the tax base is effectively c, i.e., the guidance value, since the government assesses taxes at c pending a successful appeal. Appendix Figure A8 replaces the blue circles in Figure 5b with the total r within each bin, which means that to the right of zero, these figures are identical, but will differ to the left of zero. The figure shows that the total aggregated r values 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.

[40] Note that these transactions are measured as reporting the guidance value (because when $r < c$ the tax base is assumed to be c). A caveat is that it may be that for some of these transactions, $m < c$ and buyers disputed the accuracy of the government-issued guidance value c, so the under-reporting rate at zero (i.e., $r = c$) is potentially more accurate as it is less subject to this confound.

although these estimates are based on a relatively small number of transactions (the counts per bin in this region are shown in Figure 5a), and the 95% confidence interval includes zero.[41]

B.1.3 Dealing with Biased Measurement Error

While our aggregation approach helps to assuage concerns of classical/symmetric measurement error, it is possible that measurement error can be biased. This section explores this issue further, estimating the form that measurement error would be required to take to explain the empirical patterns we observe. We conclude that for biased measurement error to drive our results, we would need to attribute the bunching patterns that we observe in r around c solely to truthful reporting, and assume that measurement error in the p proxy perfectly mirrors the "excess bunching mass" identified in Figure 5b, which seems implausible.

We have this far discussed p as a potentially noisy but unbiased estimator of m. A more general measurement error model is $p = m + \nu(r, c, m) + \epsilon$. Here the $\nu(r, c, m)$ is a bias term that can itself be a function of the transaction's reported, guidance or true market values. In reality any form of this function is possible; indeed a form of this measurement error function can always be found that fits the data and rationalizes perfectly truthful reporting under the assumption that c exactly equals m. For example, the pattern of under-reporting across bins in Figure 5c could be explained by truthful reporting at m plus a $\nu(r, c, m)$ function with the largest measurement error for bunching transactions reporting $r = c$, moderate measurement error for transactions with $r < c$, and declining measurement error as r increases relative to c. Ultimately, the question boils down to whether a specific form of measurement error (plus the assumption that m equals c for a large number of transactions) is more or less plausible than the hypothesis that agents under-report property values.

To illustrate this issue, we assume perfectly truthful reporting $r = m$, solve for ν and then present the form that measurement error would be required to take to explain the empirical patterns we observe.[42] We ignore the classical measurement error term ϵ as this averages out when we analyze bin-aggregated data. If $r = m$ then we can solve (post-averaging) for $\nu(r, c, m)$ as $\nu = p - r$. We can also decompose the total bias in measurement error $\bar{\nu}$ by reporting bins. Let $\bar{\nu}$ be the aggregate measurement error

[41] Within each bin, we sample with replacement the same number of observations. This allows us to compute the aggregate values and under-reporting rates for different samples. We then compute the bootstrapped standard errors and then construct the 95% confidence intervals reported in Figure 5c. We use the same procedure for the full sample, and also in other cases where we construct standard errors in this paper.

[42] To be more precise, the tax-base value is equal to $\max[r, c]$; under completely truthful reporting we assume that the tax-base value is the true market value, but we simply write $r = m$ here for parsimony.

bias, equal to the sum of ν across all transactions divided by the sum of true transaction values (i.e., the denominator is the sum of reported tax base values $\max[r, c]$ under the assumption of perfectly truthful reporting). Let $j \in 1, ..., J$ index $\frac{r-c}{c}$ bins (2% bins in our case), and let N_j be the number of transactions in bin j. Then:

$$\bar{\nu} = \frac{\sum_{j=1}^{J} \sum_{i=1}^{N_j} (p_{ij} - r_{ij})}{\sum_{j=1}^{J} \sum_{i=1}^{N_j} r_{ij}} \tag{13}$$

$$= \frac{1}{\sum_{j=1}^{J} \sum_{i=1}^{N_j} r_{ij}} \left(\sum_{i=1}^{N_1} (p_{i1} - r_{i1}) + \sum_{i=1}^{N_2} (p_{i2} - r_{i2}) + \ldots + \sum_{i=1}^{N_J} (p_{iJ} - r_{iJ}) \right), \tag{14}$$

where again, r stands in for the tax base value $\max[r, c]$.

Each term in the sum in the numerator of equation (13), under the assumptions made about truthful reporting, is the contribution of biased measurement error in a given bin to the aggregate (truthfully) reported value. Appendix Figure A9 plots these terms estimated in the data, and shows that under the hypothesis of truthful reporting, measurement error in p must be substantially positively biased for bunching transactions at $r = c$, and decline monotonically with the gap between r and c. Put differently, for measurement error to drive our results, one would need to attribute the bunching patterns that we observe in r around c to truthful reporting, and attribute a very specific form of measurement error in the p proxy that essentially mirrors the "excess bunching mass" identified in Figure 5b.[43]

We highlight three results from this exercise. First, the form of measurement error required for truthful reporting to explain the patterns in the data does not fit with many plausible sources of measurement error, such as one in which p is some fixed percent higher than true m, regardless of reporting behavior (as would be the case if Propequity prices were inflated versions of true market values). Second, under the hypothesis of truthful reporting plus biased measurement error we would not see bunching of r at c unless guidance values are set extremely carefully to match market prices. We investigate this issue more deeply below, and find substantial evidence to the contrary. Third, if we assume all deviations between p and tax base values (i.e., $\max[r, c]$) arise from measurement error, this sets a floor of zero on the total amount of under-reporting, and if we instead assume that p is a perfect proxy and all deviations come from under-reporting, we get an upper bound on under-reporting of 10.94% (95% bootstrapped C.I.

[43] Figure A10 plots the bin-level measurement error rate (i.e., each term on the rhs of equation (13)) under the hypothesis that measurement error explains our data. Again, we see that the measurement error rate would specifically have to be highest for bunching transactions and then decline linearly for transactions where buyers self-select in reporting $r > c$.

= [10.8%, 11.31%] in our data. A conservative interpretation of our results is therefore that under-reporting lies between zero and 11% in aggregate, which is substantially lower than the anecodotal and small-scale estimates in this market, and also low relative to tax-evasion estimates for self-employed workers in developed countries.

We pursue this investigation further below, however, where we describe additional approaches to addressing potential measurement error in our proxy p for m.

B.1.4 Exact Matching of Propequity Projects to Transactions

A different way to assess the importance of measurement error in p is to focus on a subset of the data that has plausibly less measurement error. Appendix Figures A11, A12 and A13 present counterpart results to our analyses in Figures 5a, 5b and 5c estimated on a restricted dataset of the 60% of transactions where there is an exact match between building names in the Propequity data and the Registrar transactions data. To the extent that measurement error arises from imperfect matching across buildings (e.g., transactions in older buildings being matched to new luxury building launches in Propequity) we expect this sample to be less affected by measurement error. Appendix Figure A11, the replication of Figure 5a, exhibits a very similar shape to the full sample version, indeed the bunching patterns appear slightly sharper (bins just to the right of zero do not show as much increase in mass as the bins just to the right in the full sample version). The shape of estimated levels of under-reporting by bin is essentially the same (Appendix Figure A12 versus 5b), but the exact match sample levels are lower in some bins.[44] Overall, the aggregate under-reporting rate in the exact match sample is 9.48% (95% bootstrapped C.I. = [9.37%,9.48%]), which is similar to the full sample estimate of 10.94% (95% bootstrapped C.I. = [10.8%, 11.31%]).

B.1.5 Stale Guidance Values

For measurement error to explain our results rather than under-reporting, c must perfectly track m across both time and space to explain visible bunching of r at c. However c is set in a geographically coarse manner, and infrequently updated. Over the sample period, real estate prices have grown substantially, with growth rates that vary considerably in different Mumbai sub-regions, making it implausible that c perfectly tracks m over time given the government's process for setting c. Moreover c is set at a relatively broad geographical level. Given the considerable spatial variation of property values, it is also implausible that c perfectly tracks m spatially. In the data, c per square meter values are very close or the same for all properties within each sub-zone

[44] For example, in the exact sample match the estimated under-reporting rate when $\frac{r-c}{c} = 0$ is approximately 0.2, whereas it is 0.25 in the full sample.

(the average (median) subzone in the data is 686,818 (264,136) sqm).[45] This means that a single guidance value for a large region is unlikely to be an accurate reflection of the full distribution of the true value of the assets in that region at any given point in time.

Despite these issues, it might still be the case that c values are set carefully to match the first spatial moment of m. However, if this is true, 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. This generates a concrete prediction. 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 I1).[46]

In contrast, if buyers and sellers under-report to evade taxes, and with the dates of c changes publicly announced in advance, we would expect that bunching would be greatest immediately *before* any increases in the government-assessed value. If c predictably increases, there is a predictable jump in the tax burden incurred by under-reporting the transaction value at $r = c$ 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 A14b documents behavior that is consistent with the under-reporting explanation, and inconsistent with either the measurement error or anchoring explanations. The figure shows that bunching mass at $r = c$ spikes immediately *prior* to scheduled guidance value changes but there are no corresponding spikes in the third-party estimated proxy p.[47] While this is clearly evident in the plot, it is diffi-

[45] 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.

[46] 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, given further evidence described below, 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.

[47] The guidance values were increased in 2014, 2015 and 2016. The Maharashtra government chose

cult to draw strong conclusions relative to the month-by-month variation overall from pure visual inspection. Table A2 therefore 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 (statistically significant at the 5% level) 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%.[48]

Finally, more support for the under-reporting explanation is provided by Figure A15, which 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 clear pattern of backdating agreement dates just prior to guidance value changes to take advantage of lower guidance values.

B.1.6 High-Frequency (Daily) Reporting Behavior

Figure A16 studies *daily* reporting behavior around scheduled guidance value changes (as indicated by the green vertical lines in A14a).[49] Figure A16a shows a large spike in registered transactions on the day directly before the scheduled guidance value change, Figure A16b shows that the bunching rate for the large number of transactions registered right before the guidance value increase is approximately 10% higher, and Figure A16c shows the fraction of transactions registered below the guidance value which also shows an increase in the days prior to the guidance value change. These results suggest that even non-bunching transactors prefer to avoid the new guidance values which are expected to increase their tax burden. This could happen, for example, if the new guidance values are above what buyers planned to report.

C 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

not to increase the guidance values in 2017, 2018, 2019 and 2020, and decreased the rate as a result of the pandemic in 2021.

[48] This sample average is the simple average under-reporting rates across transactions; this differs from our aggregate under-estimating estimate of 10.94% because it does not weight transactions by their size.

[49] The dates of scheduled guidance value changes have moved around over time and are not always on 1 January each year. This reduces concerns that these dates coincide with policy announcements that are routinely made at the turn of the year.

guidance value) to five percent above the kink. A household maximizes the following utility function to 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.

$$\hat{e} \approx \frac{\hat{B}}{z^* \cdot h_0(z^*) \cdot \log(\frac{1-t_0}{1-t_1})} \tag{15}$$

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.[50] To estimate the counterfactual density, we fit a polynomial of a specified degree to the observed reporting density, excluding the data in a specified window around the kink, using the following specification:

$$C_j = \sum_{i=0}^{q} \beta_i^0 \cdot (Z_j)^i + \sum_{i=R_l}^{R_u} \gamma_i^0 \cdot \mathbf{1}[Z_j = i] + \epsilon_j^0. \tag{16}$$

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 Chetty et al. (2011) is to set a symmetric bunching win-

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

dow, 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 Chetty et al. (2011) 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.[51]

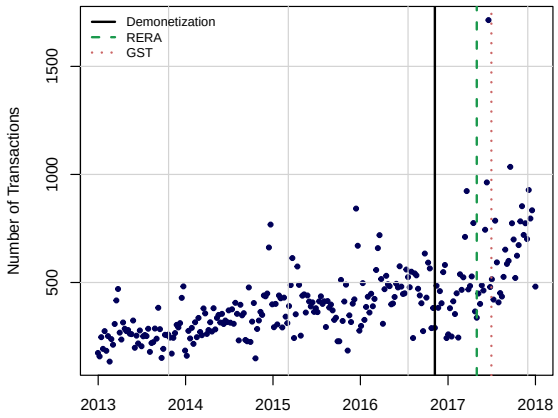
The second step is to compute the excess mass of reported values around the kink relative to this counterfactual density. Using equation (16), we compute the counterfactual mass in each bin within the bunching window, $\hat{C}_j^0$. Subtracting this predicted mass from the observed density yields the estimated excess number of individuals who report values near the kink relative to this counterfactual distribution:

$$\hat{B} = \sum_{i=R_l}^{R_u} C_j - \hat{C}_j^0 = \sum_{i=R_l}^{R_u} \hat{\gamma}^0. \tag{17}$$

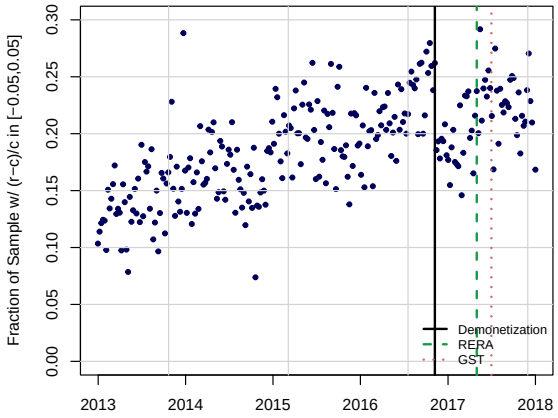
Standard errors for $\hat{e}$ are estimated using a bootstrap procedure. We resample with replacement from the underlying distribution of transactions 10 times, re-estimating the elasticity each time, and defining the standard error as the standard deviation of the distribution of $\hat{e}$ estimates.

[51] 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."

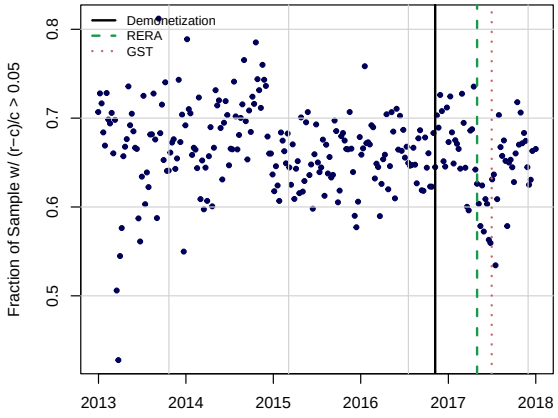
D Impact of India's Demonetization on Under-Reporting



(a) Transaction Count



(b) $-.05 \leq (r-c)/c \leq .05$



(c) $(r-c)/c > .05$

Figure D29
Reporting Behavior in Months Around Demonetization

Figures D29b and D29c present the fraction of sample in each $(r-c)/c$ bin.

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 be frequently 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. Nearly 100% of the outstanding ₹500 and ₹1,000 were ultimately deposited in banks (see, e.g., Lahiri 2020), suggesting that the policy did not expropriate 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; Karmakar and Narayanan (2020) finds the largest consumption impacts for households without bank accounts.[52]

As we have described, a potentially significant illicit use of cash is to facilitate property-value under-reporting and tax evasion via the unaccounted payment of cash from buyer to seller. Indeed, eliminating this form of evasion was a primary stated goal of the policy.[53] Post-demonetization, with less cash in the economy, we might therefore 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 of the government's intention to crack down on cash transactions. We might also therefore expect the trend of under-reporting behavior to change. On the other hand, real estate developers may also enable non-cash mechanisms to undervalue properties. For instance, real estate developers may price new apartment purchases at a lower value and simultaneously invoice buyers for services (operations, maintenance, infrastructure charges, etc.). While the purchase contract is completely under the gaze of law, service contracts do not have similar transparency, thus allowing buyers to pay for such contracts using electronic payment. This dual-invoicing mechanism, reminiscent of Fisman and Wei (2004), could allow developers to make up for any reductions in the reported value

[52] Nigeria engaged in a similar demonetization exercise in October, 2022 (*Riots erupt in Nigerian cities as bank policy leads to scarcity of cash* 2023).

[53] In February 2017 (three months after the demonetization policy was implemented) Union Finance Minister Arun Jaitley stated: "Common people have supported us graciously without asking for anything in return. I believe that demonetization was ideal. Was there even one major demonstration or unrest against demonetization? Isn't it true that 75% of the people supported it? It is [Isn't it] true that cash as a way of life has changed today? Nobody will tell you to pay 40% cash if you buy property. People have started changing themselves and the long-term impact will be seen soon" (Das 2020).

on a transaction deed through a separate electronic payment—and could explain why under-reporting continues even when cash became scarce due to demonetization.[54]

D.1 Demonetization Aggregate Effects on Under-Reporting

Figure D29a shows weekly transaction counts around the demonetization event; we do not observe a sharp reduction in the total number of transactions in the weeks surrounding demonetization, despite the large macroeconomic shock it occasioned. Appendix Figures A18, A19, and A20 show no sharp changes in the total estimated value of transactions, the total amounts under-reported, or the estimated under-reporting rates respectively. Overall, these plots suggest that demonetization did not have a major impact on aggregate real estate transactions or aggregate measures of under-reporting in Mumbai.[55] We next look deeper to see if demonetization affected the composition of real estate transactions.

D.2 Decomposing Demonetization Effects

We decompose aggregate changes around demonetization into changes within different parts of the $\frac{r-c}{c}$ distribution. We expect the greatest changes in "buncher" transactions (with $r = c$), as these had the highest estimated under-reporting rates and therefore should be most affected by the cash crunch induced by demonetization. To this end we analyze transactions falling in three sets, namely: $\frac{r-c}{c} < -0.05$; $-0.05 \leq \frac{r-c}{c} \leq 0.05$; and $\frac{r-c}{c} > 0.05$.

Figure D29b plots the fraction of transactions whose reported value is within 5% of the transaction's guidance value. There is a clear drop evident in the fraction of these "bunching" transactions immediately following the week of demonetization; an event study model estimated on the weekly data translates this drop into a 5.2% decline in the fraction of transactions reporting just around the guidance value, which is a roughly 20% decrease given that approximately 25% of transactions bunched just prior to demonetization. Figure D29c shows a corresponding increase in the fraction of transactions reporting over five percent more than their guidance value. (Appendix Figure A22 shows the weekly fraction of transactions with $r < 0.95c$; here there is no discrete jump, but this fraction drifts up slightly over the six months after demon-

[54] Our measure p of market value attempts to include costing for such services, so we do not expect this to be a major form of biased measurement error in our estimates of market prices. Appendix Table A17 presents an example of a price sheet provided by our data provider showing that such service costs are included in property valuation.

[55] This evidence is in interesting contrast to Alvarez and Argente (2022), who find large negative quantity effects of government restrictions on cash usage for Uber rides.

etization.) Two other major economic reforms affected the housing market shortly following demonetization, namely the passing of the Real Estate Regulation Act or RERA, and introduction of the Goods and Services Tax or GST. There is also a decline in the fraction of transactions bunching after the introduction of these reforms, and a corresponding increase in the fraction of transactions reporting 5% or more above the guidance value.[56] Despite this visible redistribution of transactions from bunching to non-bunching, Appendix Figures A19 and A20 confirm the aggregate results that there is no major change in the weekly amounts under-reported or the overall under-reporting rate.

Using a different approach to verify these results, Figure A21a compares the bunching behavior of transactions in the 180 days pre- and post- demonetization. Consistent with the weekly aggregate plots, approximately 3% fewer transactions bunch within 1% of the guidance post-demonetization relative to pre-demonetization. Figure A21b shows the estimated under-reporting rates pre- and post- when we aggregate all transactions within each bin. While both pre- and post-demonetization series show that bunching transactions have the highest estimated under-reporting rates, there is very little change in these estimated under-reporting rates by bin, which translates into very little change in aggregate under-reporting rates.[57] Overall, it seems that while demonetization seemingly deterred bunching transactions, these account for only a fraction of overall transactions; moreover, the bunching transactions that continued after demonetization had similar under-reporting rates to the bunching transactions prior to demonetization. This explains why we do not observe a significant change in overall under-reporting rates after demonetization. Figure A21c reconfirms this inference of small changes pre-and post-demonetization by showing the aggregated under-reporting amount in millions of USD within each bin pre- and post-demonetization periods. Indeed, this plot shows a slightly larger amount of under-reporting coming from buncher transactions after demonetization versus before, despite the fact that demonetization did reduce the number of buncher transactions overall. This is because while the number of bunching transactions post-demonetization is smaller, the average estimated market value of these transactions and their associated under-reporting rates are slightly larger, so the total amount under-reported is ultimately quite similar.[58] Furthermore, the under-reporting rate in non-buncher bins is similar across the

[56] Appendix E discusses these two major reforms during our sample period.

[57] Aggregating over all transactions, we estimate 254 million dollars of under-reporting in the 180 days prior to demonetization and 355 million dollars in under-reporting in the 180 days after demonetization; naturally this time-series difference can be explained by many factors, but overall it seems unlikely that demonetization itself reduced under-reporting amounts in a major way.

[58] The total amount of under-reporting in the 180 days before demonetization amongst buncher transactions is 134 million USD, which equals 3,321 transactions times an average estimated market value of

pre-demonetization and post-demonetization periods, which explains why we do not see an aggregate change in under-reporting in the weekly time-series figure (Appendix Figure A20).

To further quantify how the change in bunching changes the overall under-reporting rate, we can use a similar procedure to that outlined in equation (13), but now under the assumption that the proxy p is measured without error. This allows us to decompose aggregate under-reporting into bin-specific contributions both pre- and post-demonetization. Appendix Figure A23 plots each of these terms comparing pre-demonetization and post-demonetization amounts. The aggregate pre-demonetization under-reporting rate is 0.098 (95% CI = [0.090,0.120]), while the aggregate post-demonetization under-reporting rate is 0.110 (95% CI = [0.107,0.128]). The figure shows that the contribution to the total under-reporting rate from the $r = c$ bin is smaller in the post-demonetization period (consistent with the fact that there were fewer bunching transactions), however this difference is swamped by the bin by bin variation across the rest of the distribution. Overall, the results suggest that demonetization reduced the number of high under-reporting bunching transactions, but the effect on aggregate under-reporting in this market was likely small. We caveat these results noting that our results are specific to the Mumbai and Mumbai suburban districts, and that demonetization results in other locations may have been different.

E 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.[59] Overall, the purpose was to curtail the ability of developers to

0.191 million USD times a bin-level under-reporting rate of 0.211. The total amount of under-reporting in the 180 days after demonetization amongst buncher transactions is 140.7 million USD, which equals 2,593 transactions times an average estimated market value of 0.223 million USD times a bin-level under-reporting rate of 0.243.

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

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 May 1, 2017 were not affected by RERA.[60] Nonetheless, in Figure A20 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.[61] 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. We find no major changes in under-reporting rates around the GST reform in Figure A20.

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

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

F Hetereogeneity

F.1 Seller Type and Under-reporting Behavior

Anecdotally, sellers' differential ability to accept cash to fund the gap between r and m is an important factor that could help explain the extent of under-reporting. For example, households that run small businesses paying their workers and other suppliers in cash may be more likely to accept large currency payments. Conversely large developers may require honest buyer reporting because they prefer sales records to reflect the true income received from these sales, to facilitate transparent revenues 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). To understand these explanations, we condition bunching behavior on whether the seller is a corporate entity or an individual.[62]

F.1.1 Developer vs. Resales

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. Given that firms interact with many buyers, there is greater risk that one buyer may whistleblow under-reporting behavior leading to an audit or detection, similar to the logic of Kleven, Kreiner and Saez (2016), who argue 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 6a and 6b show the count of developer sales and resales by $\frac{r-c}{c}$ bin respectively. The bunching amongst developer sales is "sharper," in the sense that the increase in transactions exactly in the bunching bin is high relative to nearby bins. Bunching amongst resale transactions is also apparent, but counts decline more smoothly as we move away from the bunching bin. Figures 6c and 6d show the bin aggregated reported and market values for developer sales and resales. Figure 6e shows bunching behavior as a fraction of total transactions, so the developer and resale curves are comparable. There is more bunching for resale than for developer driven transactions, consistent with Kleven, Kreiner and Saez (2016). Figure 6f shows under-reporting rates by $\frac{r-c}{c}$ bin for developer sales and resellers. Both types of sales show substantial bunching

[62] There is also likely heterogeneity across buyers in access to such (untracked) cash. 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 (e.g., formal sector workers in say the IT sector) may have greater difficulty accessing large sums of cash without additional scrutiny, given the direct deposit of their income into bank accounts.

at $r = c$, 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, we do see more bunching and under-reporting for resale transactions.[63]

F.1.2 Heterogeneity by Transaction Amount

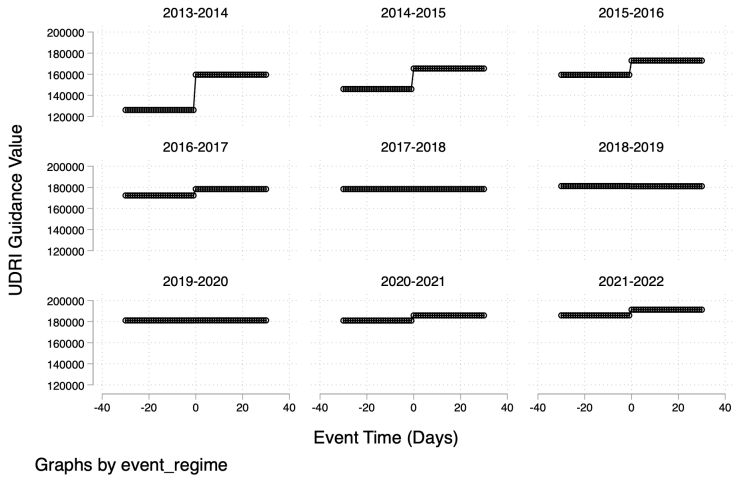
Are under-reporting rates higher for larger value transactions that may be more visible to the tax authorities? Figure A25 presents corresponding figures on densities, bin-specific aggregates of reported versus market values, and under-reporting rates, for transactions with above and below median guidance values. 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. These results are consistent with our findings in Appendix Figure A26, where we see higher reporting elasticities for low versus high valued transactions.

[63] In Figure A24, we re-estimate Figure 6 for a sample since January 2019, after which all transactions have been successfully classified as either developer or resale. The patterns are broadly similar and more precise.

G Estimating the Response of Transactions to Changes in Guidance Values

In this Section we describe our methodology for estimating the response of the quantity of housing transactions with respect to changes in guidance values (i.e. the "extensive margin" response to changes in guidance values.) Guidance values are changed annually, with the date varying across our sample. Figure G1 presents the time-series of average guidance values per square meter in the Mumbai region. The figure shows there were major average increases in 2013, 2014 and 2015 and minor increases in 2021 and 2022. These policy changes also include regional level changes, which we exploit in a difference-in-difference event-study design.

Figure G1
Guidance Value Policy Variation Over Time



We pursue two research designs. In the first, we split subzones into areas that received above versus below 10% increases in their guidance value. We treat each year as a separate "event," and create a balanced panel of subzones where we have an observation for each subzone in each of nine months before and after the guidance value change. We then pool all of the events in to one dataset and study the evolution of guidance values and transactions in the treatment and control subzones.

Figure G2 shows the the evolution of the mean logarithm of guidance values per square meter based on the legislative changes before and after the guidance value changes we study. Treatment subzones experience an average increase of approximately 18% in guidance values, whereas control subzones experience a less than 0.1%

increase. In terms of the economic magnitude, given that the transaction tax rate is 5%, the maximum impact of the guidance value increase would be on those who choose to report the guidance value (i.e. the "bunchers"), so an 18% increase in the guidance value would correspond to a 18% x 5% = 0.9% increase in the effective tax rate - at the median reported value in our sample of $210,000 this is a change in the tax-burden of $1,890.

Figure G2
Event-Study Treatment Variation in Policy Guidance Values

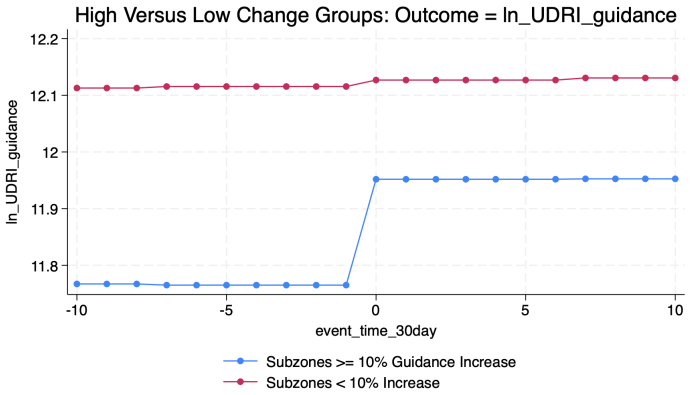


Figure G3 shows how the mean transaction's guidance value changes in response to the increase in legislated guidance values across the two groups.[64] Note that these guidance values are based on actual transactions (as opposed to just the policy guidance values). The results are noisier but we see an approximate 10% increase in the transactions-based guidance value increases.

[64] We construct the mean within each group here by first taking the median guidance value of transactions within a subzone, and then taking the mean of all of those medians within the group. Using the mean within the group is strongly sensitive to outliers.

Figure G3
Event-Study Treatment Variation in Transaction Guidance Values

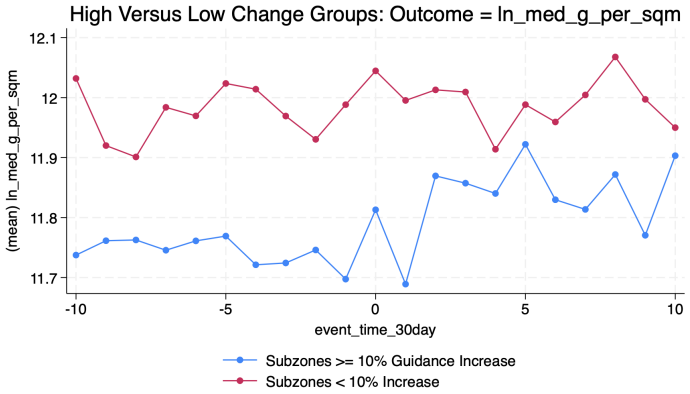
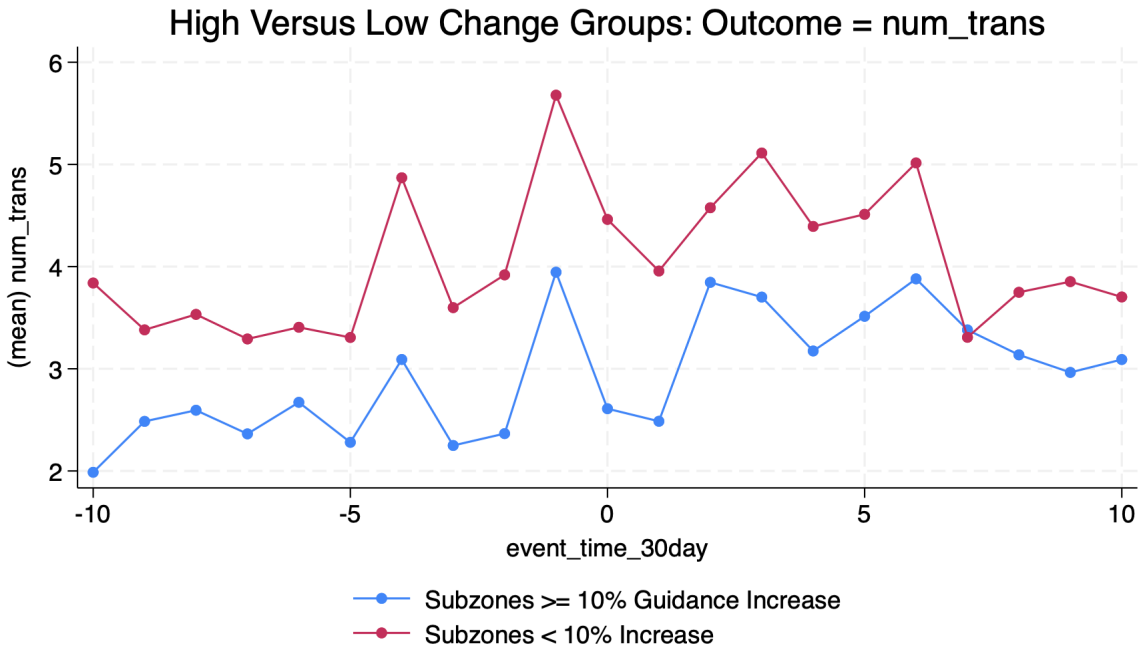


Figure G4 plots the mean number of transactions per subzone in the treatment versus control areas. The trends prior to the increase in guidance values is similar across the regions, suggesting the parallel trends assumption holds. Interestingly, both groups show a spike in transactions just prior to the announcement of the guidance value increases. This suggests that there is a timing-based elasticity. However, after the policy change, there is no discernible reduction in the quantity of transactions in the treatment versus control subzones.

Figure G4
Event-Study Effect of Guidance Value Change on Quantity of Transactions



The difference-in-difference analysis shown so far does not exploit the full distribution of guidance value variation, and also does not exploit the fact that we expect transactions whose reported value is closer to the guidance value to be affected more by changes to guidance values. To exploit this variation more fully, we now present results from a triple-difference design.

Let c and c' be the guidance value per square feet in a subzone that experiences a guidance value increase from c to c'. The intuition for our design is that under a large extensive margin elasticity, we would expect the quantity of transactions with reported values less than c' to decline substantially after the reform, whereas the quantity of transactions with guidance values reported above c' to remain relatively stable – this is because the type of buyer who reports above c' prior to the reform is (at least in theory) unaffected by the increase in guidance values because they are already reporting above the new, higher, guidance value—this type of buyer should continue to buy after the reform as well.

Each subzone has a different c and c', therefore to analyze data from all subzones together we need a standardized measure of reporting behavior relative to c and c'. For transactions prior to the guidance value change, we define this standardized measure

of reporting as $\bar{r} = \frac{r-c}{\hat{c'}-c}$. This measure reflects how much higher the transaction's reported value is relative to the upcoming change in the guidance value. If this ratio is less than 1, it captures the type of transaction that should be affected by the guidance value change; if it's greater than 1 we would not expect the transaction to be affected by the guidance value change. We use the notation $\hat{c'}$ instead of c' because we also utilize some transactions that occur prior to the guidance value change. For these transactions, we do not observe the guidance value the transaction would have had if the transaction had happened after the guidance value change. We estimate c' for these transactions by grossing up their guidance values by the percentage increase in the subzone guidance value.[65]

For transactions occurring after the guidance value change, we define this standardized measure as $\bar{r} = \frac{r-\hat{c}}{c'-\hat{c}}$. For these transactions, we do not observe what their guidance value would have been if the transaction had happened prior to the guidance value change, so we estimate $\hat{c} = \frac{c'}{1+g}$. Note that $\bar{r}$ is only defined for subzones that experience a non-zero change in their guidance value before after the policy change—this effectively limits our sample to transactions from the years 2013-2016.

To begin, we define the "treatment" portion of the reported value distribution as transactions where $\bar{r} < 1$, and the control portion as those with $\bar{r} \geq 1$. Figure G5 shows a histogram of $\bar{r}$ for transactions prior to the policy change (blue) and after the policy change (red). Bunching in the blue histogram occurs at zero, reflecting the many transactions that report $r = c \rightarrow \bar{r} = \frac{r-c}{\hat{c'}-c} = 0$ before the policy change. Bunching in the red post-period histogram is at 1, reflecting that $r = c' \rightarrow \bar{r} = \frac{r-\hat{c}}{c'-\hat{c}} = \frac{c'-\hat{c}}{c'-\hat{c}} = 1$. Qualitatively, under a large negative elasticity we would expect the blue transactions with $\bar{r} < 1$ to not appear in the red distribution, so that the red distribution would be very similar to the blue distribution above 1, and have zero mass below 1. Under a zero elasticity, we expect all of the blue mass with $\bar{r} < 1$ to shift to red mass at 1. Qualitatively, the raw data appears much closer to the low elasticity result, given the large amount of red mass at and above 1.[66]

[65] Let g be the proportionate increase in guidance values, then $\hat{c'} = (1+g)c$.

[66] Note that the distance between 0 and 1 in this graph reflects both small and large guidance value changes – i.e. this figure does not reveal differences across subzones that experienced small versus large guidance value changes.

Figure G5
Illustration of Treatment/Control Regions

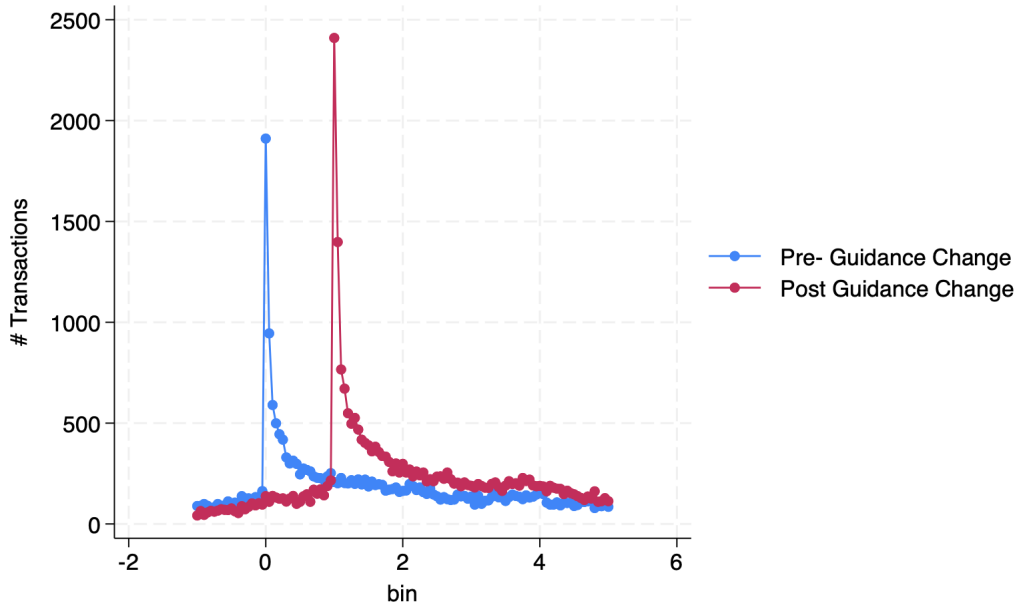


Figure G5 suggests a low elasticity for the number of transaction, however it does not exploit across-subzone variation in the magnitude of the guidance value change. We now test whether subzones with larger guidance value increases experienced a relatively larger decline in their number of transactions within the "treatment" portion of the distribution after the guidance value policy change, as compared to the relative decline in the "control" portion of the distribution after the policy change. We operationalize this with the following triple-difference regression model:

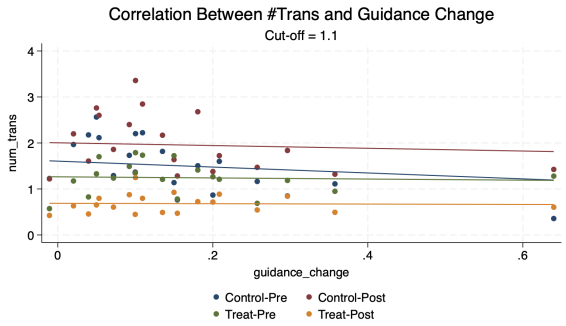
$$y_{ict} = \beta_0 + \beta_1 P_t + \beta_2 T_{ic} + \beta_3 \Delta g_i + \beta_4 P_t * T_{ic} + \beta_5 P * \Delta_i + \beta_6 T * \Delta g_i + \beta_7 P_t * T_{ic} * \Delta g_i + \epsilon_{ict} \tag{18}$$

where y_{ict} is the number of transactions in subzone i, treatment/control portion of the reported value distribution $\bar{r}$, and the event-time month t. P_t (post) is a dummy for observations after the policy-change. T_{ic} is an indicator for the treatment portion of the distribution. Δg_i is a continuous variable measuring the subzone's guidance value change. Our main coefficient of interest is β_7, which tells us how much more the quantity of transactions responds to the increase in guidance values after the policy change, especially in the treatment portion of distribution.[67]

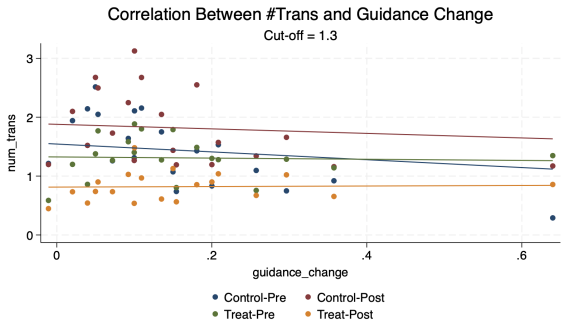
[67] This method of comparing "more" versus "less" treated groups before and after a policy change us-

Figure G6 uses a binscatter formulation to visualize our results. The x-axis is the guidance value changes, and the data is split into four groups: treatment portions of the distribution pre-policy change, treatment portions of the distribution post-policy change, control portions of the distribution pre-policy change, and control portions of the distribution post-policy change. In this figure treatment transactions are defined as those whose reported values are less than 1.1 times the guidance value. Under a negative elasticity, we would expect the slope of the quantity of transactions/guidance value change relationship to be most negative for the treatment portions of the distribution post-policy change. The other lines control for the cross-sectional relationship between guidance value changes and the quantity of transactions prior to the policy change in the treatment group, and such cross-sectional correlations in the pre- and post- periods in the control group. The figure shows that there is essentially no visible change in the correlation between a subzone's guidance value changes before and after the policy chance is implemented in the treatment portion of the distribution; nor is there much change in the control portion of the distribution. This suggests that the guidance value changes did not have a causal effect on the quantity of transactions at the subzone level.

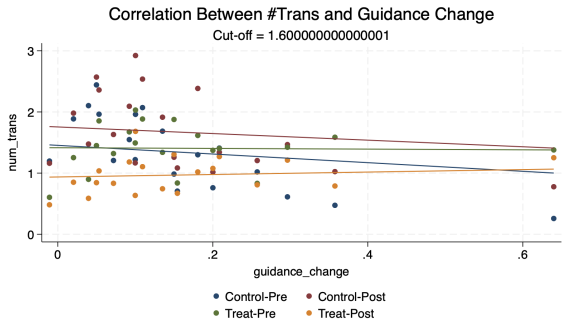
ing panel data follows Feldstein (1995). For a general description of the methodology see Saez, Slemrod and Giertz (2012).



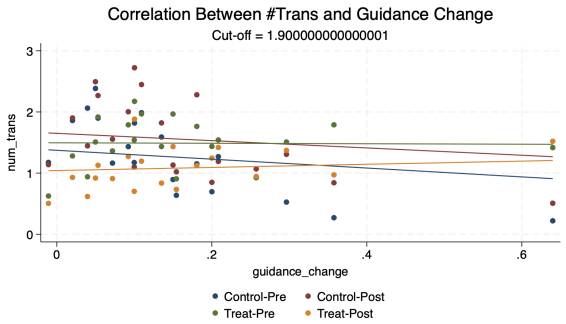
(a) Treatment Cut-off at 1.1 x Guidance Value



(b) Treatment Cut-off at 1.3 x Guidance Value



(c) Treatment Cut-off at 1.6 x Guidance Value



(d) Treatment Cut-off at 1.9 x Guidance Value

Figure G6
Guidance Value Changes and Quantity of Transactions by Treatment/Control and Pre/Post Groups

This figure shows the relationship between guidance value changes experienced in a sub-zone and the quantity of transactions. Within each panel the lines reflect this relationship separately for treatment/control portions of the standardized reported value distribution both before and after the policy change. The panels show results for varying values of cut-offs used to define the treatment/control groups.

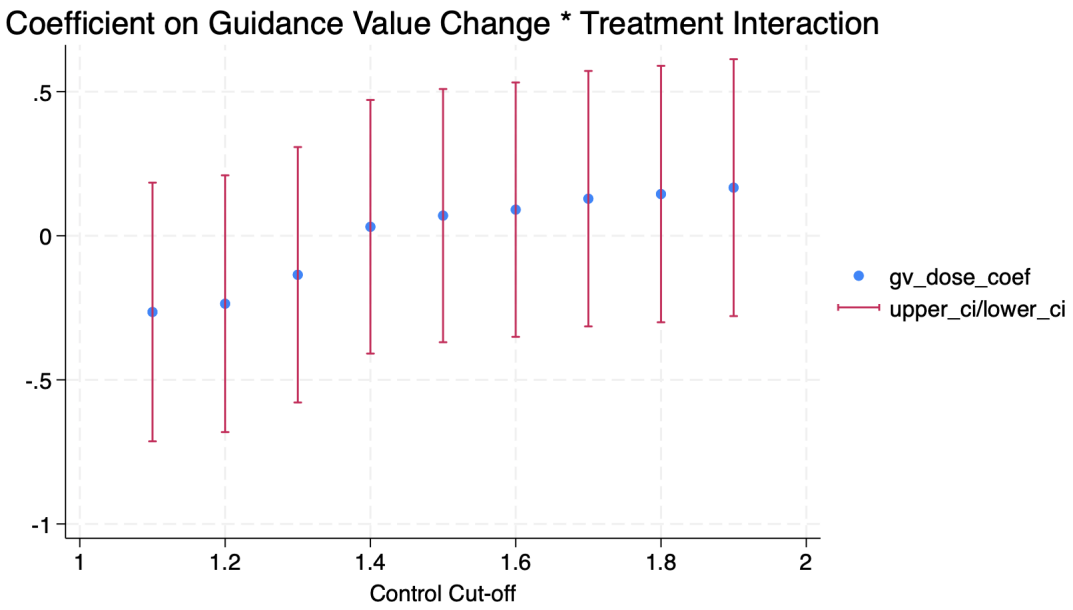
Figure G6, Panal (a), assumes that all transactions that reported less than 1.1 times the post-policy change guidance value are most likely to be affected by the guidance value change. It is possible, however, that even transactions that report above the 1.1 threshold are affected by the policy change. For example, suppose the current guidance value is 100 rupees per square meter, and the updated guidance value is 120 rupees per square meter, and buyers have a rule of thumb to report 20% above the guidance value. Prior to the policy change a buyer would report 120, and after the policy change they would report 144. In this case transactions, reporting at 144 are affected "treated," in that their behavior was affected by the policy change, but using a cut-off of 1.1 x 120 = 132 would not count these transactions as affected.

Panels (b)-(d) of Figure G6 show how the correlation between treatment status and guidance value variation changes as we increase the cut-off used to define the treat-

ment portion of the distribution. The slope of the quantity of transactions to guidance value variation is similar for the treatment portion of the distribution before and after the policy change—suggesting once again a small elasticity of quantity of transactions to changes in guidance values.

To quantitatively assess robustness to the cut-off choice, Figure G7 plots the coefficient on the $P_t * T_{ic} * \Delta g_i$ (β_7) variable as we increase the cut-off definition. Note that this coefficient is equal to the difference in the responsiveness of transactions quantity to the guidance value change between the treatment and control portions of the distribution before versus after the change (i.e., these coefficients summarize the information in the slopes in Figure G6.) Regardless of the cut-off, we cannot reject the null hypothesis that β_7 equals zero. Further, the figure shows that increasing the cut-off leads to even smaller estimates of the responsiveness of the quantity of transactions to the changes in guidance values.

Figure G7
Guidance Value Change Effects for Different Treatment/Control Cutoff Choices



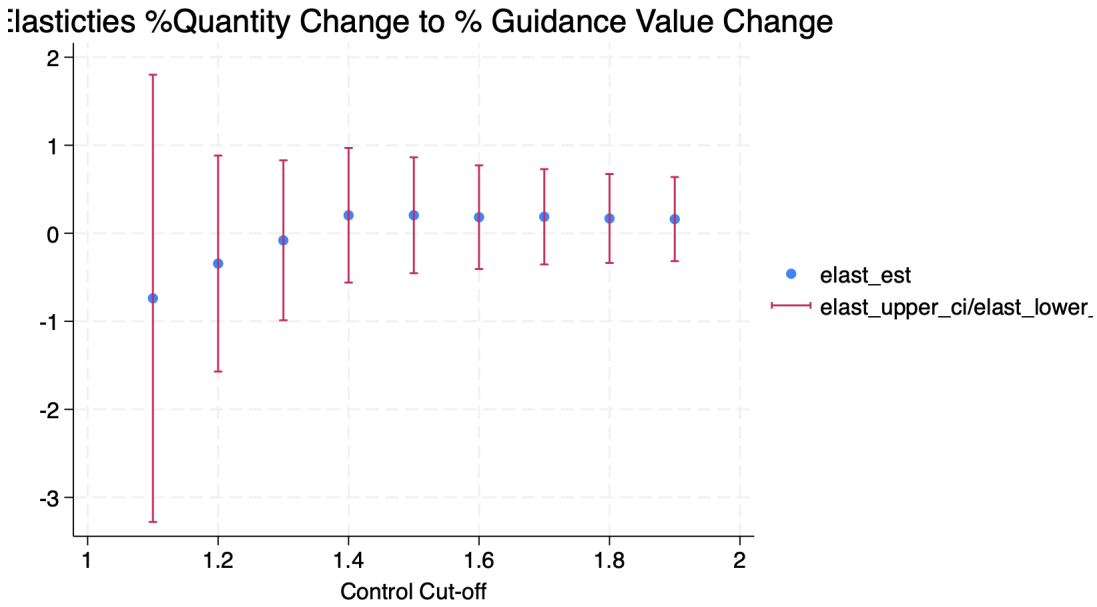
To translate these coefficient estimates into an elasticity estimate, we estimate how many more transactions a 10% increase in guidance values leads to in the treatment group in the post-period, subtracting out the corresponding number for the treatment group in the pre-period, and the relative increase for the control group. In terms of the parameters from our triple-difference model, our elasticity estimate is:

$$\epsilon = \frac{0.1(\beta_3 + \beta_5 + \beta_6 + \beta_7)}{\beta_0 + \beta_1 + \beta_2 + \beta_4)} / (.1) \\ - \frac{0.1 \cdot (\beta_3 + \beta_6)}{\beta_0 + \beta_2} / 0.1 \\ - \frac{0.1 \cdot (\beta_3 + \beta_5)}{\beta_0 + \beta_4} / 0.1 \\ - \frac{\beta_3 \cdot 0.1}{\beta_0} / 0.1 \quad (19)$$

Starting with the simplest term of this equation, the last line corresponds to the control portion of the distribution in the pre-period. $\beta_3 \cdot 0.1$ is our estimate of how much a 10% increase in the guidance values decreases the quantity of transactions for this group. We divide by β_0 to translate this to a percentage decrease in transactions, and then divide by the 0.1 change in the guidance value so we get an elasticity (whose units are percentage change in transactions over percentage change in guidance values). The remaining three lines of this equation indicate the same calculation but for (starting from the top line) the treatment group in the post-period, the treatment group in the pre-period, and the control group in the post-period. The whole equation shows how much larger the elasticity implied from the triple-difference model is for treatment group in the post-period, relative to the treatment group in the pre-period, as well as differencing out any elasticity changes in the control group after the policy change.

Figure G8 plots these elasticity estimates against the cut-offs, along with 95% confidence intervals calculated using the delta method. Once we increase the cut-off level to above 1.4, which seems reasonable given the number of transactions above this level in the red-distribution in Figure G5, we cannot reject an elasticity estimate that is different from zero. Further, the largest elasticity (i.e. the most negative) within the 95% confidence interval is approximately -0.4, suggesting that we can rule out elasticities larger than this based on this triple-difference specification. Overall, from both research designs, we find little evidence to suggest meaningful extensive margin responses to the guidance value changes that have occurred over our sample period.

Figure G8
Elasticity Estimates from Triple-Difference Design



Optimal Tax Policy with Misreporting: Theory, and Evidence from Real Estate

For Online Publication

Santosh Anagol Vimal Balasubramaniam
Benjamin B. Lockwood Tarun Ramadorai Antoine Uettwiller

March 20, 2025

I Data

I.i Transactions Data

Our primary dataset on reported values and guidance values comes from Propstack Analytics and is described in the main text. For our analysis of transactions associated to mortgages, however, we required more detailed information on buyer characteristics to match between transactions data and mortgage. We obtained the underlying transaction documents for this sub-sample analysis for the purpose of matching transactions to mortgages.[68]

The main underlying data source on transactions 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.

[68] This section also describes the underlying documents from which the Propstack Analytics data is ultimately sourced.

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 I1 presents an example of the original document downloaded from the IGR *eSearch* facility. Figure I2 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 a 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[69]. The first row of the table in Figure I2 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

[69] 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.

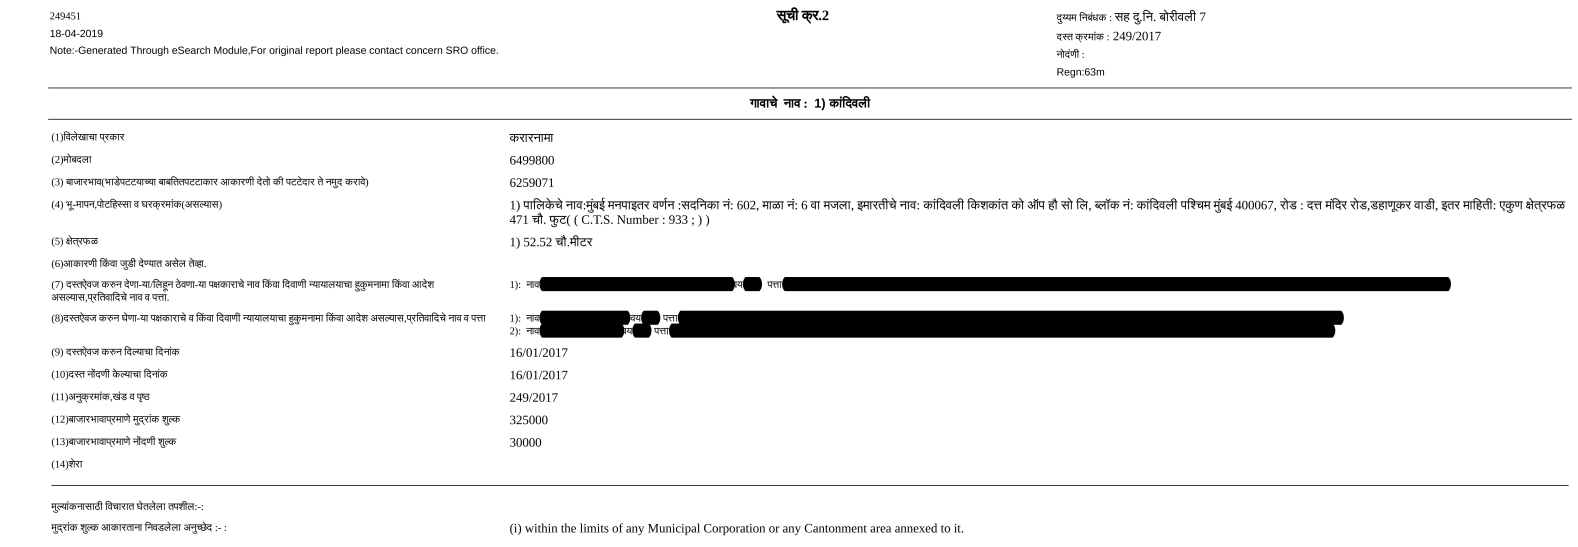
the registrar. The last row provides data on the registration fee paid which is capped at ₹30,000 or 1% of the reported value, whichever is lower.

We validate the coverage of our transaction reports data from the IGR *eSearch* facility by matching the total real-estate stamp duty number of documents filed and revenue generated in each year from our transactions data to the official aggregate numbers. Figure I3 presents this comparison. The top panel of the figure shows the official number of documents filed with the IGR in each month in green against the number of documents in our transaction data (orange for the registrar data we manually sourced for matching to mortgages and blue for Propstack). The time-series are very highly correlated. The bottom panel of Figure I3 shows the official aggregate tax revenue collected from stamp duty in each month against the aggregate stamp revenue from our transaction data. The total revenue figures from both sources include stamp revenues and the registration fees for all transactions in a given month. Once again, the time-series are very highly correlated, especially in the second half of our sample period

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.

Figure I1
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
नोंदणी :
Regn:63m

गावाचे नाव : 1) कांदिवली

(1)विलेखाचा प्रकार	करारनामा
(2)मोबदला	6499800
(3) बाजारभाव(भाडेपटटयाच्या बाबतिपटटाकार 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] पत्ता[redacted]
(8)दस्तऐवज करुन घेणा-या पक्षकाराचे व किंवा दिवाणी न्यायालयाचा हुकुमनामा किंवा आदेश असल्यास,प्रतिवादिचे नाव व पत्ता	1): नाव[redacted] पत्ता[redacted] 2): नाव[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 I2
Translated Transaction Document: An Example

This figure presents the translated version of the original document in Figure I1 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

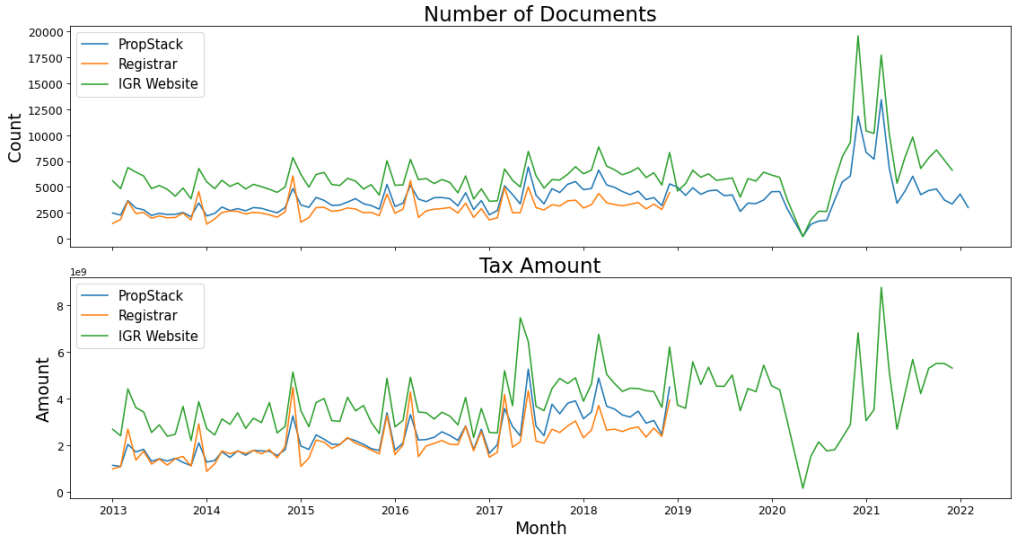
(1) Type of document	Agreement
(2) Reward	6499800
(3) Quotes (The leaseholder leases the details of the rent that the sergeant should specify)	6259071
(4) Land measuring, portals and home number (if any)	1) Name of the corporation: Mumbai Manipatayar Description: House No. 602, Malala No. 6th Floor, 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.

Figure I3
Coverage Validation

This figure reports the total real-estate stamp duty revenue generated in each year from our transactions data to the official aggregate numbers.



I.ii Circle Rate Scheduled Changes

Figure A14 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 I1 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.[70] In the early years of our sample, 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 I4 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

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

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.

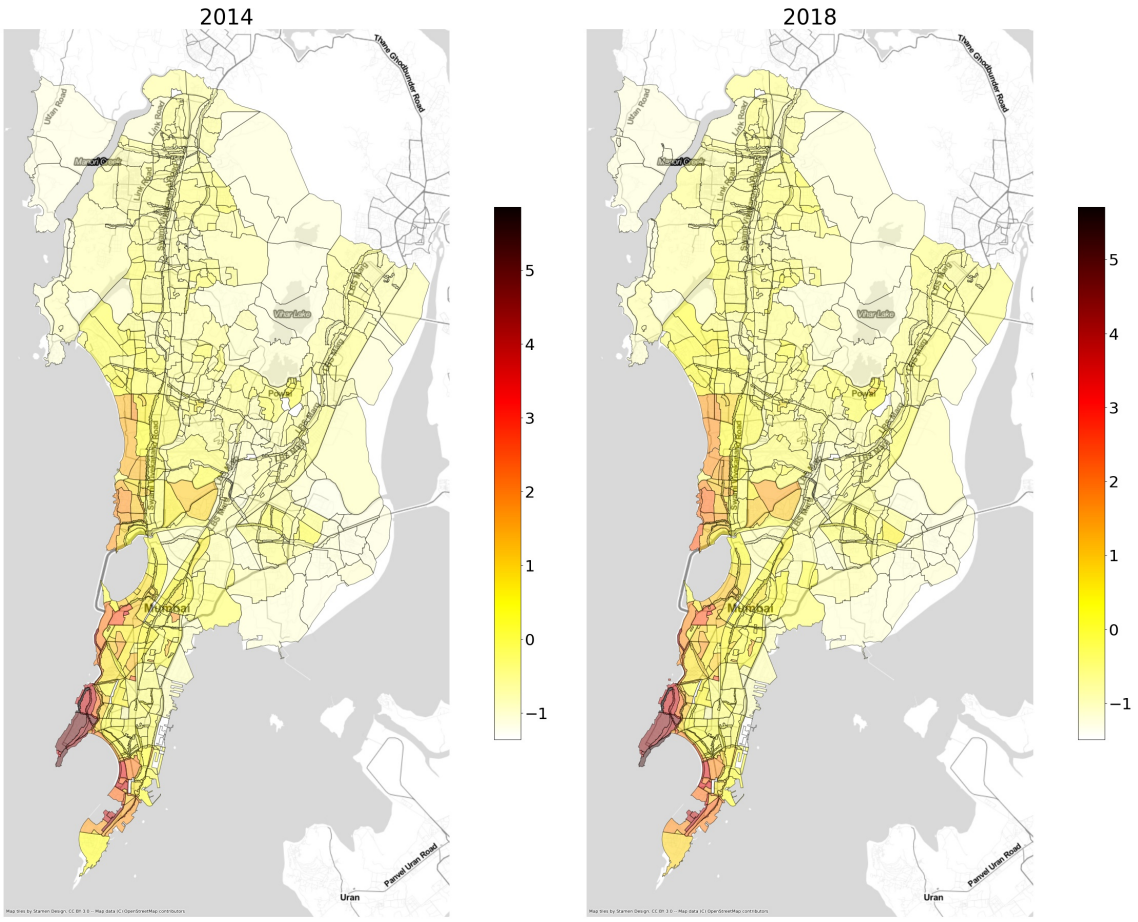
Table I1
Circle Rate - Summary Statistics

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 ($\times$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

Figure I4
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).



II Matching Transactions to Mortgages

As noted above, our mortgage analysis required us to match individual transactions data we obtained separately from the Propstack database.[71] For these transactions data that we independently source, we do not have access to cleanly labeled real estate development names, so we must match them to the estimated market prices using location. Here we describe the matching of our independently sourced transaction

[71] The Propstack database of transactions does not enough identifying information for matching to mortgages.

documents to Propequity estimated market prices.

Our main approach is to match these independently sourced transaction documents to the "CTS" level. A CTS is the smallest administrative geo-spatial unit in Greater Mumbai for the IGR.[72] 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 of our IGR transactions.

Panel A of Appendix Table II2 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.

[72] 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.

Table II2
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

We now describe how we matched these independently source transactions documents to mortgages. We start with the 215,121 independently sourced transactions (see Table II2) and 125,195 mortgage transaction documents in the Mumbai central and Mumbai suburban districts.[73] Table II3 describes our matching procedure. We 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

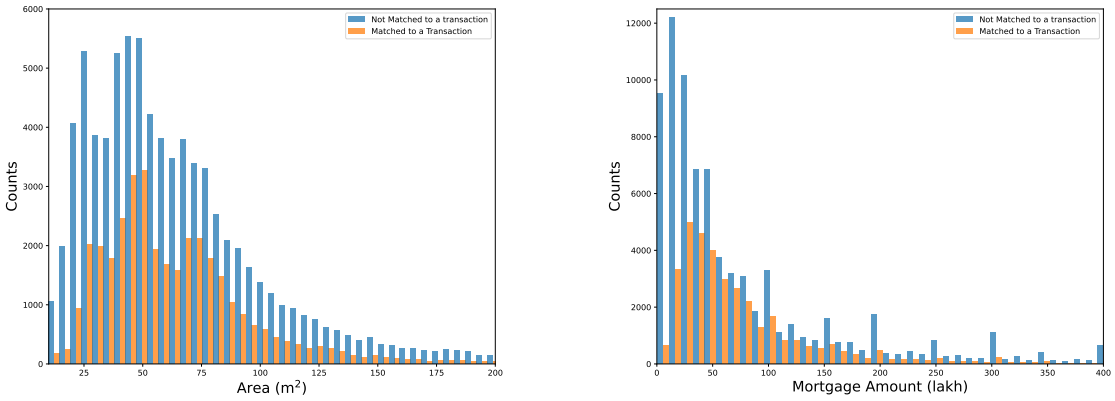
[73] Note for this matching exercise we start with the full number of sales transactions we downloaded for this region before removing observations without a property description and without a CTS.

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. In order to asses whether we are systematically missing underlying transactions, we present a comparison of the distribution of successful and unsuccessfully matched mortgages in Figure II5. We find that the coverage is not disproportionately missing any part of the property size and value distributions, but caveat these results with the fact that there may be measurement error in the loan-to-value ratio of transactions that we were unable to find a mortgage match for.

Figure II5
Mortgage Balance Tests

Panel A (B) represents an histogram of property size (value) for mortgages that either match or do not match to a transaction in the sample.



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 II3
Mortgage Matching Summary Statistics