Essays in Urban and Development Economics

Tanner Regan

Department of Geography and Environment

London School of Economics

A thesis submitted to the Department of Geography and Environment
of the London School of Economics for the Degree of Doctor of Philosophy

London, October 2020

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I declare that my thesis consists of 37,027 words including footnotes but excluding bibliography and appendices.

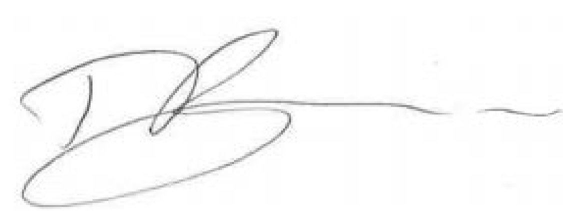

Statement of co-authored work

I certify that Chapter 3 of this thesis is co-authored with Martina Manara and is part of a larger research project. I contributed 75% to the writing of this output in its current form. Chapter 4 of this thesis is co-authored with Julia Bird and Piero Montebruno, and I contributed 75% of the work. Chapter 5 of this thesis is co-authored with Vernon Henderson and Anthony Venables, and I contributed 33% of the work.

Acknowledgements

The work laid out in this thesis benefited greatly from the support of my colleagues, friends, and family.

I was fortunate to have guidance from excellent mentors throughout my studies. In this, Vernon Henderson was exceptional; the feedback I received from him was invaluable, and his unfaltering energy and determination is an admirable model of work ethic. I am also grateful for the supervisory support from Henry Overman, for the learning experience working with Guy Michaels, and for many fruitful discussions with Gharad Bryan. Finally, I thank Matthew Turner for inspiring me to start a career in research.

My research involved travelling to Kenya, Tanzania, and Uganda, and in each of the places I benefited from the support of many people. I am particularly grateful to Priya Manwaring; Chapter 2 would not have been possible without her collaboration. Monica Karangi introduced me to local life in Nairobi, and was also helpful in collecting data. In Dar es Salaam, I am particularly indebted to Charles Mariki for his support in data collection over the years. I also thank Martin and Temu at Ubungo Municipality, and our team of surveyors from Ardhi University, especially Denis Ngole and Robi Kisiyombe. In Kampala, I am grateful for collaboration with the research team at the KCCA led by Henry Kikonyogo, Robert Raikes Mugangiazi, and Daniel Muhumuza Nuweabine. In particular, Jones Ahabwe, Moses Kyazze, Julius Mutebi, and Fred Wabendo. I also gratefully acknowledge the financial support for my research from the International Growth Center.

There were many other interactions that I benefited from. I learned a great deal working with with Tony Venables, Julia Bird, and Ferdinand Rauch. I am thankful to Felipe Carozzi for his support and engagement in the RGUE reading group, as well as, Olmo Silva, Lindsay Relihan, Christian Hilber, and Gabriel Ahlfeldt for joining and sharing their expertise on occasion. I had helpful discussions with Joana Naritomi, Nava Ashraf, Clare Balboni, Maitreesh Ghatak, Niclas Moneke, Erica Pani, Di Song Tang, Justine Knebelman, and Daniel McMillen. I am also grateful for the comments and feedback I received at seminars, workshops, and conferences.

Many friendships were made throughout the course of my PhD. Jamila was there from start to finish, through thick and thin. She has helped me grow as a person and as a researcher. Andreas, with a split banana and a black coffee, gave me reason to wake up and go to campus every morning. Martina taught me how to do fieldwork and made me wake up much earlier than Andreas did. Margarida was my own personal ethics review when IRB just wouldn't cut it. Many fun times were had with Matt, Louise, Raphaela, and Tiernan. I also learned from and was supported by everyone at the SERC office, in particular; Cong, Vivian, CK, Seb, Stephan, Amanda, Ule, Kalle, Simon, Olivier, Caps, Rosa, Giulia, Maria, Neeraj, Brian, and Piero.

Finally, I am grateful to Liz for supporting me through the final years, and encouraging me to finish in a timely manner. Ultimately I am most indebted to the support of my parents. Mom always told me not to make so many assumptions in life, so I learned to channel this into work as an economist. From my Dad, I learned to do a job right the first time around. He also taught me that its easier to cut down something that's too long than extend something that is too short, which surprisingly works as well for writing as it does for lumber.

Abstract

This thesis studies the economics of cities in developing countries. It combines empirical methods using satellite and administrative data, controlled experiments, and primary surveys to study the economics of property development, urban planning, and public finance in developing country cities. The thesis is organised into four independent chapters. The first chapter studies how valuation for the purpose of property tax affects revenues and equity in Kampala, Uganda. I show generally that any assessment procedure that consistently estimates mean value conditional on observables will be regressive. I document that this regressivity is also empirically important in Kampala; on average, the poorest landlords would face an effective tax rate that is 42% higher than the policy rate. The second chapter studies the demand for title deeds in Dar es Salaam, Tanzania. There is significant willingness-to-pay for formal documents that protect property rights. This willingness-to-pay can be predicted at the individual level by traditional local leaders. The third chapter describes the role of slums in Nairobi. While slums are very dense areas with poor-quality buildings and lacking access to key services, measures such as child health and school attendance have caught up or are on pace to catch up with the formal sector. The fourth chapter models the building of a city with informal and formal construction technologies. Using data on individual buildings in Nairobi for 2003 and 2015, a novel set of facts are developed that support assumptions of the model, estimate all parameters of the model, and calculate welfare losses of conversion frictions. For slums in older areas near the centre, even after buying out slumlords, overcoming institutional frictions would yield gains amounting to about $18,000 per slum household, 30 times typical annual slum rent payments.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis consists of four independent chapters on the economics of cities in developing countries.

The first chapter studies how valuation for property tax affects revenues and equity with empirical evidence from Kampala, Uganda. Many governments in developing countries are turning to mass valuation in order to efficiently administer property taxes. Errors in valuation result in effective tax rates at the individual property level that differ from the policy rate. With empirical evidence from Kampala, Uganda I study how methods of valuation impact revenue and equity through effective tax rates. I show generally that any assessment procedure that consistently estimates mean value conditional on observables will be regressive. This helps explain a stylized fact in the property assessment literature that ratios tend to be high for low-priced properties. I document that this regressivity is also empirically important in Kampala; on average, the poorest landlords would face an effective tax rate that is 42% higher than the policy rate. In developing countries compliance is a key concern for tax policy. I estimate the causal effect of an increase in billed liability on compliance, and with this quantify the impact of statistical valuation on revenue. I find that statistical valuation would lower total annual revenues by about 4.2%. This is driven by effective tax breaks for the rich outweighing effective tax hikes for the poor; revenues collected from the bottom eighty percent of landlords by wealth would rise by 6% and those from the top twenty percent of landlords would fall by 10%.

The second chapter (co-authored) studies the demand for title deeds in Dar es Salaam, Tanzania. The inability to recover public project costs discourages much needed infrastructure investments in the developing world. Effective price discrimination could help cover project costs, so we consider whether neighbourhood leaders have information about how local individuals value title deeds. To study this we conduct a lab-in-the-field experiment with 90 local leaders and elicit willingness-to-pay by the Becker-DeGroot-Marschak (BDM) method from 146 property owners. We find that leaders are able to identify the aggregate distribution of demand well, but are less effective at discriminating among plot owners in their neighbourhood. One notable exception is their ability to identify the property owner with the lowest willingness-to-pay. Aggregate demand predictions deteriorate under an environment where the responses of leaders are used to allocate subsidies, but an incentive scheme of cash prizes was able to mitigate this. To keep leaders from misreporting according to personal preferences, an appropriately designed policy will compensate leaders for accuracy.

The third chapter (co-authored) overviews the role of slums in urban Africa, focusing on Nairobi. Spatially disaggregated data show that slum areas are very dense with poor-quality buildings, lacking access to key services such as sewage disposal and electricity. However, improvements to building quality, public-service provision, and socio-economic characteristics are mostly outpacing those seen in the formal sector. Measures such as child health and school attendance have caught up or are on pace to catch up in the near future with the formal sector, while improvements in building quality and service provision are advancing more slowly. We find significant heterogeneity across the city, and in particular that central slums look to be 'stuck' with low-quality buildings and poor service provision, though not with low socio-economic indicators. We explore potential explanations for why slums located on highly prized land near the centre may be stuck with poor infrastructure.

The fourth chapter (co-authored) models the building of a city, estimate parameters of the model, and calculate welfare losses from institutional frictions encountered in changing land-use. We distinguish formal and slum construction technologies; in contrast to slums, formal structures can be built tall, are durable, and nonmalleable. As the city grows areas are initially developed informally, then formally, and then redeveloped periodically. Slums are modelled as a technology choice; however, institutional frictions in land markets may hinder their conversion to formal usage that requires secure property rights. Using unique data on Nairobi for 2003 and 2015, we develop a novel set of facts that support assumptions of the model, estimate all parameters of the model, and calculate welfare losses of conversion frictions. We track the dynamic evolution of the city and compare it with model predictions. In the core city formal sector, about a third of buildings were torn down over 12 years and replaced by buildings on average three times higher. For slums in older areas near the centre, even after buying out slumlords, overcoming institutional frictions would yield gains amounting to about $18,000 per slum household, 30 times typical annual slum rent payments.

Chapter 2

Property valuation for taxation: Vertical equity and revenue efficiency

2.1 Introduction

Property taxes are an important potential source of public revenue, but are underutilized in many developing countries. OECD countries collect property tax revenues worth one or two percent of GDP, while countries in sub-Saharan Africa collect one tenth of this (Moore & Wilson 2017). Faced with limited municipal revenues and rapidly growing populations, cities in developing countries are increasingly turning to taxes on property which can offer a significant source of funding for the provision of local public goods and services.[1] In addition to the importance of property taxes as a revenue source, their consequences for vertical equity are a fundamental consideration since the majority of the rich in Africa continue to pay substantially less income and property taxes than their legal obligations (Kangave et al. 2018).

In these contexts, the tax liability of properties is traditionally determined on a case-by-case basis by trained professional assessors. The administration of property taxes in such a way is costly to scale. Therefore mass valuation, which determines a statistical or 'points-based' prediction of property value based on observable characteristics, is a common approach to

[1] There are many examples of cities expanding their property tax net around sub-Saharan Africa in the past decade; in 2017 Hargeisa, Somaliland expanded their property tax registry from 59,000 to over 200,000 properties (Hass 2017); in 2014 expansion of the property tax net in Kampala, Uganda raised potential revenue from the two central districts by almost threefold (Manwaring & Regan 2020); in Dakar, Senegal only an estimated 20% of plots are on the current tax roll which has motivated the implementation of a modernised property tax management system (Knebelmann 2019); in Freetown, Sierra Leone the city council has recently increased the number of properties on the tax register from 57,000 to about 110,000 (Prichard et al. 2020); in Nairobi, Kenya the valuation roll is only now being updated for the first time since the 1980s (The Standard 2019); in Kananga, DRC a tax campaign registered citizens mostly for the first time (Weigel 2020); and Rwanda recently completed a nationwide program that registered 11.6 million parcels for the first time (Ali et al. 2018).

address the rapid expansion of the tax base (Jibao & Prichard 2016). It is also a common method of valuation employed in developed countries. Mass valuation keeps costs low, but errors in valuation will lead to effective tax rates that differ from the policy rate; for example undervalued properties will face a lower effective tax rate.[2] This effective taxation could have important consequences for equity and overall revenues.

Despite the large push for expanded property tax bases in developing countries, the consequences of mass valuation are largely unstudied. Pomeranz & Vila-Belda (2019) note that 'information is king' when it comes to tax administration, however the public finance literature in developing countries has mostly focused on one aspect of information; improving self declared liabilities through third party reporting (Naritomi 2019, Pomeranz 2015). This paper considers the issue of information in the administrative assessment of liability where the tax payer has little influence. In section 2.2.1 below I expand on the related literature.

This paper makes three contributions. First, I prove that any assessment procedure that consistently estimates mean value conditional on observables will be regressive. This helps explain a stylized fact in the property assessment literature that assessment rates tend to be higher for low-priced properties (McMillen & Singh 2020). Second, I document the extent of assessment regressivity in Kampala, Uganda. While this has long been a topic in the property assessment literature for developed countries, e.g. (McMillen & Singh 2020), (McMillen 2013), (Oates & Fischel 2016), and (Payton 2012), my work is the first to document these effects in a developing country. The developing country context differs in two important ways; first, there are minimal records of property transactions and so market value is largely unobserved, and second compliance rates can be extremely low making the decision to comply a primary margin to consider. Third, I quantify the aggregate and distributional consequences of statistical valuation on compliance and revenue. While this issue has been conceptualised by Zebong et al. (2017) who qualitatively compare alternative valuation models, this paper is, to the best of my knowledge, the first to quantify the effect of statistical valuation on revenues and equity.

I show that statistical valuation will mechanically lead to an effective tax schedule that is regressive, and in the data this is quantitatively important. On average, landlords in the bottom decile of annual income would face an effective tax rate that is 42% higher than the policy rate. Taxing the poor at higher effective rates than the rich is undesirable in

[2] Mass valuation could also be less prone to corruption if property valuers are easily persuaded to undervalue properties. Though, even with statistical valuation, enumerators could also be persuaded to report, for example, a smaller property size thereby undervaluing the property indirectly. This is an interesting question, but outside the scope of this paper.

its own right, but it will also have consequences for the revenue that the government can raise. I compare how revenues change given a regime switch from assessor based valuation to statistical valuation. I show that statistical valuation raises the effective tax rate above the policy rate. Despite this, the effective regressivity leads to a net loss in revenues from current compliers since wealthier landlords are more likely to comply in the first place. In principle this could be counteracted by a behavioural effect, for instance if the response to lower rates encouraged the wealthy to more than proportionally increase their compliance. Therefore, I estimate the effect of the tax liability on compliance. As expected I find that compliance is negatively related to the tax liability, but that this is not strong enough to overcome the direct effect, i.e. effective tax breaks for the rich lead to net losses in revenue. Specifically I find that statistical valuation lowers total annual revenues by about 4.2%, that revenues collected from the bottom eighty percent of landlords by wealth rises by 6%, and that revenues collected from the top twenty percent of landlords falls by 10%. This suggests that there are important consequences for both overall revenue and equity when considering the tradeoffs between assessor valuation and statistical valuation.

The remainder of the paper is outlined as follows. In Section 2.2 I expand on the related literature, and give an overview on the context of property tax in Kampala, Uganda. In Section 2.3 I describe the data that is used throughout the paper and the setting in which properties were statistically valued in Kampala. In Section 2.4 I give a general result about assessment regressivity, and show that this is quantitatively important in Kampala. Next I want to compare how revenues change given a regime switch from assessor based valuation to statistical valuation. In Section 2.5 I estimate the compliance response to a change in log billed valuation, which will be an important parameter in in the regime switch analysis. In Section 2.6 I combine the distribution of effective tax rates with my estimate of the compliance response to calculate the aggregate and distributional changes in revenue due to statistical valuation. Section 2.7 concludes.

2.2 Context

In this section I expand on the literature related to property tax in developing countries, and give context on the setting for the empirical analysis of this paper in Kampala, Uganda.

2.2.1 Literature

This paper is related to a long literature on the bias and distributional consequences of property assessment (Payton 2012, Behrens 1977). This literature has been focused on developed country settings. McMillen & Singh (2020) document a pattern of assessment regressivity in the US, and note that the source of this form of regressivity is unclear with previous work having showed that it is not due to correlation with the success of appeals (McMillen & Weber 2010, McMillen 2013). They argue that standard measures are biased towards a finding of regressivity because even valuations based on unbiased regressions with simulated data necessarily will be regressive. In this paper I make the related, but opposite, point that, rather than just a sign of poor measurement, it is the assessment procedure that necessarily produces regressivity.[3] In another related paper based on property taxation in the US, Avenancio-Leon & Howard (2020) find that the effective tax burden can be 10-13% higher for black and Hispanic residents. They argue that this is mostly due to an 'assessment gap' because property valuations do not capture neighbourhood-level amenities that are negatively correlated with race. To contribute to this literature I provide a fundamental explanation for these empirical findings; that regressivity will naturally result when the assessment objective is to estimate the conditional expectation of property value. I further contribute by providing empirical evidence on the vertical equity consequences of valuation in a developing country setting.

This paper is also related to a growing literature on tax compliance in developing countries (Naritomi 2019, Pomeranz 2015, Del Carpio 2014, Castro & Scartascini 2015, Brockmeyer et al. 2019, Slemrod et al. 2020). Property taxes are important for developing countries. Taxes on labor (and business) income represent a small share of revenues in developing countries in part because the opportunity of informality leads to a high elasticity of taxable income (Gordon & Li 2009). Unlike labour, property is immovable and so easier to keep on formal records. Two recent papers estimate the compliance response to a change in tax liability. Bergeron et al. (2019) provide experimental evidence that increases in the tax rate lower revenues in the DRC, but that revenue-maximizing tax rates are endogenous to the enforcement environment, and Brockmeyer et al. (2020) shows that tax rate increases, enforcement, and late fees raise revenues in Mexico City, but that enforcement and excessive late fees can be bad for welfare relative to their gains in revenue.

Finally this paper is related to a growing literature on the equity of taxation in developing countries. Developing countries are unequal and largely lack progressive tax systems and so it is also important to focus on the equity effects of taxation per se (Alvaredo & Gasparini 2015). In a recent working paper Bachas et al. (2019) find that informality makes

[3] An important note is that, as suggested by McMillen & Singh (2020), both horizontal equity and regressivity within subsets of the distribution could be as important if not more than regressivity across the full distribution.

consumption taxes progressive; by spending a smaller share of their income on informal consumption, rich households in developing countries face a higher effective rate of taxation. This paper is similar in spirit and contributes to this growing literature by quantifying the distributional consequences of property valuation through effective tax rates.

2.2.2 Kampala context

Since its establishment in 2011, the Kampala Capital City Authority (KCCA) has placed a strong emphasis on raising tax revenues (Andema & Haas 2018).[4] A key reform was the introduction of eCitie, an automated system of revenue collection in 2013. The eCitie platform allowed the city to move away from manually recorded tax records to a system that provides up-to-date information on billing and payments to both taxpayers and the city authority. Property tax is one of the most important aspects of the city's finances, contributing to over 30% of own source revenues in 2019.

The KCCA property tax is, by law, a tax on the annual flow of rental income, or 'rateable value', of property in the city (Parliament of Uganda 2005). This type of 'Annual Rental Value System' has widespread use across Africa, being employed in Algeria, Côte d'Ivoire, Egypt, Mali, Mauritania, Morocco, Niger, Sierra Leone, and Uganda (Franzen & McCluskey 2017). Notably, owner-occupied residences are exempt from this tax and so only about 220,000 properties of the 350,000 total are eligible for the property tax.[5] For the remainder of the paper I refer to this set of 220,000 tax-eligible properties unless otherwise noted. The law allows for the rateable value to be determined by two methods. The first is a standard formula that determines the rateable value based on actual records of the annual income of the property. The second is a valuation method based on general features of properties, or particular categories of properties, in a given area. In practice, there are rarely actual income records, and so the vast majority of properties are valued by the second method.[6] Following valuation, the KCCA levies a property tax of 6% on rateable value.

[4] The KCCA acts on behalf of the central government of Uganda, and is the legal entity responsible for the operations of the capital city of Kampala. The KCCA oversees five divisions in the city (Central, Nakawa, Makindye, Kawempe, and Lubaga), which until recently were governed separately. Together they contain about 121 parishes, roughly 940 villages, and over registered 350,000 properties. There are about 1.5m residents in the city, but with a very large daytime population estimated over 4.5 million people (Hass n.d.).

[5] Other exemptions include; Official residences of the president and traditional, and cultural, or religious leaders; Places of public worship; Public outdoor sports or recreation facilities; Cemeteries, burial grounds, and crematoriums; Public charitable and educational institutions supported by endowments or voluntary contributions; Properties of institutions with which the government has contractual obligations not to levy taxes; Properties of organizations that Uganda is obliged to exempt from taxes under international treaties and diplomatic privileges; and Properties owned by local councils.

[6] Within a month of the valuation of properties, all values are publicly posted. It is possible for owners to

To facilitate property taxation through the eCitie platform, the KCCA implemented the CAMV project which expanded the property tax net through an enumeration of all properties in the city collecting data for every property on ownership identity, local characteristics, property-specific attributes, GIS coordinates, and even photographs. Valuation was then conducted by expert valuers who visited localities to determine a suitable range of values for the area, and then used data and photos from the enumeration process to determine the rateable value for each individual property. The Local Governments (Rating) Act (Parliament of Uganda 2005) stipulates that valuation must be repeated every five years, but since it is costly to do in such a way the KCCA is exploring the potential to implement a statistical procedure for estimating rateable value (Manwaring & Regan 2020).

These reforms expanded the tax base, and were an important reason that revenues more than doubled in the city between 2012 and 2015 (Kopanyi 2015). Despite the significant expansion of the property tax net however, the compliance rate is extremely low; 12% of billed properties paid their taxes in the 2019/2020 financial year, resulting in only 34% of potential revenue being raised. Evidence from taxpayer consultations suggests low compliance may be driven by a number of factors, most notably high tax liabilities, limited credibility of enforcement, poor relationship management between KCCA officials and citizens, misinformation, and weak linkages between taxes and service delivery (Ahabwe et al. 2020). Each of these factors are important to understand tax compliance, however in this paper I focus on the role of the billed liability in determining compliance.

While tax payers report that liabilities are too high, properties with larger tax burdens are more likely to be in compliance as can be seen in figure 2.1. Compliance rates for the highest liabilities (2-200mn UGX) range from about 20 to 40%, those between 0.2 to 2mn UGX comply at rates around 12 to 15%, and the bulk of remaining properties comply at a rate around 10%. Obviously this figure confounds the compliance effect of higher liabilities with that of higher property income. The willingness to pay may vary with income of a property due the opportunity cost of enforcement (locking up properties is a commonly cited enforcement method that has higher opportunity cost for high income properties), or simply because high income properties are owned by individuals with a higher willingness and ability to pay. Compliance also varies by property type; conditional on income, residential properties comply at a rate around 10%, and commercial and institutional at 16%.[7] For political reasons, the KCCA is more willing to lock up delinquent properties

protest this valuation by submitting evidence of their actual property income. However, in practice it is very uncommon to have valuations overturned in such a way.

[7] These three types of properties make up the majority. As noted later in Table 2.1 about three quarters of all properties are residential, 18% are commercial, and 5% are institutional.

if they are commercial than if they are residential, which could be one reason why their compliance rate is higher even conditional on income.

2.3 Data

In this section I discuss the Kampala property tax data and describe the setting with respect to statistical valuation in Kampala.

2.3.1 CAMV data

Starting in 2014, the KCCA conducted an enumeration of all properties in the city, called the CAMV (Computer Assisted Mass Valuation) project as described in section 2.2. This project was finally complete for all districts in the city by 2018, and data had been collected for every property in the city. Over 300 characteristics were collected for each property, though many have missing values.

Table 2.1 shows a selection of the property characteristics found in the CAMV data. The first column corresponds to all properties in the city. Properties are most commonly rented for residential use (73% of the time), with commercial (18%) and institutional (5%) making up most of the remainder. Properties are also overwhelmingly in single-storey buildings, with only 8% located in a multi-storey building. There is little heterogeneity in the measures that the KCCA chose to capture building level characteristics. The vast majority of buildings are categorised as 'good' (93%), having brick walls (92%), or concrete floors (96%). On the other hand, properties are spread across areas with varying local characteristics. Most properties are on customary land (51%), with Mailo (33%) and leasehold (15%) making up most of the rest.[8] Most properties are on land categorised as middle class (57%), with a large share also on slum land (38%), and a minority on high class (5%). Public service delivery is also varied; with sanitation split between pit latrines (67%) and sewer or septic tank (33%), a significant share of properties with garbage collection (28%), but most without streetlighting (96%).

The CAMV dataset also has two key characteristics that are measured numerically; built area and billed rateable value of the property. Figure 2.2 plots the distribution of properties by log rateable value and log built area. Both characteristics roughly follow a log-normal distribution, so there is a skew towards higher value and larger properties. The median

[8] Mailo land is a traditional land tenure system in Kampala, where land is held privately by the Bugandan Kingdom. For more on the land tenure system in Kampala see Bird & Venables (2020).

property has a rateable value of 2.53mn UGX and a built area of 49m^2. While 223,260 properties have data on rateable value, only 207,237 have data on built area.

2.3.2 eCitie register, billing, and payment data

Since the CAMV project was only complete in 2018, many properties were first billed using the current system in the 2019/2020 financial year.[9] Therefore, this paper focuses on billing and payments data for that financial year only.

The taxpayer registry includes a unique tax payer number (COIN) for each property owner, along with their contact information. The COIN can be matched to bills, which includes a unique property number associated with the bill. Attached to each bill are records of the property's rateable value, liability, current balance, payments made, penalties, and waivers. Using the unique property identifier, all bills can also be matched to the CAMV database, as well as to the full history of payments made towards the property, including date and amount.

Each property is billed for their annual rates once per year on July 1st, and the deadline for payments is December 31st. However, there is significant confusion among taxpayers as to when payments are actually due (Ahabwe et al. 2020). The confusion is understandable, since the formal deadline for payments is December 31st, but the KCCA has been politically unable to impose any penalties or fines on payments unless they are delayed beyond June 30th. Throughout this paper I use this de facto deadline to define compliance, i.e. properties are compliant if the annual liability on their bill was paid-off before June 30th 2020.

Overall compliance and compliance by property value is discussed above in section 2.2, however it is also important to consider the distribution of compliance with respect to tax payers as they are the relevant decision makers. By matching landlords with their property through the eCitie registry, I create a measure of property wealth that is simply the total rateable value for all properties that an individual owns. Wealthier landlords are more likely to comply, as can be seen in figure 2.3, the top ten percent comply at a rate around 20% and the rest comply at a rate around 5-10%. However figure 2.3 also shows that the wealthiest landlords represent a significant amount of uncollected revenue. In fact, the top ten percent of landlords are delinquent 37bn UGX while the bottom 90 percent combined is only delinquent 29bn UGX. In other words, the relatively high compliance rate of the wealthy is not enough to compensate for their higher liability. This is reflected in tax payer

[9] This project was roughy rolled out by division so that Central division has been billed through the eCitie system since 2017/2018 and Nakawa division since 2018/2019. Most properties in the remaining divisions (Makindye, Lubaga, and Kawempe) were first billed in 2019/2020.

sentiment, as reported in focus groups, that while the rich are perceived as more likely to pay, they are also not thought to be paying their fair share (Ahabwe et al. 2020).

From the registry, there are 106,012 unique landlords so the average landlord rents out just over two properties.[10] The distribution of landlords by property count is skewed, as seen in Figure 2.4 so that about half of all landlords have only one property, one quarter have two, and less than 1% of landlords own more than six properties. About 4,231, or 4% of landlords are identified as legal entities, while the remainder are individuals. Legal entities represent a disproportionate share of the market, owning about 7.8% of the rented properties and 31% of the rateable value in the city. Compliance of legal entities is also higher; conditional on income, properties owned by legal entities comply with the tax at a rate of 23%, while those owned by individuals comply at a rate of 11%.

2.3.3 Statistical valuation in Kampala

Here I outline the setting in which the KCCA statistically valued properties. Some definitions will be useful. I refer to a property i's individual assessment as z_i, and the prediction of a their assessment as $\tilde{z}_i$. The individual assessment was done by a professional valuer with knoweledge of the location, and full details of the property. For the remainder of the paper I treat the assessment as the true value of the property. I discuss the implications of this further in Section 2.C below. I break the population of properties into four relevant samples. First is the *calibration* set of properties; these are 54,810 properties that were billed z_i. The calibration set of properties were used by the government to estimate parameters of a hedonic model based on property characteristics. Second is the *billed* set of properties; these are 14,925 properties that were billed $\tilde{z}_i$. The billed properties were predicted out of sample based on the calibrated hedonic model and do not have data on true property income. Third is the *counterfactual* set of properties; these are 4,065 properties that were billed z_i. The counterfactual properties were also predicted out of sample, and so have data on $\tilde{z}_i$. Finally the remaining 147,183 properties in the city were not used in any part of the statistical valuation. The geographical distribution of properties can be seen in Figure 2.5. In Appendix Section 2.G I discuss the relevance of each of these samples for the external validity of my results.

[10] This is likely to be an overestimate as there are only 82,665 unique phone number contacts listed in the database. This would lead to an overestimate of the number of landlords if landlords are being registered under different COINs across their properties. However, it may also be due to landlords sharing contacts, e.g. the village leader will often be a contact point for a group of landlords, in which case the number of landlords by COIN may be correct. To ensure that my results are not sensitive to this definition of landlords, I redefined landlords by unique mobile phoone numbers instead. This does not make any significant difference to my analysis.

The government assessors were working to complete the valuation of the full city before the beginning of the 2019/2020 financial year. Nearing the deadline, there were technical delays to the data uploads and it was unclear if the remaining properties would be valued on time. In order to complete valuations on time, the remaining properties (billed and counterfactual samples) were predicted using estimates from a hedonic model run on the calibration sample. More specifically, they used a linear model of log assessor value based on log built area and 14 categorical property characteristics.[11] They grouped their estimations into 48 strata by local geographical neighbourhoods and whether the property was on land zoned for commercial, or residential. This allows the hedonic value of each characteristic to vary across local markets, an important consideration as emphasized by Rosen (1974). The breakdown of properties by strata and sample can be seen in Table 2.2. For each strata s they ran regressions of the form

$$ln(z_i^s) = \mathrm{b}^s x_i^s + u_i^s \tag{2.1}$$

where x_i^s is the vector of property characteristics for property i in strata s. So predictions of property value are

$$ln(\tilde{z}_i^s) = \tilde{\mathrm{b}}^s x_i^s \tag{2.2}$$

where $\tilde{\mathrm{b}}^s$ is the vector of OLS estimates of the parameters in model 2.1 for strata s. These predicted values were applied to properties in the billed and counterfactual samples. However, properties from the counterfactual sample were some of the last to be valued, but ended up being downloaded to the system in time to be billed. From author conversations with the valuation team there was no prioritization of which of these properties were downloaded first.

2.4 Effective tax rates

In this section I first provide a general result that, if property valuation is based on the conditional expectation of true value, effective taxation will be regressive. Then, using data from the calibration sample of properties that were both statistically and individually

[11] These were; road access, neighbourhood type (high class, midddle class, and slum), building condition (dilapidated, poor,good, or renovation), sanitation type sanitation (sewer/septic tank, or pit latrine), water connection, power connection, walltype, wall finish, ceiling type, parking type, security type, and fence type. Notably, the sub-categories of these 14 categories were assigned subjective weights and treated as continous variables. For instance, the road access category assigned main road a value of 0.06, a side road value of 0, and a footpath value of -0.05.

assessed, I find that this bias is empirically significant; both properties and landlords face effective tax rates that are on average declining in true property income.

The effective tax rate is the fraction of tax liability in true income of a property, and liability is simply the policy rate times billed valuation. Therefore if the billed valuation is not equal to the true income, then the effective tax rate will differ from the policy rate. For the remainder of the paper I define the effective tax rate of a property i as

$$t_i = \frac{\tau v_i}{z_i} \tag{2.3}$$

where τ is the proportional policy rate, v_i is the billed valuation of property income, and z_i is the true property income. Here I leave the billed valuation v_i to be generic, but throughout the empirical analysis v_i can be either equal to z_i (e.g. for the calibration and counterfactual samples that were billed their true income), or it can be equal to $\tilde{z}_i$ (e.g. in the billed sample that was billed their statistical valuation, or when considering hypothetical scenarios for the calibration and counterfactual samples).

2.4.1 Mechanical bias in effective tax rates

The effective tax rate has two important properties when the liability is based on the conditional expectation of income; first, the mean effective tax rate will be biased upwards relative to the policy rate, and second, the effective tax schedule will be regressive. This is summarized formally as follows:

Proposition 2.1 *If the billed valuation of income is equal to the mean true income conditional on observables; $v_i = E[z_i|x_i]$, and both valuation and incomes are positive; $v_i > 0$ and $z_i > 0$, then:*

(i) The mean effective tax rate will be biased upwards relative to the policy rate, $E[t_i] > \tau$.

(ii) The effective tax rate will be negatively correlated with true income, $Cov(t_i, z_i) < 0$.

The proof for the first part of this proposition follows from the definition of the effective tax rate, the billed valuation being deterministic given observables, Jensen's Inequality giving $E[\frac{1}{z_i}|x_i] > \frac{1}{E[z_i|x_i]}$, and finally the assumption that the billed valuation is equal to the mean true income conditional on observables:

$$\begin{aligned} E[t_i] &= E\left[\frac{\tau v_i}{z_i}\right] = E\left[\tau v_i E\left[\frac{1}{z_i}|x_i\right]\right] \\ &> E\left[\tau v_i \frac{1}{E[z_i|x_i]}\right] = E\left[\tau \frac{E[z_i|x_i]}{E[z_i|x_i]}\right] = \tau. \end{aligned}$$

The proof for the second part of this proposition follows from the definition of covariance, the definition of the effective tax rate, the assumption that the billed valuation is equal to the mean true income conditional on observables, and the result from the first part that $E[t_i] > \tau$:

$$
\begin{aligned}
Cov(t_i, z_i) &= E[t_i z_i] - E[t_i]E[z_i] = E[\frac{\tau v_i}{z_i} z_i] - E[t_i]E[z_i] = \tau E[v_i] - E[t_i]E[z_i] \\
&= \tau E[E[z_i|x_i]] - E[t_i]E[z_i] = E[z_i](\tau - E[t_i]) \\
&< E[z_i](\tau - \tau) = 0
\end{aligned}
$$

The intuition for the first result is that properties with relatively low unobserved value will face disproportionately high effective tax rates. To see why let us consider a cross-section of properties where observables are fixed at $x_i = x$, so that the billed valuation is a constant equal to the mean true income $v_i|x = E[z_i|x]$. It follows that, in this cross-section, the effective tax rate is a constant divided by the true income of the property; $t_i|x = \frac{\tau E[z_i|x]}{z_i}$. Therefore the effective tax rate is a convex function of true property income, and the mean effective tax rate will be larger than the effective tax rate at the mean; $E[t_i|x] > \frac{\tau E[z_i|x]}{E[z_i|x]} = \tau$. Figure 2.6a plots a graphical example of the effective tax rate versus true income in a cross-section where observables are fixed. In this example, true incomes are uniformly distributed along the x-axis.[12] The median property has income $z_i = E[z_i|x_i]$ and so their effective tax rate is equal to the policy rate $\tau \frac{E[z_i|x_i]}{z_i} = \tau$. Since effective tax rates are a convex function of true income, they are skewed towards higher values and so their mean lies above the policy rate on the y-axis. We have shown that for any arbitrary set of observables the effective tax rate will be biased upwards, and therefore integrating over all cross-sections the unconditional mean effective tax rate will be biased upwards.

The intuition for the second result is that since the mean effective rate is high, the only way to get the mean valuation equal to the mean true income (i.e. an unbiased valuation) is to have high value properties face lower effective rates, and low value properties face higher effective rates. To see why consider the a graph of effective tax rates plotted against true income, as shown in Figure 2.6b. In this figure, each point can be represented by the box bounded between it's coordinates and the axes. The covariance can be visualised as the average size of the dashed-grey boxes (example data points) minus the size of the solid-black box (the coordinate of average effective tax rate and true income). Each box will

[12]The uniform distribution was chosen to easily visualise the bias, is not a necessary assumption and this result will hold true regardless of the distribution.

have a width equal to the true value z_i, a height equal to the effective tax rate $\frac{\tau v_i}{z_i}$, and so an area equal to the billed liability $\frac{\tau v_i}{z_i} \times z_i = \tau v_i$. Therefore, the average size of the dashed-grey boxes will be the liability of the average prediction $E[\tau v_i] = \tau E[z_i]$, which can be represented as the rectangle to the left of the solid-black line (base equal to $E[z_i]$) and below the solid-grey line (height equal to τ). The solid-black box will have the same base, but with height $E[t_i]$. Therefore, the sign of covariance can be summarized by the relationship between the mean effective tax rate and the policy rate: if the policy rate lies above the mean effective rate, then the covariance is positive and vice-versa. As we know, due to bias in every cross-section, the mean effective tax rate will always be larger; $E[t_i] > \tau$, and so covariance will always be negative.

To summarize, if billed valuation is based on an unbiased estimate of true property value, then unexplained variation in true income will lead to a biased mean effective tax rate relative to the policy rate. Since the billed valuation is an unbiased estimate of true property value, the effective tax schedule will mechanically be regressive.

Proposition 2.1 relied on the assumption that the billed valuation is equal to the conditional (on observables) expectation of true income; $v_i = E[z_i|x_i]$. This is a valid assumption for a few reasons. First is that this is the specific object that most statistical prediction exercises attempt to estimate. For example, under standard OLS assumptions the condition $v_i = E[z_i|x_i]$ will hold true for large samples.[13] Second is that this is the intuitive aim of most valuation exercises. For example, a points-based system of valuation, that assigns 'expert' determined points to characteristics of properties, attempts to approximate the typical value of a property based on observables. In fact, the official standards of the International Association of Assessing Officers (2017) state that "Appraisers should then compute normal or typical gross incomes, vacancy rates, net incomes, and expense ratios for various homogeneous strata of properties. ... Alternatively, models for estimating gross or net income and expense ratios can be developed by using actual income and expense data from a sample of properties and calibrated by using multiple regression analysis." Therefore, Proposition 2.1 will hold true for many valuation exercises in general. In Appendix Section 2.D I show that my results hold for a standard measure of regressivity

[13] An important caveat to this is when the prediction model is calibrated to predict log property income as it is in many cases including Kampala. In this case the prediction is equal to the mean log true income, and the predicted value itself is a biased estimate of the mean true income because $ln(v_i) = E[ln(z_i)|x_i] < ln(E[z_i|x_i])$ implies that $v_i < E[z_i|x_i]$. However, the results from Proposition 2.1 still hold; for (i) $E[t_i] = E[t\frac{exp(E[ln(z_i)|x_i]}{z_i}] = tE[exp(E[ln(z_i)|x_i] - ln(z_i)] > texp(E[E[ln(z_i)|x_i] - ln(z_i)]]) = texp(0) = t$, and for (ii) $Cov(t_i, z_i) = E[t_i z_i] - E[t_i]E[z_i] = tE[v_i](1 - \frac{E[t_i]}{E[t]}\frac{E[z_i]}{E[v_i]} < 0)$ using $E[t_i] > t$ and $E[z_i] > E[v_i]$ to get the final inequality.

in the assessment profession, the Price Related Differential, and is independent of the statutory assessment rate.[14]

2.4.2 Observed bias and regressivity in effective tax rates

Even though there is a mechanical bias in the effective tax rates, in practice the bias could be inconsequential. Using the calibration sample of properties that were both individually and statistically valued, I can calculate what the effective tax rate would have been for each property if they had been billed with their predicted income rather than the individually assessed income.

The distribution of effective tax rates in the calibration sample of properties is plotted in Figure 2.7. The majority of properties face an effective tax rate that is around the policy rate of 0.06; the effective tax rate is between 0.047 and 0.072 for half of all properties and between 0.02 and 0.18 for 98% of all properties. The median effective tax rate is slightly below the policy rate at 0.058 while the average is higher at 0.065. Therefore the average effective tax rate is about 8% higher than the policy rate.

As suggested above, this bias means that effective tax rates are regressive. In this sample, the correlation between the hypothetical effective tax rates and true income is -0.0734. A related measure of vertical equity used by the IAAO is the Price-Related Differential (PRD) which suggests regressivity if above one and progressivity if below one (International Association of Assessing Officers 2018). The PRD in this sample is 1.19, suggesting regressivity well outside the 0.97-1.03 range advocated by the IAAO. Figure 2.8 plots the relationship between effective tax rates and true property income for the calibration sample of properties. The figure plots the mean, 25th, and 75th percentiles by bins of true property value.[15] It is clear that the mean effective tax rate is negatively related to the true income of a property. For properties earning around 0.1mn UGX the average effective tax rate is 0.18, or about three times higher than the policy rate. On the other end of the distribution, for properties earning about 140mn UGX the average effective tax rate is about 0.03, or about half of the policy rate. We can see that this is not being driven by outliers that skew the means. The grey whiskers plot the 25th and 75th percentiles, and these show that for bins below 1mn UGX at least three quarters of properties are facing effective rates above the policy rate, and for bins above 10mn UGX at least three quarters of properties are

[14] In many jurisdictions assessment is required to be a fraction of market value, rather than 100%. This is referred to as the statutory assessment rate.

[15] Showing statistics by bin improves the visualisation of the graph as opposed to a scatter since a few outliers would otherwise skew the y-axis significantly. Confidence intervals are very tight and so not displayed on the graph in order to improve clarity.

facing effective rates below the policy rate. The whiskers also show that the range of the effective tax rate is also much higher at low values of true property income than at high values. This suggests a second type of vertical inequity; that low earning properties are more likely to face wildly different rates from their peers than are high earning properties. This reflects findings in Chicago where assessment rates tend to be high at very low sales prices, and also much more variable (McMillen 2013).

We are likely concerned with the effective tax rate faced by individuals rather than properties. Since landlords can own multiple properties, the effective rates for them may not be as regressive. Figure 2.9 shows the effective tax rates by city-wide deciles of landlord property wealth. The effective rate for a landlord is the valuation-weighted average effective rate across all of her properties. Here we can see that the poorest decile of landlords would face effective rates that average around 0.085. With the policy rate at 0.06 the poorest landlords are paying an average effective tax rate that is 42% higher than the intended rate. Further, the grey whiskers show that about three quarters of this bottom decile would pay an effective tax rate above the policy rate. Looking at the other end of the distribution, landlords in the top decile would pay an average effective rate that is about equal to the policy rate. However, almost three quarters of these wealthy landlords would pay an effective rate below the policy rate.

2.5 The compliance response to billed valuation

In Section 2.6 I want to compare how revenues change given a regime switch from assessor based valuation to statistical valuation. An important consideration is that changes in liabilities due to statistical valuation will lead landlords to adjust their compliance. Therefore in this section I estimate the causal effect of log billed valuation on compliance which will be an important parameter in my final analysis. In this section I first describe my estimation approach and then give results.

2.5.1 Identification strategy

We are interested in estimating the effect of a percentage change in billed valuation on the probability of compliance. However, billed valuation will be highly correlated with true property income which itself is likely to affect compliance.[16] Therefore I start with

[16]There are good reasons to think that true income has a direct impact on compliance, for instance due the opportunity cost of enforcement (locking up properties is a commonly cited enforcement method that is more costly for high income properties), or simply because preferences may be such that compliance is rising with income (e.g. the wealthier are more likely to pay to avoid embarrassment). There are also good reasons to think that the billed valuation will affect compliance directly. The pecuniary effect of

a simple linear model which considers the effect of both log billed valuation and log true property income on compliance:

$$c_i = \alpha + \beta ln(v_i) + \delta ln(z_i) + \epsilon_i \tag{2.4}$$

Here β is the parameter we are interested in estimating, but as written this model is impossible to estimate; in all cases, properties are either billed their true income (i.e. $v_i = z_i$) and so the two terms are colinear, or the bill differs but I do not observe true income (i.e. in the billed sample when $v_i = \tilde{z}_i \neq z_i$). However, the setting in Kampala does provide a unique opportunity where both the billed and counterfactual samples of properties were statistically valued, but ultimately only the billed sample was charged on this basis. Therefore, to control for the unobserved effect of true income on compliance, I estimate the effect of predicted liability in differences relative to the counterfactual sample. Specifically, I run linear regressions of the general form:

$$c_i = \widehat{\beta}_0 ln(\tilde{z}_i) + \widehat{\beta} ln(\tilde{z}_i)\mathbb{1}(i \in b) + \mathbf{x}_i'\gamma + \widehat{\alpha}_s + \widehat{\epsilon}_i \tag{2.5}$$

Where $\mathbb{1}(i \in b)$ is an indicator equal to one if the property is in the billed sample and equal to zero if it is in the counterfactual sample, $\mathbf{x}_i$ is a vector of property-level controls, and $\widehat{\alpha}_s$ are sample fixed-effects. Controls are included since, as can be seen in Table 2.1, not all observables are balanced across the billed and counterfactual samples. If properties with these characteristics differ in compliance and log predicted value, then they could bias my estimates. In the results section I show that my estimates are robust to the inclusion of a wide set of controls.

The estimate $\widehat{\beta}_0$ captures the common relationship between the log predicted income and compliance, while $\widehat{\beta}$ captures the differential for the billed sample. There are two identifying assumptions. First I assume that the correlation between predicted income $\tilde{z}_i$ and the idiosyncratic propensity to comply ϵ_i is the same in the billed sample as in the counterfactual sample. Second I assume that, conditional on the property level controls, the correlation between predicted income $\tilde{v}_i$ and the true property income z_i is the same in the billed sample as in the counterfactual sample. The second assumption is related to the first since true property income is treated as unobserved. Both are reasonable assumptions

the tax burden is the most obvious and would have a negative effect. Another negative effect is that the perceived unfairness of being over valued may lower morale and discourage tax compliance. The fairness effect may be important, as incorrectly valued properties was a primary concern in focus groups with tax payers (Ahabwe et al. 2020). A potential positive effect of the billed liability is that landlords may believe (correctly) that enforcement is targeted to higher valued properties by tax collectors suggesting a positive effect.[17] All three of these channels could explain the effect of changes in billed valuation, I do not try to differentiate between them.

as the assignment of properties to samples was due to unforeseen time pressure to finalise valuation before the end of the financial year, and so the correlation between predicted income and confounders should be similar in each.[18] Finally, given these assumptions, the estimate is identified with a specific bias. The fundamental issue is that the differencing removes part of the true effect of billed valuation on compliance, because statistical valuation is correlated with billed valuation in the counterfactual sample (not just true income). I can estimate this bias in the counterfactual sample where I know both the statistical valuation and billed valuation. Using this estimate of the bias I correct for it directly. In Appendix Section 2.E I explain the two identifying assumptions and the bias correction in depth.

2.5.2 Results for the compliance response to billed valuation

The main results are in Table 2.3 and are based on the estimating equation given by model 2.5. The sample is the properties that were statistically valued out of sample by the KCCA, i.e. combination of the billed and counterfactual samples. The outcome is always an indicator equal to one if the property complied with the tax. Since landlords can own multiple properties I cluster by landlord ID. The top row reports the main estimate of interest; the differential effect of predicted income when it is billed ($\widehat{\beta}$). The first three columns show results for residential properties only and the last three include all property types.[19]

In columns 1 and 4 the marginal effect of log predicted income is given for the billed sample only, i.e. not in differences. Therefore here we expect the effect of log predicted income to be biased upwards due to it's correlation with unobserved true income. Here we can see that log predicted income has significantly positive effect. This highlights the concern laid out in the section above that the billed valuation is correlated with true property income which itself has a positive impact on compliance. This confirms the intuition that, on the margin, higher income properties are more likely to comply, and the the relationship between compliance and unconditional billed valuation will be confounded by this opposing force.

[18]From author conversations with the valuation team at the KCCA, the statistical valuation was done out of necessity when it was realised that there were technical delays to the data uploads and not all properties would be valued on time. The counterfactual sample is therefore created of properties that were some of the last remaining to be valued, but ended up being downloaded to the system in time to be billed. There was no prioritization of which of these properties were downloaded first.

[19]For specifications with all property types (columns 5 and 6) I include property type dummies and their interaction with log predicted income. This is to control for the imbalance in property types across the billed and counterfactual samples seen in Table 2.1.

Columns 2 and 5 deal with the confounding issue of true property income by differencing the effect of log predicted income of properties that were billed their statistical valuation compared to those that were not. Here we expect the differenced estimate to be negative since higher bills, all else equal, should lead to lower compliance. Once differenced, the compliance response to a change in log predicted income is negative. In the full sample of properties (column 5) this is significantly different from zero, while the residential only sample (column 2) it is borderline non-significant at the ten percent level. For residential properties (column 2), compliance falls by 1.3 percentage points for a unit increase in log predicted income relative to residential properties that were not billed their statistical valuation. For properties in general (column 5) this rate is 1.6 percentage points, though estimates are within one standard error of each other.

Columns 3 and 6 add the same set of property level controls that were used in the prediction model to statistically value the properties. The full set of controls are listed in the table footnote. The coefficients are stable to the addition of these controls, and the controls improve precision for the residential only sample. For residential properties (column 3), compliance falls by 1.6 percentage points for a unit increase in log predicted income relative to residential properties that were not billed their statistical valuation. For properties in general (column 6) this rate is 1.5 percentage points.

Finally, as mentioned in Section 2.5.1 the estimate $\widehat{\beta}$ has a specific bias that can be empirically estimated using the counterfactual sample where both predicted and true income are observed. The fundamental issue is that the differencing removes part of the true effect of billed valuation on compliance, because statistical valuation is correlated with billed valuation in the counterfactual sample. The bias will then depend on this correlation, which I can estimate in the counterfactual sample where I know both the statistical valuation and billed valuation.

From Table 2.3 I take column 3 as my preferred specification because it is estimated with the most precision and the point estimate is similar to that of the other specifications. After correcting for this bias, I get an estimate of -0.0214 with standard error (0.0084) which I consider the causal estimate of the average compliance response to a change in log billed valuation. So for an exogenous 10% increase in billed valuation we expect compliance to fall by about 0.2 percentage points. Since the overall compliance rate in this sample is 5.8% this translates into a 3.4 percent decline in compliance.

2.6 Statistical valuation and revenue

I want to compare how revenues change given a regime switch from assessor based valuation to statistical valuation. There are two broad considerations. First, holding compliance fixed, changes in liability will affect revenues. Second, changes in liabilities will lead landlords to adjust their compliance. Therefore to quantify the effect of the regime switch, in this section I combine the distribution of hypothetical effective tax rates in the calibration sample with my estimate of the compliance response to a change in log billed valuation from Section 2.5.1.[20] First, I conceptualise how this relates to revenue changes. Second, I explain how I use the data to estimate the aggregate and distributional impact on revenue. Finally, I present my results.

2.6.1 Conceptual motivation

This section outlines a simple model to conceptualise how tax revenue would change if properties were billed their statistical valuation rather than their true property income

Each property i is billed a liability based on the policy rate and their valuation τv_i. The revenue raised from a given property is $r_i = \tau v_i \cdot c_i(v_i)$, where $c_i(\cdot)$ is the individual probability of compliance as a function of the property valuation. Switching to bills based on statistical valuation, the change in revenue from an individual property will be:

$$\Delta r_i = \tau v_i c_i(v_i) - \tau z_i c_i(z_i) \tag{2.6}$$

Here we can see that changes in the billed valuation will impact revenue through changes to liability as well as to changes in compliance. Note that I have assumed that switching to statistical valuation has no effect on true income. On the contrary, statistical valuation may encourage landlords to alter their properties to lower the tax burden.[21] This real effect should be negligible in the medium run because valuation is fixed; it took a decade for properties to be revalued in the latest iteration and it is unclear when they will be revalued next. Also due to the extreme rates of delinquency observed empirically, compliance appears to be the more relevant margin to consider.

[20] By hypothetical I mean the effective tax rates that would result if this sample was billed their statistical valuation.

[21] For instance, in historical England, the imposition of a tax on windows led them to be boarded up and new properties built without windows (Oates & Schwab 2015). However, the statistical valuation in Kampala is significantly more complex to avoid than a lump-sum based on a single property feature.

In order to aggregate the change in revenues over properties I make two assumptions. First, I assume that individual compliance as a function of the billed valuation can be approximated with a first order taylor series around the property's true income: $c_i(v_i) = c_i(z_i) + \frac{dc_i(z_i)}{dv_i}[v_i - z_i]$. This is a reasonable assumption given that the majority of hypothetical effective tax rates are small deviations from the policy rate as shown in Figure 2.7 where most properties fall within three percentage points of the policy rate. Given this first assumption, the individual change in revenue from Equation 2.6 can be re-written as:

$$\begin{aligned}
\Delta r_i &= \tau v_i \left[c_i(z_i) + \frac{dc_i(z_i)}{dv_i}[v_i - z_i]\right] - \tau z_i c_i(z_i) \\
&= \tau[v_i - z_i]c_i(z_i) + \tau v_i \frac{dc_i(z_i)}{dv_i}[v_i - z_i] \\
&= \tau[v_i - z_i]c_i(z_i) + \tau[v_i - z_i]\frac{dc_i(z_i)}{dln(v_i)} \qquad (2.7)
\end{aligned}$$

Here we see that the change in effective tax rate for the individual can be decomposed into two effects on revenue. First, holding compliance fixed, the change in billed valuation affects revenue through its direct impact on the liability. Second, faced with a new liability, the probability of compliance will adjust and therefore impact revenue.

The second assumption is that, the compliance response to a change in the log billed valuation is uncorrelated with the difference in valued and true income.[22] This assumption would be violated if, for instance, the marginal compliance response for properties was asymmetric for those that were overvalued compared to those that were undervalued. I explore the implications of this issue in Section 2.F and show that this is a conservative assumption for my estimate of the impact of statistical valuation on revenues. Without data to test this specific assumption, I consider this the pragmatic approach. With this assumption I can aggregate the change in revenue for a given landlord wealth decile:

$$\begin{aligned}
\Delta R_d &= \int_{i \in d} \Delta r_i di \\
&= \int_{i \in d} \tau c_i(z_i)[v_i - z_i]di + \int_{i \in d} \frac{dc_i(z_i)}{dln(v_i)}\tau[v_i - z_i]di \\
&= \int_{i \in d} \tau c_i(z_i)[v_i - z_i]di + \frac{1}{n_d}\int_{i \in d} \frac{dc_i(z_i)}{dln(v_i)}di \cdot \int_{i \in d} \tau[v_i - z_i]di + 0 \\
&\equiv \Delta C_d + \beta_d \cdot \Delta L_d \qquad (2.8)
\end{aligned}$$

The aggregate change in revenue is then composed of a direct effect $\Delta C_d \equiv \int_{i \in d} \tau c_i(z_i)[v_i -$

[22]With this assumption $\int_{i \in d} \frac{dc_i(z_i)}{dln(v_i)}[v_i - z_i]di = \frac{1}{n_d}\int_{i \in d} \frac{dc_i(z_i)}{dln(v_i)}di \cdot \int_{i \in d}[v_i - z_i]di$.

$z_i]di$, which is compliance-weighted aggregate change in liability, and a behavioural effect $\beta_d \cdot \Delta L_d$, where $\beta_d \equiv \frac{1}{n_d} \int_{i \in d} \frac{dc_i(z_i)}{dln(v_i)} di$ is the mean compliance response to a change in the log billed valuation and $\Delta L_d \equiv \int_{i \in d} \tau[v_i - z_i]di$ is the aggregate change in liability.

Finally the aggregate change in revenue for the population is simply the sum across all landlord wealth deciles: $\Delta R = \sum_d \Delta R_d$. These assumptions allow the aggregate change in revenue to be decomposed into terms that can be empirically calculated. In the next section I relate the terms laid out here with their empirical counterparts.

2.6.2 Relating to the empirical approach

Here I derive empirical counterparts of the three terms (ΔC_d, β_d, and ΔL_d) from Equation 2.8 to calculate the change in revenue if all properties were billed their statistical valuation rather than their true property income.

The direct effect of statistical valuation on aggregate revenue, ΔC_d, has a straightforward empirical counterpart. It is the sample total difference between statistically assessed liability and true liability for all compliant landlords in decile d of landlord wealth. This measure is calculated using the calibration sample of properties where I observe both a statistical and true liability of each property.

The behavioural effect of statistical valuation on aggregate revenue is composed of two parts. The aggregate change in liability, ΔL_d, has a straightforward empirical counterpart. It is the sample total difference between statistically assessed liability and true liability for all landlords in decile d of landlord wealth. This measure is calculated using the calibration sample of properties where I observe both a statistical and true liability of each property. The mean compliance response to a change in the log billed valuation, β_d, is more difficult to measure empirically. I use my causal estimate of the average compliance response to a change in log billed valuation, from Section 2.5.1 which was -0.0214. This treats the compliance response as constant across landlord wealth deciles. Again this is a pragmatic approach because I do not have data that measures true landlord wealth in the billed sample, so cannot distinguish their compliance behaviour. However, it is reasonable to think that landlords behave with a consistent response to a percentage change in their liability. To summarize, I calculate the behavioural effect of statistical valuation on aggregate revenue in decile d as the product of the compliance response estimate and the sample aggregate change in liability; $-0.0214 \cdot \Delta L_d$.

2.6.3 Results

This section gives results for the empirical counterparts of the three terms (ΔC_d, β_d, and ΔL_d) from Equation 2.8 to calculate the change in revenue if all properties were billed their statistical valuation rather than their true property income. I calculate these on the calibration sample of properties where I observe both a statistical and true liability of each property.

Results are given in Table 2.4. All units are in millions of UGX. The first row gives the current revenue raised from each decile of landlord wealth. Revenues rise with landlord wealth, e.g. the bottom decile currently contributes 10mn UGX while the top decile contributes 843mn UGX. This is because wealthier landlords face larger bills and are likely to comply (Figure 2.3).

The second row gives the direct effect which is the total difference between statistically assessed liability and true liability for compliant properties. We can see that statistical valuation of properties results in an increase in the revenue contributed by compliers with low property wealth, and a decrease in revenue contributed by compliers with high property wealth. For example, there would be an increase in 9mn UGX being billed to compliers in the bottom ten percent of landlord wealth, and a decrease of 128mn UGX being billed to compliers in the top ten percent. On aggregate, the regressivity of the effective tax rates leads to a loss in revenue from compliers. In the final column we see that the total direct effect is negative; revenue collected from compliers would fall by 121mn UGX. Driving this result is the regressivity of the effective tax rates. Statistical valuation is essentially giving a tax break to compliers who are most likely to comply (Figure 2.3).[23]

While the direct effect is large, the behavioural effect acts as a mitigating force because the compliance response is negative. The third row gives the behavioural effect on revenue which is the product of the average compliance response to a change in log billed valuation (set at -0.0214), and the total difference between statistically assessed liability and true liability for all properties in the calibration sample. At lower levels of landlord wealth the change in revenues from the behavioural response is negative because the average change in liability is positive. At higher levels of landlord wealth this is reversed. Since the total decrease in liability is large for the high wealth landlords, the overall behavioural effect is positive.

[23] If everyone was equally likely to comply, then the direct effect would balance out across the distribution of landlords because the average billed valuation is equal to the average true income, and so their sums are also the same.

The final row is the sum of the second and third rows, and gives the total effect of statistical valuation on revenues. We can see that, on average, revenue contributions will rise from all parts of the distribution except the wealthiest twenty percent. The effective tax break given to the top 20 percent of landlords outweighs the increased revenue from the bottom.

In total revenues will fall by 77m UGX as a result. The total revenue raised for this sample of properties is currently 1.855bn UGX, so statistical valuation lowers annual revenue by about 4.2%. Again, not only does revenue fall, but also the distributional consequences are unfair. The top 20 percent of landlords by wealth currently contribute 1182mn UGX, so under statistical valuation they would contribute about 10% less in revenue. The bottom 80 percent of landlords by wealth currently contribute 672mn UGX, so under statistical valuation they would contribute about 6% more in revenue. This inequality is particularly pronounced for those in the absolute bottom decile of property wealth who would be contributing 80% more in revenue.

2.7 Conclusion

Property tax is a key source of revenue for cities in developing countries that struggle to raise own-source revenues and yet face growing costs of public service delivery. At the same time inequality in these cities is high and efforts to expand the tax base disproportionately bring the poor into the fold. In this paper I show that property tax administration that relies on statistical valuation can have important consequences for revenue and equity. This can improve understanding of the equity and revenue efficiency tradeoffs between different models of assessment.

2.A Figures

Figure 2.1: Property-level compliance by liability

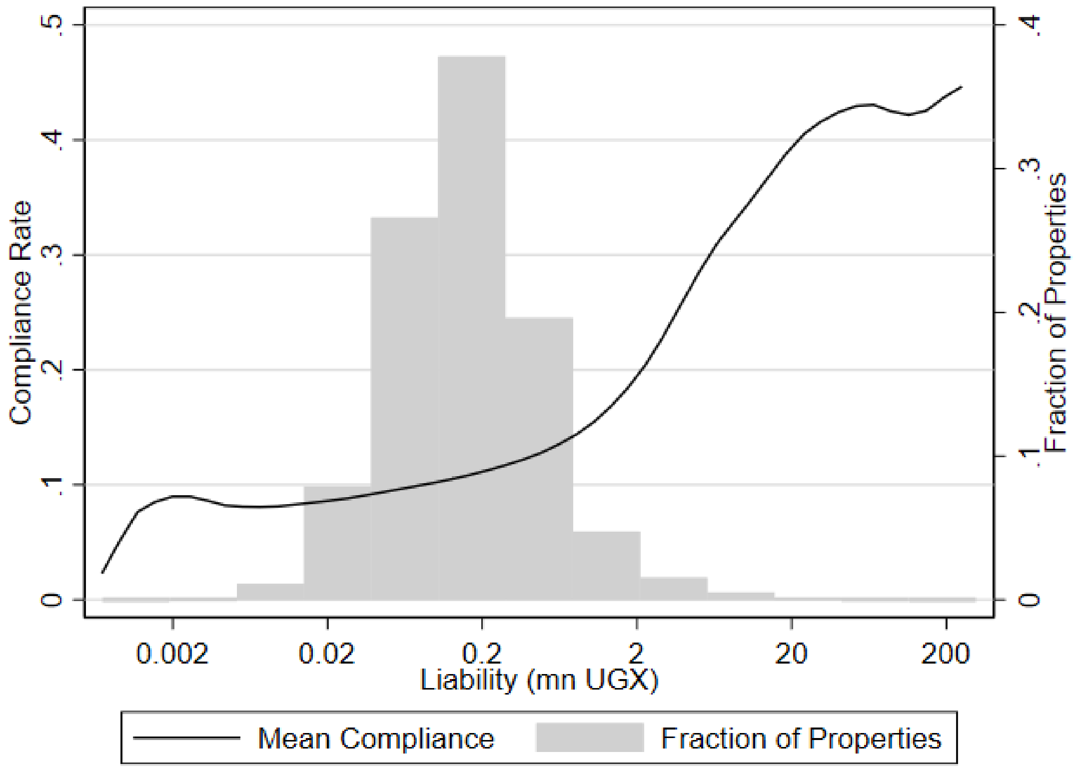

Notes: This figure shows the mean compliance rate by property liability. The compliance rate is estimated with a local polynomial epanechnikov kernel over log liability with bandwidth of one log-point. The fraction of properties for log-point bin is shown by the gray shaded bars in the background. The x-axis is on a log scale and shows the liability in millions of ugandan shillings.

Figure 2.2: Distributions of characteristics from CAMV data

(a) Rateable Value

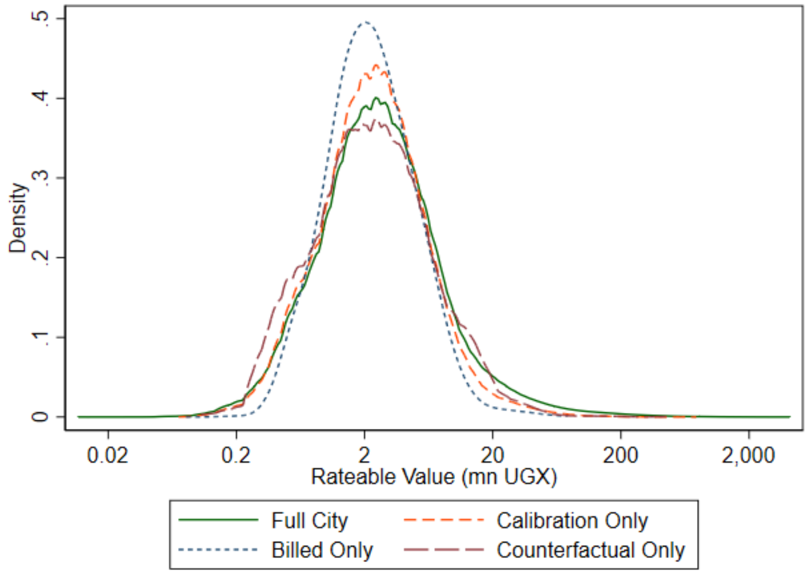

(b) Built Area

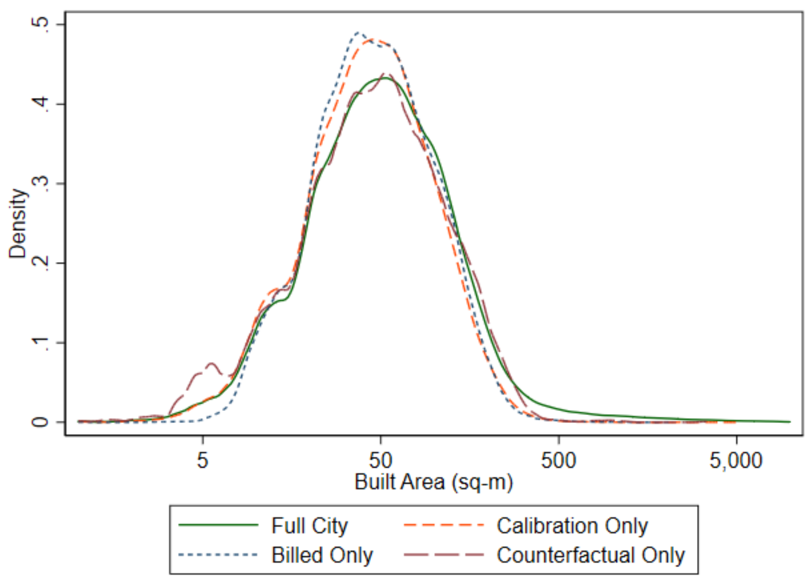

Notes: This Figure plots the distributions of properties by rateable value and by built area for four groups of properties. The first group is of all properties in the city, of which 223,260 have data on rateable value and 207,237 have data on built area. The three other groups are; the calibration sample (properties used to calibrate the KCCA prediction model), the billed sample (properties for which the KCCA billed liabilities based on predicted rateable value), and the counterfactual sample (properties for which the KCCA predicted property value but did not use it for billing). Figure 2.2a plots the distribution of properties by rateable value. The x-axis is on a log scale, and smoothed with and epanechnikov kernel with bandwidth of one-fifth of a log-point. The median property has a rateable value of 2.53mn UGX. Figure 2.2b plots the distribution of properties by built area for three samples. The x-axis is on a log scale, and smoothed with and epanechnikov kernel with bandwidth of one-tenth of a log-point. The median property has a built area of 49m^2.

Figure 2.3: Compliance and Property Wealth

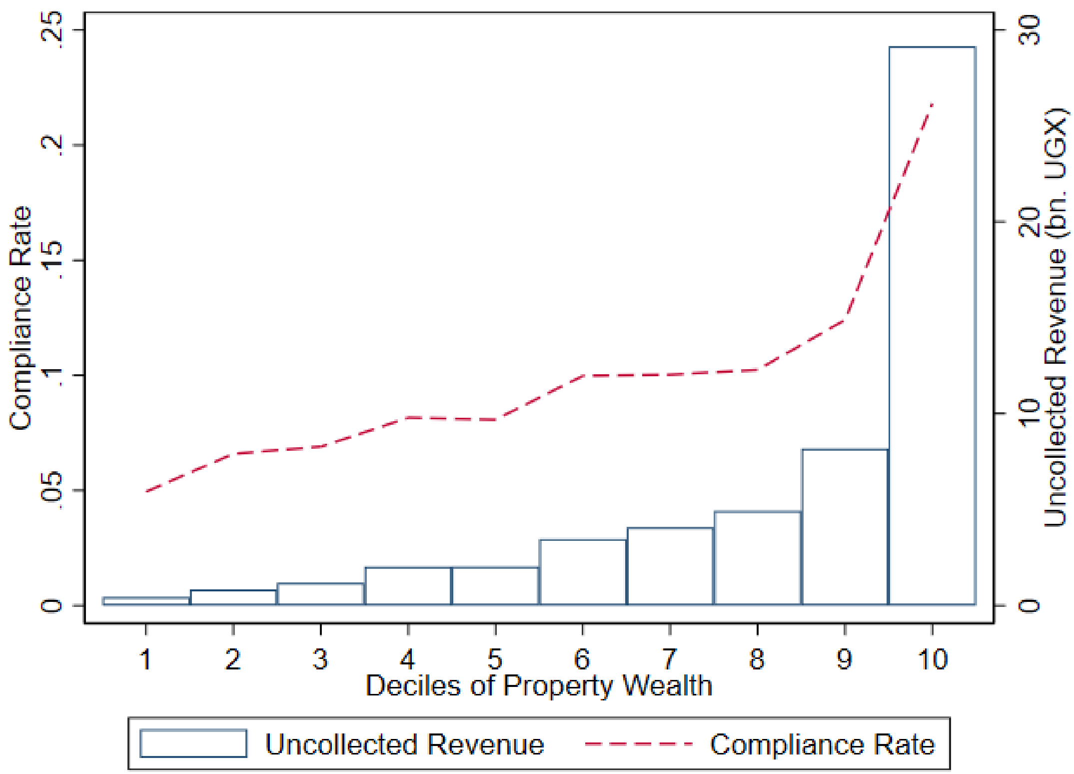

Notes: This figure shows compliance (share of properties that are fully paid this year) and uncollected revenue (total annual payments delinquent this year) by deciles of landlord wealth. Landlord wealth is the sum of rateable values for each landlord.

Figure 2.4: Distribution of landlords

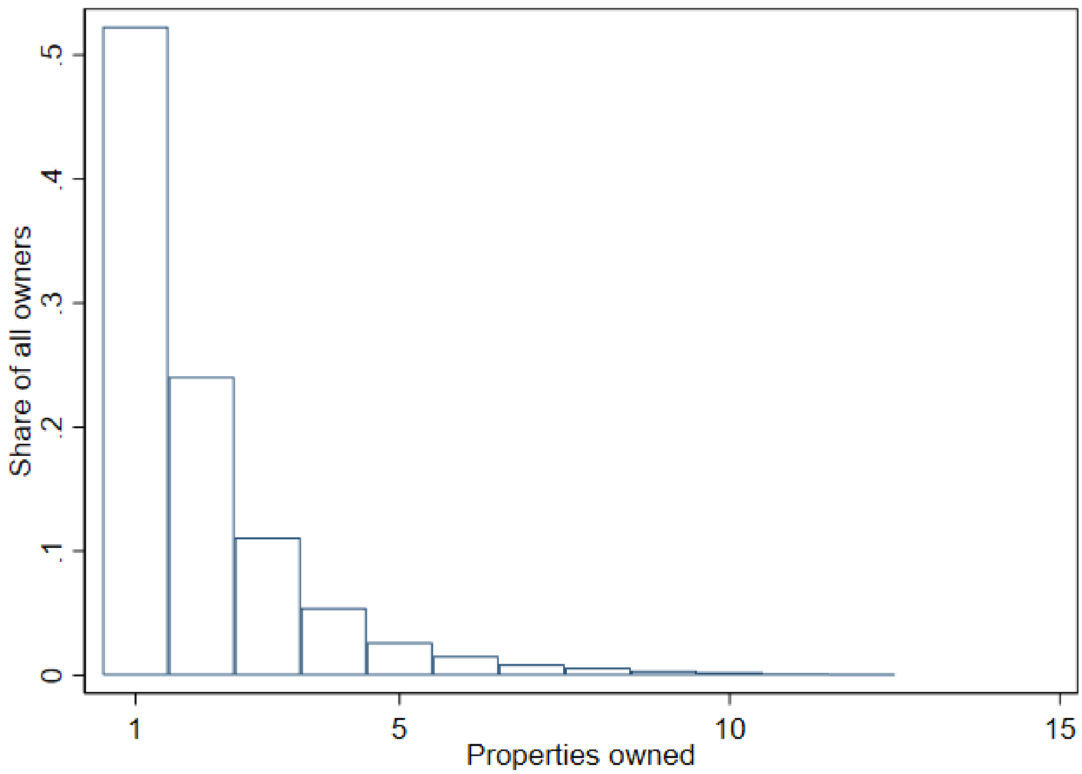

Notes: This figure shows the distribution of landlords by number of properties owned. The right tail of the graph is truncated; 500 landlords not shown here own more than 15 properties.

Figure 2.5: Map of properties

(a) All properties

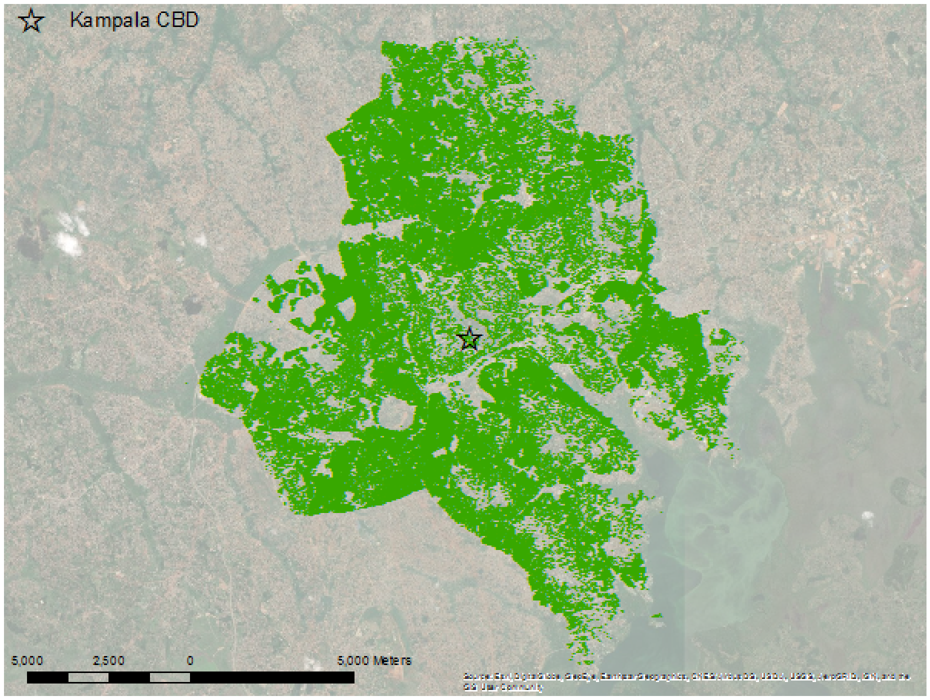

(b) Properties by sample

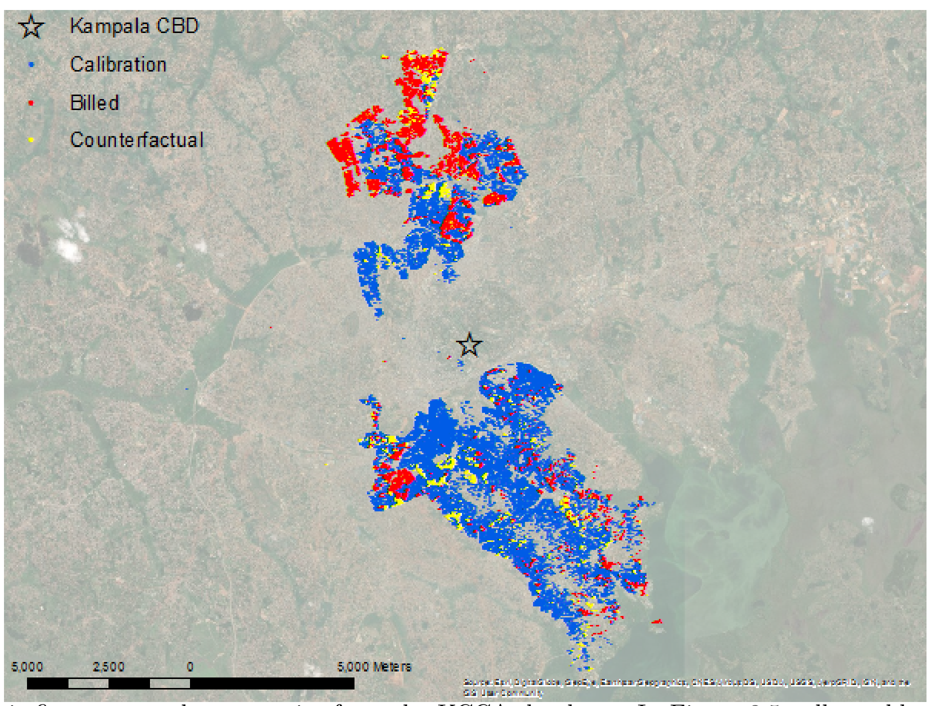

Notes: This figure maps the properties from the KCCA database. In Figure 2.5a all taxable properties are mapped as green dots. In Figure 2.5b only those properties that were use as part of the statistical valuation exercise are mapped. The calibration sample is shown in blue, the billed sample in red, and the counterfactual sample in yellow. In both figures the city centre is shown with a star and a scale is given in the bottom left corner.

Figure 2.6: Graphical examples for Proposition 2.1

(a) The mean effective tax rate is biased upwards for any cross-section with fixed observables

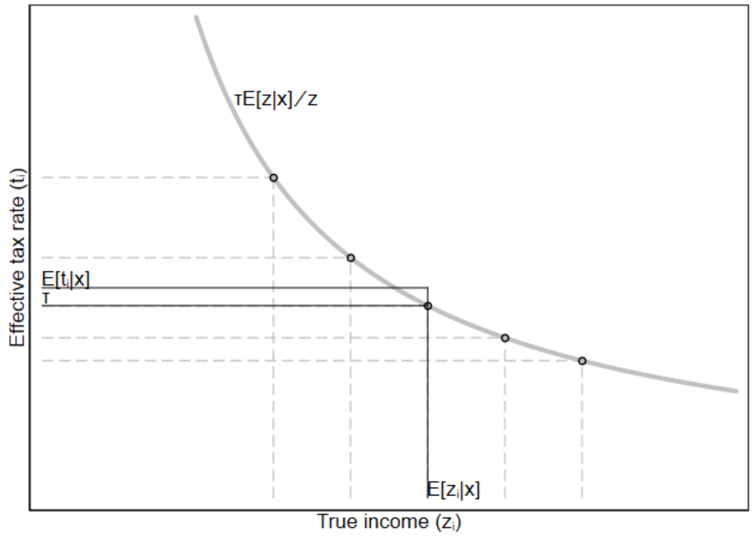

(b) Graphical example for negative covariance between the effective tax rate and true income

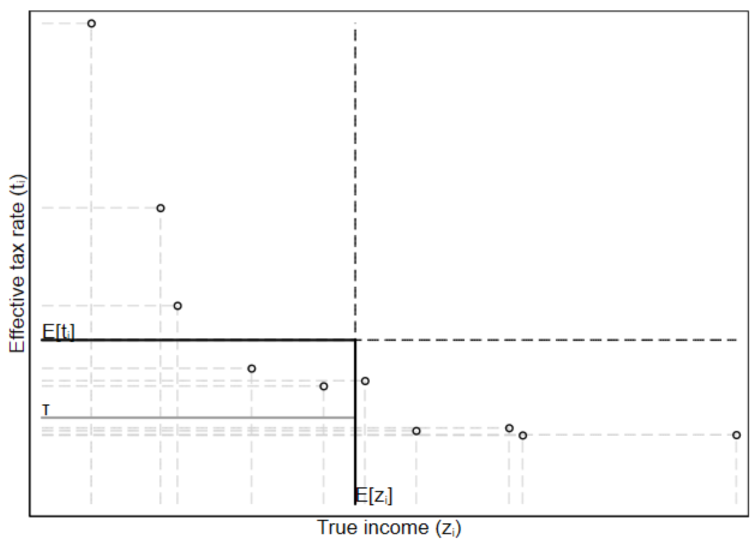

Notes: This figure gives graphical intuition for Proposition 2.1. Figure 2.6a plots the effective tax rate vs. true income in a cross-section where observables are fixed for an example distribution of true income. True incomes are uniformly distributed, so their mean lies at the median along the x-axis. Since effective tax rates are a convex function of true income, they are skewed towards higher values, and so their mean lies above the median. Figure 2.6b plots the effective tax rate vs. true income in the full sample, for an example distribution of true income. By definition, covariance will be negative/positive if the average area of the rectangles below and to the left of each point (dashed gray lines) is less/more than the area of the rectangle below and to the left of the the means along each axis (solid black lines). Since each dashed grey rectangle has an area equal to predicted liability; $\frac{\tau v_i}{z_i} \times z_i = \tau v_i$*, then their average will be* $\tau E[v_i] = \tau E[z_i]$*. So their average area can be represented as the rectangle to the left of the solid black line and below the solid grey line* $\tau \times E[z_i]$*. The rectangle inside the black lines will have the same base, but with height* $E[t_i]$*. Therefore the black rectangle will always be larger, and covariance will always be negative.*

Figure 2.7: Distribution of Effective Tax Rates

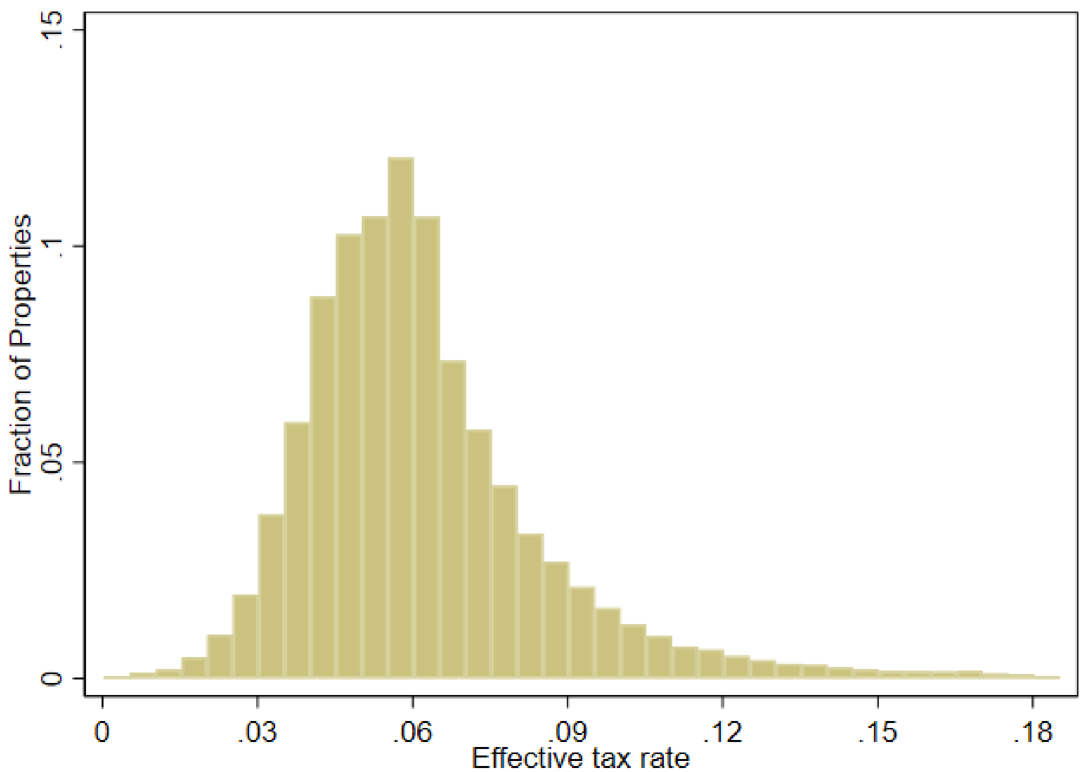

Notes: This figure shows the distribution of the effective tax rate for the calibration sample of properties that were valued both statistically and individually. The fraction of properties in each half-percentage point bin is given on the y-axis. The right tail of the graph is truncated; properties in the top one percent are not shown. The majority of properties not shown have rates below 0.5. The highest effective rate is 7.7 and only six of the properties not shown have effective rates above 1.

Figure 2.8: Effective Tax Rate by True Property Income

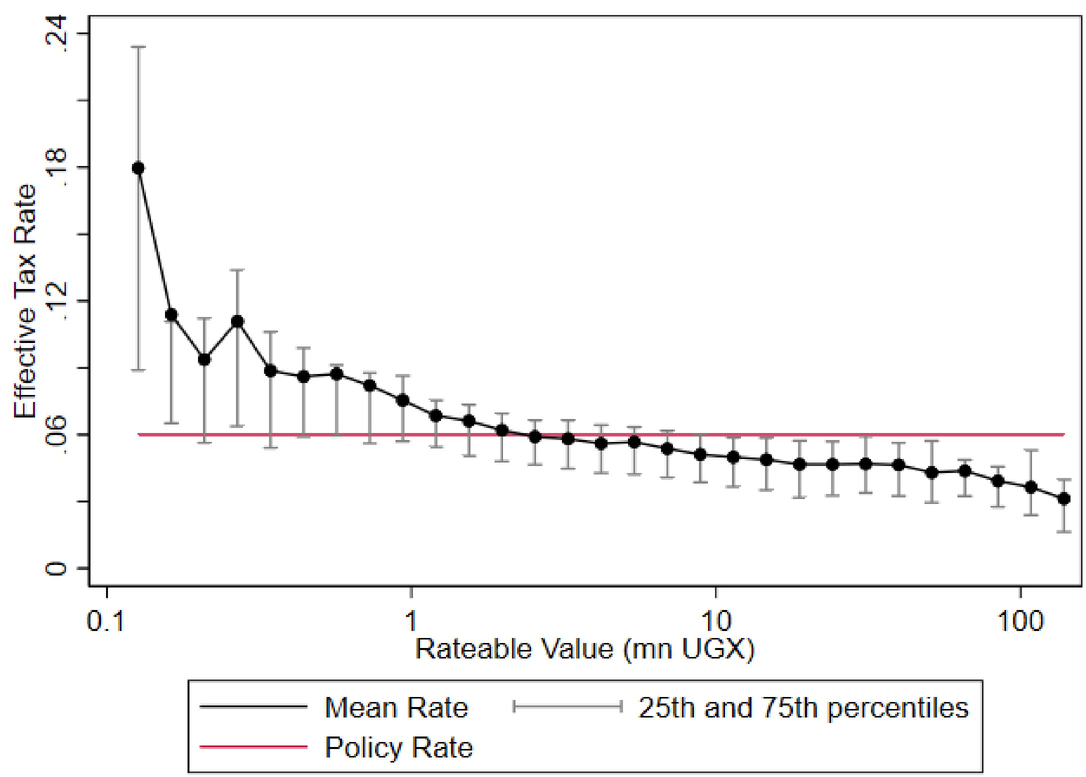

Notes: This figure shows the mean, 25^{th}, and 75th percentiles of the effective tax rate by property liability bins for the calibration sample of properties that were valued both statistically and individually. The x-axis is on a log scale and shows the liability in millions of ugandan shillings. Bins are taken for ever quarter-log-point in rateable value, and bins with less than 15 observations are discarded.

Figure 2.9: Effective Tax Rate by City-Wide Landlord Wealth

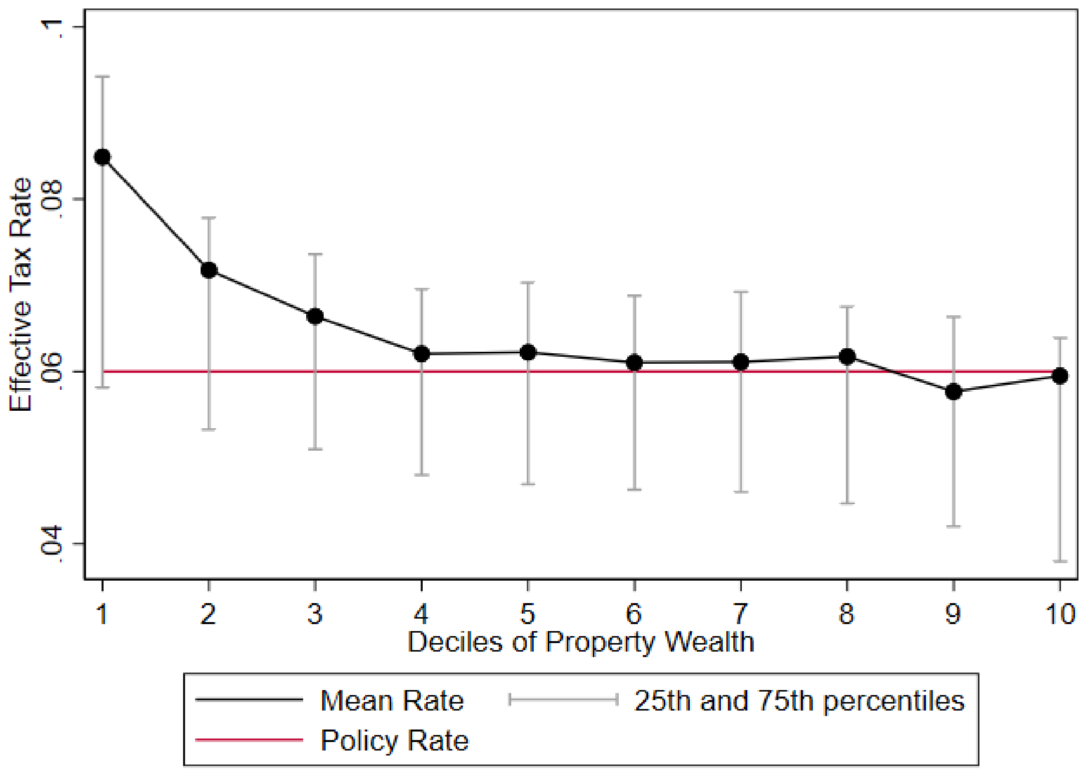

Notes: This figure shows the mean, 25^{th}, and 75th percentiles of the effective tax rate by city-wide deciles of landlord wealth for the calibration sample of properties that were valued both statistically and individually. Deciles are calculated across the entire city, rather than just the calibration sample, so that landlords who own properties across samples have all of their income accounted for. The effective rate for a landlord is the liability-weighted average across all of her properties.

2.B Tables

Table 2.1: CAMV Summary statistics

			Samples		
		Full City	Calibration	Billed	Counterfactual
Property Type		223,298	54,810	14,884	4,039
	commercial	0.18	0.14	0.03	0.18
	institutional	0.05	0.00	0.01	0.01
	residential	0.73	0.86	0.96	0.81
Building Type		220,809	54,810	14,884	4,039
	single	0.92	0.96	0.98	0.95
	storied	0.08	0.04	0.02	0.05
Building Condition		220,833	54,810	14,884	4,039
	dilapidated	0.01	0.01	0.01	0.01
	good	0.93	0.93	0.94	0.95
	poor	0.05	0.06	0.05	0.04
	underconstruction	0.01	0.00	0.00	0.00
Walls		218,631	54,638	14,835	4,025
	brick/concrete	0.92	0.93	0.95	0.90
	cladding	0.01	0.00	0.00	0.00
	iron	0.02	0.02	0.01	0.06
	mud	0.01	0.02	0.01	0.01
	timber	0.03	0.02	0.03	0.04
Floors		218,195	54,776	14,880	4,038
	concrete	0.96	0.97	0.99	0.97
	earth	0.03	0.02	0.01	0.03
Tenure		220,983	54,810	14,884	4,039
	freehold	0.02	0.00	0.02	0.00
	customary	0.51	0.58	0.41	0.39
	leasehold	0.15	0.14	0.08	0.11
	mailo	0.33	0.28	0.49	0.49
Neighbourhood		215,110	54,810	14,884	4,039
	high class	0.05	0.02	0.02	0.03
	middle class	0.57	0.54	0.65	0.66
	slum	0.38	0.43	0.33	0.30
Sanitation		218,143	54,810	14,884	4,039
	pit latrine	0.67	0.73	0.71	0.71
	sewer/septic tank	0.33	0.27	0.29	0.29
Garbage Collection		220,983	54,810	14,884	4,039
	no	0.72	0.82	0.85	0.82
	yes	0.28	0.18	0.15	0.18
Street Lighting		220,983	54,810	14,884	4,039
	no	0.96	0.99	1.00	1.00
	yes	0.04	0.01	0.00	0.00

Notes: this table summarizes a selection of the property characteristics found in the CAMV data. Each panel corresponds with a categorical characteristic, the first row of a panel gives the number of non-missing property records for the characteristic, and the following rows give the share of properties associated with each category. The first column is the sample of all properties in the city, the second column is the sample of properties that the KCCA used to calibrate their rateable value billed model, the third column is the sample of properties that were billed based on the KCCA predicted property values, and the fourth column is the sample of properties that the KCCA predicted property value for but were billed based on their true income.

Table 2.2: Number of properties by KCCA calibration strata

		Number of properties by sample		
Division	Strata Name	Calibration	Billed	Counterfactual
Makindye	Buziga Commercial	273	54	22
Makindye	Buziga Residential	1668	223	102
Makindye	Ggaba Commercial	547	0	33
Makindye	Ggaba Residential	2172	741	278
Makindye	Kabalagala Commercial	384	1	7
Makindye	Kabalagala Residential	1146	24	28
Makindye	Kansanga Muyenga Commercial	903	149	52
Makindye	Kansanga Muyenga Residential	2938	286	97
Makindye	Katwe 1 Commercial	341	0	4
Makindye	Katwe 1 Residential	379	0	0
Makindye	Katwe 2 Commercial	243	0	0
Makindye	Katwe 2 Residential	1032	1	0
Makindye	Kibuli Commercial	354	8	6
Makindye	Kibuli Residential	1753	244	157
Makindye	Kibuye 1 Commercial	153	0	0
Makindye	Kibuye 1 Residential	1904	18	0
Makindye	Kibuye 2 Commercial	480	2	1
Makindye	Kibuye 2 Residential	3723	2	0
Makindye	Kisugu Commercial	409	6	1
Makindye	Kisugu Residential	1502	136	21
Makindye	Lukuli Commercial	312	2	1
Makindye	Lukuli Residential	0	21	14
Makindye	Luwafu Commercial	312	0	8
Makindye	Luwafu Residential	1752	56	235
Makindye	Makindye 1 Commercial	265	21	57
Makindye	Makindye 1 Residential	1701	58	429
Makindye	Makindye 2 Commercial	397	69	240
Makindye	Makindye 2 Residential	1375	13	38
Makindye	Nsambya Central Commercial	326	0	1
Makindye	Nsambya Central Residential	3426	0	4
Makindye	Salaama Commercial	842	28	36
Makindye	Salaama Residential	3583	137	181
Makindye	Wabigalo Commercial	413	12	1
Makindye	Wabigalo Residential	1147	32	0
Kawempe	Kawempe 1	2840	3559	436
Kawempe	Kikaaya	2273	1593	67
Kawempe	Kyebando	3485	5155	422
Kawempe	Makerere 1	587	0	38
Kawempe	Makerere 2	1684	0	9
Kawempe	Mpererwe	198	800	340
Kawempe	Mulago 1	232	0	4
Kawempe	Mulago 2	999	0	22
Kawempe	Mulago 3	1378	0	52
Rubaga	Najjanankumbi 1	847	816	172
Rubaga	Najjanankumbi 2	560	441	14
Rubaga	Ndeeba Commercial	813	62	380
Rubaga	Ndeeba Residential	759	155	55
	Total	54810	14925	4065

Notes: This table counts the number of properties involced in statistical valuation by sample (calibration, billed, and counterfactual) and by strata. There are 48 different strata, and the KCCA ran separate hedonic models for each one.

Table 2.3: Compliance response to billed valuation

	(1)	(2)	(3)	(4)	(5)	(6)
Log Predicted Income	0.0133**	-0.0128	-0.0162*	0.0114**	-0.0166*	-0.0147+
$\times \mathbb{1}(\text{Billed} = 1)$	(0.00298)	(0.00791)	(0.00782)	(0.00301)	(0.00736)	(0.00779)
Log Predicted Income		0.0260**	-0.0483		0.0158	-0.0655+
		(0.00735)	(0.0341)		(0.0119)	(0.0377)
$\mathbb{1}(\text{Billed} = 1)$		0.168	0.220+		0.225*	0.189+
		(0.115)	(0.114)		(0.107)	(0.114)
Billed Only	✓			✓		
Residential Only	✓	✓	✓			
Property-Type F.E.				✓	✓	✓
Prop. F.E. × Log Prd. Inc.					✓	✓
Prediction Model Controls			✓			✓
R^2	0.002	0.005	0.037	0.007	0.013	0.050
Observations	14288	17566	17566	14884	18909	18467
Landlord Clusters	7479	9175	9175	7685	9635	9500

Notes: Standard errors in parentheses are clustered at the landlord level, p-value thresholds are denoted for the test that the estimate is equal to zero ($^+$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$). This table reports estimates of the marginal compliance response to a change in log billed valuation. The outcome is always an indicator equal to one if the property complied with the tax. The top row reports the main estimate of interest; the differential effect of predicted income when it is billed ($\widehat{\beta}$). The first row shows the differential effect of log predicted income for the group of properties that were billed based on their predicted income. Columns 1 and 4 restrict the sample to properties that were billed based on predicted income. Columns 1-3 restrict the sample to residential only, while columns 4-6 include all property types (residential, commercial, and institutional) and include property-type fixed effects. Columns 5 and 6 in addition add property-type effects, i.e. interactions between property type fixed effects and log predicted income. Columns 3 and 6 include controls based on the KCCA Prediction Model, i.e. all property characteristics used in to predict property income; Log built area, road access, neighbourhood type (high class, midddle class, and slum)l, building condition (dilapidated, poor, good, or renovation), sanitation type sanitation (sewer/septic tank, or pit latrine), watter connection, power connection, wall type, wall finish, ceiling type, parking type, security type, fence type, and a fixed effect for each calibration model.

Table 2.4: Statistical Valuation Effect on Revenue

	Deciles of Property Wealth										
	1	2	3	4	5	6	7	8	9	10	Total
Baseline Revenue	10	22	37	64	69	127	153	190	340	843	1855
Direct Effect (ΔC_d)	9	4	5	3	3	7	1	-2	-23	-128	-121
Behavioural Effect ($\beta_d \cdot \Delta L_d$)	-1.0	-0.5	0.0	0.7	0.8	2.1	2.9	4.4	8.8	25.5	44
Total Effect	8	4	5	4	4	9	4	2	-15	-102	-77

Notes: This table shows the statistical valuation effect on revenue by decile of property wealth decomposed into it's component parts. Units are millions of UGX. The first row gives the current revenue raised from each decile of landlord wealth. The second row gives the direct effect which is the total difference between statistically assessed liability and true liability for compliant properties in the calibration sample. The third row gives the behavioural effect on revenue which is the product the marginal compliance response to a change in log billed valuation (set at -0.0214), and the total difference between statistically assessed liability and true liability for all properties in the calibration sample. The final row is the sum of the second and third rows, and gives the total effect of statistical valuation on revenues.

2.C Assessors valuation as measure of true income

Throughout the paper I treat the assessor's value interchangeably with the true property income. Obviously the individual assessment is also measured with error. In fact, as long as the assessors follow the international standard to "... compute normal or typical gross incomes... for various homogeneous strata of properties" then even their assessments will tend to be regressive as shown in Proposition 2.1 (International Association of Assessing Officers 2017).

The first reason for this choice is that the assessor's value of a property is the predicted object in the statistical valuation procedure. In this sense, the assessors value of the property income is the intended tax base in the valuation exercise. It is therefore important to see how errors in valuation relative to this standard affect equity and revenues. A second reason for this choice is that the measurement error in the assessor's value is unlikely to be correlated with observables in such a way that it overturns the regressivity of the statistical assessment. For instance, the assesor would need to be significantly biased in overvaluing properties with high value. This is despite the fact, that as I have shown, even the assessors valuation itself will tend to be regressive if it is unbiased. In the figure below, I run monte carlo simulations of the measurement error in assessors valuation. Specifically, I assume that the log assessors value is measured with classical measurement error, i.e. $z_i = z_i^* + u_i$. I consider thirty normal distributions with mean zero and standard deviation as a fraction of observed standard deviation in log assessors value. So that the x-axis on the graph is the ratio of the standard deviation of measurement error, to the standard deviation in log assessors income. I take 500 random draws of u_i from each normal distribution, and plot the mean correlation between the 'true' effective tax rate and 'true' income. As the graph shows, for $var(u) < var(z)$ the correlation between the effective tax rate and true income will be even more negative. Then for the case of classical measurement error, my estimate of regressivity can be seen as conservative.

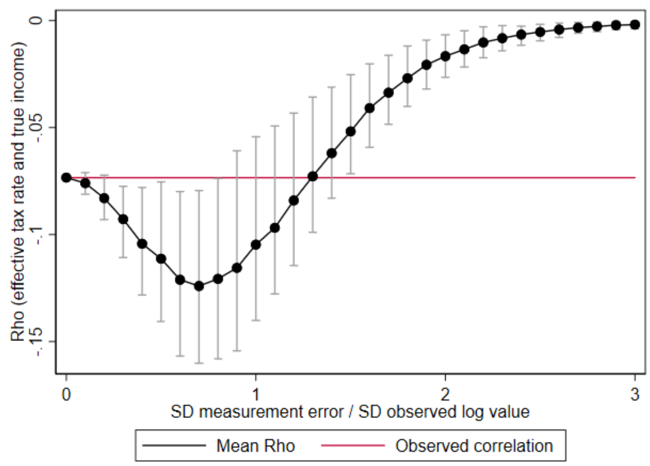

Finally true property income can never be known with certainty, and so without evidence of any specific bias using the assessor's value is a pragmatic approach.

2.D Alternate measure of regressivity

In Section 2.4 I showed that valuation based on the mean property value conditional on observables will be regressive. To do so I proved that the covariance of the effective tax rate and true income must be negative. There are other measures of the regressivity of a full distibution, the standard being the Price Related Differential (PRD) (International Association of Assessing Officers 2017). For properties assessed at value a_i with market price p_i the PRD is defined as the average assessement ratio $\frac{a_i}{p_i}$ divided by ratio of the averages $\frac{a}{p}$. The International Association of Assessing Officers (2017) suggests that the PRD fall in the range of 0.97 to 1.03, where values above one are regressive and values below one are progressive. The literature typically finds evidence of regressivity using these measures (McMillen & Singh 2020). Under the same assumptions outlined in Proposition 2.1, we can show that the PRD must be greater than 1.

Let us suppose that the statutory assessment rate is α, (e.g. by law assessments target $\alpha\%$ of market value). A reasonable assessment procedure takes all observable characteristics and assigns the expected average market value conditional on observables (e.g. this is the objective of OLS). Therefore the objective of assessment is to predict the mean price conditional on observables weighted by the statutory rate:

$$a_i = \alpha E[p_i|x_i]$$

Where a_i is the assessment, p_i is the market value, and x_i are observable (to the assessor) characteristics of property i.

The PRD is then

$$PRD = \frac{1}{n}\sum_i \frac{a_i}{p_i} / \frac{\sum_i a_i}{\sum_i p_i}$$

In large samples we get convergence to the population moments:

$$PRD = E[\frac{a_i}{p_i}] / \frac{E[a_i]}{E[p_i]}$$

The denominator will be α, so

$$PRD = \frac{1}{\alpha} E[\frac{a_i}{p_i}]$$

Now since the assessment is deterministic with respect to observables

$$\begin{aligned} PRD &= \frac{1}{\alpha} E[a_i E[\frac{1}{p_i}|x_i]] \\ &> \frac{1}{\alpha} E[a_i \frac{1}{E[p_i|x_i]}] \\ &= \frac{1}{\alpha} E[\alpha E[p_i|x_i] \frac{1}{E[p_i|x_i]}] \\ &= 1 \end{aligned}$$

Where the inequality comes from Jensen's inequality and the fact that all assessments are positive.

2.E Identification and bias correction explained

In section 2.5.1 there were two identifying assumptions. First, is that the correlation of predicted income and idiosyncratic propensity to comply divided by the variance of the KCCA estimate is equal across both samples. More specifically we require that $\frac{cov(ln(\widehat{v_p}),\epsilon_p)}{var(ln(\widehat{v_p}))} = \frac{cov(ln(\widehat{v_c}),\epsilon_c)}{var(ln(\widehat{v_c}))}$, where subscript p denotes observations in the billed sample and subscript c denotes observations in the counterfactual sample. The second assumption is that the correlation of predicted and true property income is equal across both samples. More specifically we require that $\frac{cov(ln(\widehat{v_p}),ln(z_p))}{var(ln(\widehat{v_p}))} = \frac{cov(ln(\widehat{v_c}),ln(z_c))}{var(ln(\widehat{v_c}))}$. Finally, the estimate is identified with a specific bias that I correct for empirically, i.e. $\widehat{\beta}_p \to \beta[1 - \frac{cov(ln(\widehat{v_c}),ln(z_c))}{var(ln(\widehat{v_c}))}]$.

To simplify the analysis, let us treat Equation 2.5 as the difference between the simple regressions from each sample an so $\widehat{\beta} = \widehat{\beta}_p - \widehat{\beta}_c$.[24]

Starting with the billed sample we can run the simple regression of compliance on log predicted income. Since the billed value and predicted income are equal in this sample, i.e. $\tilde{z}_p = v_p$, the estimate can be written as:

$$\begin{aligned}
\widehat{\beta}_p &= \frac{cov(ln(\tilde{z}_p), c_p)}{var(ln(\tilde{z}_p))} \\
&= \frac{cov(ln(\tilde{z}_p), \beta ln(v_p) + \delta ln(z_p) + \epsilon_p)}{var(ln(\tilde{z}_p))} \\
&= \beta \frac{cov(ln(\tilde{z}_p), ln(v_p))}{var(ln(\tilde{z}_p))} + \delta \frac{cov(ln(\tilde{z}_p), ln(z_p))}{var(ln(\tilde{z}_p))} + \frac{cov(ln(\tilde{z}_p), \epsilon_p)}{var(ln(\tilde{z}_p))} \\
&= \beta + \delta \frac{cov(ln(\tilde{z}_p), ln(z_p))}{var(ln(\tilde{z}_p))} + \frac{cov(ln(\tilde{z}_p), \epsilon_p)}{var(ln(\tilde{z}_p))}
\end{aligned}$$

Moving to the counterfactual sample we run the same simple regression of compliance on log predicted income. Since billed value and true income are equal in this sample, i.e. $v_c = z_c$, the estimate can be written as:

$$\begin{aligned}
\widehat{\beta}_c &= \frac{cov(ln(\tilde{z}_c), c_c)}{var(ln(\tilde{z}_c))} \\
&= \frac{cov(ln(\tilde{z}_c), \beta ln(v_c) + \delta ln(z_c) + \epsilon_c)}{var(ln(\tilde{z}_c))} \\
&= \beta \frac{cov(ln(\tilde{z}_c), ln(v_c))}{var(ln(\tilde{z}_c))} + \delta \frac{cov(ln(\tilde{z}_c), ln(z_c))}{var(ln(\tilde{z}_c))} + \frac{cov(ln(\tilde{z}_c), \epsilon_c)}{var(ln(\tilde{z}_c))} \\
&= (\beta + \delta) \frac{cov(ln(\tilde{z}_c), ln(z_c))}{var(ln(\tilde{z}_c))} + \frac{cov(ln(\tilde{z}_c), \epsilon_c)}{var(ln(\tilde{z}_c))}
\end{aligned}$$

Now, noting that our differenced estimate $\widehat{\beta}$ is the difference of the two simple regression estimates. We use the first assumption that $\frac{cov(ln(\tilde{z}_p), \epsilon_p)}{var(ln(\tilde{z}_p))} = \frac{cov(ln(\tilde{z}_c), \epsilon_c)}{var(ln(\tilde{z}_c))}$ and take differences to get:

$$\begin{aligned}
\widehat{\beta} &= \widehat{\beta}_p - \widehat{\beta}_c \\
&= \beta + \delta \frac{cov(ln(\tilde{z}_p), ln(z_p))}{var(ln(\tilde{z}_p))} - (\beta + \delta) \frac{cov(ln(\tilde{z}_c), ln(z_c))}{var(ln(\tilde{z}_c))}
\end{aligned}$$

[24]Similar results hold if we instead first condition both the left and right hand side variables on the observable controls in Equation 2.5.

Now we use the second assumption that $\frac{cov(ln(\tilde{z}_p),ln(z_p))}{var(ln(\tilde{z}_p))} = \frac{cov(ln(\tilde{z}_c),ln(z_c))}{var(ln(\tilde{z}_c))}$, and we have:

$$\begin{aligned}
\widehat{\beta} = \beta + \delta\frac{cov(ln(\tilde{z}_p),ln(z_p))}{var(ln(\tilde{z}_p))} - (\beta+\delta)\frac{cov(ln(\tilde{z}_c),ln(z_c))}{var(ln(\tilde{z}_c))} \\
= \beta - \beta\frac{cov(ln(\tilde{z}_c),ln(z_c))}{var(ln(\tilde{z}_c))} + \delta[\frac{cov(ln(\tilde{z}_p),ln(z_p))}{var(ln(\tilde{z}_p))} - \frac{cov(ln(\tilde{z}_c),ln(z_c))}{var(ln(\tilde{z}_c))}] \\
= \beta - \beta\frac{cov(ln(\tilde{z}_c),ln(z_c))}{var(ln(\tilde{z}_c))} + \delta\cdot 0 \\
= \beta[1 - \frac{cov(ln(\tilde{z}_c),ln(z_c))}{var(ln(\tilde{z}_c))}]
\end{aligned}$$

So there is a specific bias on the difference estimate. This bias itself can be estimated in the counterfactual sample where I have data on both (out-of-sample) predicted income and true property income. To calculate the bias I run the simple regression of true property income on predicted valuation, and use one minus the coefficient as an estimate of the bias. With this estimate of the bias $\widehat{\text{bias}}$ and the true value of β can be recovered as $\beta = \frac{\widehat{\beta}}{\widehat{\text{bias}}}$. The standard error is estimated by the delta method using the nlcom command in stata.

2.F Assumptions of the behavioural response: robustness to over vs. under valuation

In Section 2.6 I make anim portant choice about the compliance response to billed valuation; I assume that the compliance response is uncorrelated with the difference in billed valuation and true income. I discuss the reasoning behind this choice here

The assumption that the compliance response is uncorrelated with $v_i - z_i$ implies that the response to over valuation is the same as under valuation. However it is reasonable that people react more strongly to over valuation. For instance, a landlord who billed a valuation that is larger than their true income may reduce their compliance more strongly than they would have raised their compliance upon receiving a billed valuation smaller than their true income.

To show that this is a conservative assumption, I consider the extreme case where there is only a compliance response for overvalued properties, and no response for undervalued properties. Then the average compliance response is $\beta = \frac{1}{N}(N_o\overline{\beta} + N_u \cdot \underline{\beta}$ where N_o and N_u denote the number of properties that are over and under valued respectively. I set the compliance response for all those properties who were undervalued at zero, $\underline{\beta} = 0$ and so the average compliance response for those that are overvalued is $\overline{\beta} = \frac{N}{N_o}\beta$. For the calibration sample this is close to half at 49% of properties, and since $\beta = -0.0214$ we have $\overline{\beta} = -.0435$.

I apply these new estimates of the compliance response to properties that were under and over valued and repeat the exercise from Section 2.6.3. The table below gives results. First, since the compliance effect is shut off for undervalued properties the behavioural effect is always negative (row three). Second, the total revenue loss must be larger because there is no compensating behavioural effect from lower bills for the rich. Under this scenario, statistical valuation would reduce total revenues by 177mn UGX or 9.5%, the top twenty percent of landlords would contribute 15% less in revenue, and the bottom eighty percent of landlords would pay about the same amount. Therefore my key results are unchanged; revenues would fall on aggregate and that this would disproportionately benefit the wealthiest landlords.

Statistical Valuation Effect on Revenue (differential response for over and under valued)

	Deciles of Property Wealth										
	1	2	3	4	5	6	7	8	9	10	Total
Baseline Revenue	10	22	37	64	69	127	153	190	340	843	1855
Direct Effect (ΔC_d)	9	4	5	3	3	7	1	-2	-23	-128	-121
Behavioural Effect ($\beta_d \cdot \Delta L_d$)	-2.5	-2.4	-2.4	-3.2	-3.2	-4.9	-5.4	-6.2	-9.2	-15.3	-55
Total Effect	6	2	2	0	0	2	-4	-9	-33	-143	-177

2.G External validity discussion

One issue with analysing the subset of properties that were statistically valued is that these properties may be different in ways that are important for my analysis. The properties that were statistically value were certainly not chosen at random. In fact, the CAMV data collection and statistical valuation were done by division in waves. These waves were prioritized so that Central division, with higher value properties, was completed first, followed by Nakawa division the next year, and only in the final wave did they complete the remaining divisions, Makindye, Lubaga, and Kawempe. It is only in this final wave that the KCCA did statistical valuation because they were pressured to complete on time for the beginning of the financial year.

Figures 2.2 and 2.5 and Table 2.1 describe properties and compare across the three samples; calibration, billed, and counterfactual samples and the remaining taxable properties in the city. In Table 2.1 we can see that property characteristics are similar across samples, though notably fewer properties are commercial in the billed sample. In Figure 2.2b, we see that the distribution of building size in the calibration and billed samples have significant overlap, while the distribution for the whole city is more skewed to the right; it seems some very large buildings are missing from the statistically valued sample of properties. In Figure 2.2a, we see that the full city distribution in terms of rateable value is also skewed to the right compared to the calibration sample of properties. Finally, in Figure 2.5 we can see that the calibration, billed, and counterfactual properties are sampled from similar geographic locations, while missing some areas of the full city entirely.

There are two cases where external validity is important. First, when I analyse hypothetical effective tax rates (i.e. Figures 2.7, 2.8, and 2.9 and Table 2.4) I use the calibration set of properties because here I have data on both z_i and $\tilde{z}_i$ and also enough data to compare across the full distribution of landlord wealth (the counterfactual sample is smaller and so the estimates across the distribution are noisier). If the distribution of hypothetical effective rates in the calibration set differs from the distribution that would have happened in Central and Nakawa then my estimates on the effects on statistical valuation will not be externally valid. Since, by definition, all predictions for the calibration set of properties are estimated in-sample, it is likely that there is less unobserved variation in this sample. In that case, my estimates of regressivity would be under estimated. In addition, Central and Nakawa have a distinct underlying distribution of true property values, they are closer to the city center and are more likely to be high value than those in the three remaining divisions. In this case, if high value properties are measured with more error, then again I would be underestimating the regressivity of satistical valuation, however if properties in Central and Nakawa were better predicted than regressivity may not be as bad.

Second, when I estimate the compliance response to billed liability I assume that the relationship between idiosyncratic propensity to comply and the statistical valuation is the same in the billed sample as in the counterfactual sample. This is a condition for internal validity and to address it I show that the estimate is robust to the inclusion of controls that are unbalanced across the samples, most notably whether the property is residential or commercial. The external validity issue is that, if the behavioural response to the tax liability differs across these samples, then the estimates below will not be externally valid. At least across the calibration, billed, and counterfactual samples there is good reason to believe that the behavioural responses and distribution of effective tax rates will be similar. As explained by the government valuation team, the billed and counterfactual samples were created out of time pressures. With the beginning of the financial year approaching, the valuation team realised they would not complete the individual valuations of all properties in time. The solution was to statistically predict those remaining properties. Since, this was not planned from the beginning the assignment of properties to calibration or predictions samples was quasi-random. Accordingly, it is reasonable to assume that the behavioural estimates from the billed sample and the distribution of effective rates in the calibration sample are at least externally valid within the sample of properties that were statistically valued. This is still a meaningful population to consider as these three districts represent 70% of all properties in the city.

Chapter 3

Eliciting demand for title deeds: Lab-in-the-field evidence from urban Tanzania

3.1 Introduction

Much of urban land in Africa is allocated low values of built capital, remains unplanned, and is settled under informal property rights (UN Habitat 2016*a*, Lall et al. 2017). The development of these cities depends on the formalisation of property rights (Henderson et al. 2020). Formalisation creates transparency in prices enabling functional urban land markets, and improved property records facilitating taxation (Collier et al. 2017). More generally, property rights can reduce expropriation risk, lower the cost of property protection, and remove barriers to credit (Besley & Ghatak 2010).[1]

Establishing property rights is key for the development of cities, but the process is costly for cash-strapped governments in sub-Saharan Africa.[2] To recover program costs once neighbourhoods are surveyed and entered into a town plan, plot-specific fees are charged for title deeds. This stage of formalisation (the uptake of title deeds) remains a bottleneck in many African cities (Omar 2017, Sheuya & Burra 2016, Moses & Chiwambo 2018, Bezu & Holden 2014). This is observed in Dar es Salaam where formal titles account for only 20-25% of residential surveyed plots.

Integrating traditional local leaders in the formalisation process has the potential to raise

[1] There is empirical evidence that property rights have a positive impact on investment in rural Africa (Besley 1995, Goldstein & Udry 2008). For urban land, evidence is concentrated in South America finding impacts on household investment, education and labour supply (Field 2007, Galiani & Schargrodsky 2010).

[2] Formalization requires surveying and town planning to meet the standards of formal law. There are scale economies to surveying, and so governments and development agencies alike make efforts to coordinate land demarcation (surveying) en masse.

the uptake of title deeds. Institutions in Africa have long relied on traditional local leaders (Michalopoulos & Papaioannou 2013). While formalisation can be seen as eclipsing their role in the land tenure system, these leaders are complementary to state capacity when they are formally integrated (Henn 2020). For instance, property tax collection by local leaders raises more revenue than collection by state agents because of their knowledge of local individual's payment propensities (Balan et al. 2020).

If leaders know and will reveal information on the willingness-to-pay for title deeds, it can be used by the state to better target fees, i.e. by charging less to owners with lower willingness-to-pay. This can raise uptake while still covering project costs. However, extracting this information accurately can be difficult if leaders have private reasons to favour some property owners in their neighbourhood. Two obvious questions arise which are the focus of our paper; are leaders informed about the willingness-to-pay for title deeds? and if so, will they share this information accurately when they are able to influence the prices faced by their neighbours?

This paper is related to a literature that studies the use of agents to target subsidies for the poor, to identify individuals with high returns to loans, and to report corruption (Olken 2009, Niehaus et al. 2013, Basurto et al. 2018). Agents may have different preferences from the social planner and strategically give misleading information. In a closely related paper to ours, Rigol et al. (2020) test whether cash incentives can encourage entrepreneurs to report which of their peers have the highest marginal returns to a loan. Our paper is, to the best of our knowledge, the first to study whether agents can be used to extract information on willingness-to-pay.

There is also a long literature on eliciting willingness-to-pay for non-market based goods. In a related paper, Ali et al. (2016) estimate the demand for title deeds in a neighbourhood of Dar es Salaam using a take-it-or-leave-it randomisation of title fees. Their method estimates mean compliance conditional on fee size, and so cannot be used to determine individual willingness-to-pay. In another related paper, (Berry et al. 2020) elicit the willingness-to-pay for water filters using the Becker-deGroot-Marschak (BDM) method. This method does allow the researcher to estimate individual willingness-to-pay, however the policy maker cannot use it to set fees in practice.[3] Our paper provides a method (by eliciting third party information) that both identifies individual willingness-to-pay, and can be practically implemented.

[3]The BDM cannot be used in practice by the policy maker because it relies on the credible incentive that the customer will be able to buy the good at a random price.

In another related study, Balan et al. (2020) show that tax collection by local elites can raise more revenue than collection by state agents. Their evidence suggests that the primary mechanism trough which this works is informational advantages of chiefs that enabled them to better target tax visits based on households' underlying payment propensities. They test this with a treatment arm where state collectors meet with local chiefs and indicate, address by address, ability and willingness-to-pay. Our paper sheds light on this particular mechanism by directly measuring the ability of local leaders to predict willingness-to-pay, and by studying the conditions under which this information can be accurately extracted.

Our paper makes three contributions. First we challenge the view that the low uptake of title deeds is due to plot owners not recognising, or not needing, the benefits from tenure formalisation Briggs (2011). We provide evidence of significant demand for title deeds, albeit at lower prices than the government is currently charging. We estimate the demand for property titles using the Becker-deGroot-Marschak (BDM) method which incentivises respondents to truthfully reveal their willingness-to-pay.[4] Roughly 40% of plot owners in our study are willing to pay fees equal to the monthly income of a typical household. This is much higher demand than is found in previous work in Dar es Salaam (Ali et al. 2016). Yet, demand remains lower than current fees. We suggest that the government could set lower prices so as to raise the uptake of titles.

Our second contribution is to show that local leaders have accurate information about both the aggregate demand curve in their neighbourhoods, as well as, the ability to distinguish variation in willingness-to-pay across owners in their neighbourhood. This is true even when conditioning on the fee size, or property value. Therefore, the local knowledge of community leaders can be used to set prices of land titles so as to raise uptake and collect sufficient revenue. This would help to make formalisation inclusive for the urban poor and financially viable for the government.

Our third contribution, is to show that, when predicting willingness-to-pay, leaders are influenced by an environment where their predictions are used to allocate subsidies, but also that almost all of these distortions can be mitigated with a simple cash prize for ex-post accuracy. Notably, we find that leaders only distort their average response; we find no evidence that this environment affects leader predictions when it comes to discrimination across different plot owners in their neighbourhood.

Finally, forty-two of our respondents were selected, at random, to undertake in-depth in-

[4]The BDM method was originally developed by Becker et al. (1964) and is still used at the frontier of applied work (Berry et al. 2020).

terviews digging into the determinants of their willingness-to-pay, including the expected benefits and costs of tenure formalisation, alongside other factors. Results from this qualitative work are discussed in a companion paper (Manara & Regan 2020).

The paper is outlined as follows. The experimental setting and design are described in section 2 which covers the study context and data collection methods. The data is described in section 3 including a description of the demand elicited by BDM. Results in section 4 show leaders' ability to predict demand on the aggregate and for price discrimination. In section 5 we conclude.

3.2 Context and Setting

3.2.1 Conceptual motivation for extracting willingness-to-pay

In this paper we propose that, by collaborating with leaders who have local knowledge, the central government can more effectively target fees to both neighbourhoods and individual plots and owners. In doing so, this price discrimination could raise uptake while still meeting the fee requirements to cover cost. By price discriminating, the government can cross-subsidize and thereby raise the revenue required to make a titling program cost effective. The potential gains are twofold. First, conditional on the project being complete, price discrimination can be used to recover the Harberger triangle deadweight loss. In a simple example this would mean waiving fees for particular individuals who are identified as having a low ability and willingness-to-pay. The second set of potential gains can be much larger. In a context where projects are only built if the expected revenue is above the fixed cost, then price discrimination can make the entire project viable which can lead to significant gains (Kremer & Snyder 2018). This reasoning is similar in spirit to Romer (1994) who shows the potential for large gains from trade when 'new' goods are introduced in the market by raising enough revenue to cover a fixed cost of entry.

An alternative response is simply for the government to cover these project costs and give away the titles for free. There are two reasons why this may not be possible. First, is that the government may not be able to secure the necessary funds to do so, or the efficiency of raising public funds may be so low as to make it unviable. This may be especially true if channels of raising revenue are limited or wasteful as is the case in many developing countries (Pomeranz & Vila-Belda 2019). Secondly, a growing body of research underscores that building capacity for revenue collection is important for state development (Besley & Persson 2014). From this perspective, governments should not universally subsidise formalisation, but rather encourage those with higher private benefits to make more of a contribution to the public fund. Finally, while it might seem radical to advocate that the government price discriminates when allocating property titles, it is important to observe that the current fee structure already employs a price discrimination strategy. Invoice fees are calculated based on location, land use, and individual plot size. Further, in the private market for survey services, a basic version of leader-elicited price discrimination is already employed. The largest surveying and planning company in Tanzania, offers a 'free lunch' to individual plot owners that can not afford to pay the survey fees. To do this, they hold discussions with local leaders who help them identify the plot owners with the lowest willingness-to-pay. As long as these individuals do not own plots above 800m^2, they are offered the service for free.[5]

3.2.2 Experimental Setting

Our study was conducted in Dar es Salaam where the Tanzanian Ministry of Lands, Housing and Human Settlements Development (MLHHSD) designed and implemented a pilot project to formalise land. Here uptake has been less than 13% after the first two years. Since the government has fronted the fixed cost of surveying and planning they have lessened coordination issues, and now plot owners can simply pay their invoice to complete the process of acquiring a title.

The title that we study is a legal document of ownership, Certificate of Right of Occupancy (CRO), that is supplied by the MLHHSD and provides the highest protection by law in the country. A CRO formally recognizes a 66 year lease of a plot of land from the government. Legally a CRO provides private benefits in three ways; protection from government-led expropriation, use as collateral with mainstream banks, and legal transferability of land.[6]

A plot of land must be surveyed and approved by the municipal town planning office to be eligible for a CRO. There are scale economies to surveying; the survey of a standalone plot may cost around 6 million TSh while the average cost drops to 17% of this when 10

[5]From author's discussions.

[6]More specifically, owners of a CRO who are expropriated by the government are entitled to higher compensation, and since surveying is a pre-requisite, documentation of exact plot boundaries mitigates potential conflict with neighbours (Wolff et al. 2018). While ownership of a plot without CRO can be enough for access to small loans, these typically have a maximum ceiling of 20 million TSh, when the informal sale agreement is used as collateral. Instead, there is no ceiling for loans pledged against CRO: in this case, the loan amount is only limited by the collateral value and the bank's single borrower limit. (Manara & Pani 2020). For land sales the CRO provides the buyer a guarantee of the seller's rightful ownership. While land is often sold informally these types of sales are susceptible to scamming with land being sold to multiple people. Wolff et al. (2018) describe a case in Kigamboni, Dar es Salaam, where a single plot was sold to over 30 individuals.

plots are surveyed at once, 5.8% for 100 plots, and for large scale projects with more than 1,000 plots the average cost is about 0.2 million TSh.[7] For this reason the MLHHSD has run a pilot program in the ward of Kimara, Dar es Salaam, surveying plots at scale. The costs of this program include survey and administrative fees.[8] Facing these fees, demand for CROs has been low. This motivates concerns over the effectiveness of the government's pricing strategy for raising revenues and supplying affordable CROs.

We conduct a survey and experimental sessions with 90 local leaders from fifteen neighbourhoods where surveying was complete and invoices issued to plot owners.[9] The neighbourhoods were all located in Kilungule A and B, two Mtaas [10]in Ubungo Municipality of Dar es Salaam. For simplicity we refer to our study area as Kilungule and it is shown in Figure 3.3. In the following weeks we conducted information sessions with plot owners. Finally, we surveyed 146 owners and offered them a discounted invoice through the BDM method. These owners had yet to pay their invoice by the start of the intervention (April 2019); over three years since the formalization project commencement. The full time line of the study is depicted in Figure 3.1.

3.2.3 Data Collection

Sample Selection

We collected CRO invoice records of all 1,482 invoiced plots in our study area and matched 1,401 of these to geo-located plot boundaries. Of these, only 13% had purchased their title deed, even though 28% had been invoiced over two years earlier, and only 3% had been invoiced within the last six months. From this population we randomly sampled fifteen invoiced plots from each neighbourhood in our study area, for a total of 225 plots. We stratified our sampling so that low, medium, and high value plots were represented in each neighbourhood. We then conducted a rapid survey of the selected plot owners in order

[7]From author's discussions with two of Tanzania's leading survey companies.

[8]Figure 3.2 shows an example of an invoice from the Kimara program with charges that include costs for surveying, administrative costs and also land value capture 'Premium' and 'Revolving Fund'. The latter is a mark-up used to subsidise future surveying projects. Some fees are fixed (Application, CRO, and Deed Plan), while all others vary with plot size and land value.

[9]The specific title of these leaders is 'mjumbe', or 'wajumbe' (plural), and while they represent a political party they are unofficial and unpaid positions and so bear a quasi-formal status (Manara 2020). The neighbourhoods, or 'shinas', typically contain 250 plots, and there are roughly six wajumbe per shina. In our study area both the ruling CCM party and the opposition Chadema party were represented by wajumbe and, in each neighbourhood, a given party will have one mjumbe and a collection of assistant wajumbe.

[10]Mtaa is the smaller administrative unit and the lower level local government in Tanzania. Typically referred as *street or neighbourhood.*

to gather their contact information, occupancy (i.e. owner or tenant), and their social connection to each leader (i.e. whether they knew or ever interacted with their leaders).

Following the invoice collection, we conducted a census of the 96 leaders in our study area which allowed us to match them to neighbourhoods, classify their party affiliation (CCM or Chadema) and hierarchical position (main leader or leader assistant), and geo-locate their residence. From this population we randomly assigned leaders to one of three treatment groups. We stratified the randomisation based on political affiliation and physical and social proximity to sampled invoiced plots.

All randomisation was done mechanically during a series of public meetings where the process could be observed. Despite the potential for mechanical error, this was an important procedure to garner trust with the community. It also provided a practical experience with randomisation so that those who also participated in the BDM at the end of the study were already familiar with the lottery process.

Owner Information Sessions

We invited all 225 sampled plot owners to attend an information session to introduce them to our project, two to three weeks before their actual research session. The focus of the information session was on familiarising the respondents with the BDM procedure. They were told that, during the research session, they would be asked "What is the maximum price that you would and could pay in the next 10 days for your invoice towards your title deed?". We then explained the concept of willingness-to-pay both in theory and with examples. They were told that on research day they would have an opportunity to commit to pay their invoice if it was offered at a price they could afford, and so it was important that they thought carefully over the following weeks about their willingness-to-pay for the title deed. We then explained the specifics of the BDM method and that their best strategy was to determine for themselves their true willingness-to-pay and then reveal exactly that price to the surveyor. We used theory and examples to show why this was the best strategy for them. We finished the session by practising with volunteers for either a soda or an aerial photo of their plot. Throughout the session we asked for feedback from respondents until it was clear they understood.

Owners were given at least two weeks between the information session and the price elicitation session. During this period they were encouraged to consult others (family, joint plot owners, friends, etc.) on their willingness-to-pay and plan out a strategy for gathering the funds they may need if they won a discount. This time was also used to sort out individual issues with each invoice. Some of these issues were simple for us and the Ubungo Municipal Office to accommodate, such as the misspelling of names, partial payments already made, and the addition of spouses to invoices. On few exceptional circumstances, we allowed 'decision makers' to participate on behalf of the true owner on the invoice.[11] For other issues we had to drop invoices from our sample. This was the case where, by the time of starting the study, invoices had already been fully paid or the plots sold (19 cases), where at least one owner had deceased (4 cases), when the owner lived out of country and could not be reached (13 cases) or had conflicts (5 cases) over the rightful ownership. After discarding these issues there were 184 remaining invoiced plots.

Owner Survey and Price Elicitation

We invited 184 eligible plot owners to participate in a survey and price elicitation lottery, and 146 of these attended. We also invited a leader to each session in order to establish trust with the respondents. The survey collected information on demographics, a CRO knowledge test, sentiments towards tenure security, and perceived costs and benefits of a CRO. Following the survey each respondent participated in the BDM price elicitation. This began with a practice round where the respondents were randomly assigned the opportunity to purchase either a soda or an aerial photo of their plot (see Figure 3.4 for an example) through the BDM mechanism. Following the practice, they were offered the opportunity to acquire the title deed for their plot at a discounted invoice price, again through the BDM mechanism. If the respondent won the discount, they were scheduled to make their payment within ten days.[12]

The BDM procedure that we implement closely follows that of Berry et al. (2020) with slight adjustments to our context. Respondents stated their willingness-to-pay (bid) and participated in a lottery extracting a new invoice price (draw). According to standard BDM procedure, if the draw was lower or equal to their bid, they would be offered the title deed at the new discounted price; if the draw was higher, they would not be offered a new price. Each BDM session began with a description of the procedure followed by a practice for either a soda or an aerial photo of their plot before proceeding with their

[11] This was done in two exceptional cases, one where the plot owner had been living in Canada for over thirty years and his brother was the de-facto owner of the plot, and a second where the plot owner was disabled and her son took on responsibility for the plot. In both cases the decision maker was responsible for paying the invoice, and in neither case did we change the name of the plot owner on the invoice.

[12] Because the title deed cost was high for many households we did not ask for immediate payment. First, asking respondents to bring the full amount of cash necessary to cover their bid value would be a significant wasted effort in the case that they did not win. The second reason was to allow enough time to gather funds from family, friends, or micro-lending groups. In fact, 15% of respondents admitted asking the financial support of family and friends to make higher bids (Manara & Regan 2020).

invoice. Practice rounds enabled respondents to understand that their bid should represent the maximum price they could and would like to pay; their bid could not be changed after the lottery; and, upon winning, they must make the according payment within ten days. Once the bid for the invoice was finalized, a price was drawn which determined whether the respondent would pay for the invoice at the drawn price.[13]

There were 39 respondents who drew prices lower than their bid and so won a discounted invoice value. For each, we confirmed that they could pay and that they had a plan to collect the necessary funds, and had them sign off on their bid value and draw outcome. All participants received a 10,000TSh cash allowance for their participation, and winners were required to use this as a down-payment in order to discourage overstating their willingness-to-pay. Still, five (12.8 percent) of the winners did not complete the purchase.[14]

Leader Survey and Experiment

We conducted surveys with the leaders one month in advance of the first plot owner price elicitation session. All 96 leaders in our study area were invited to participate and 90 (93.8 percent) attended and completed the survey. The questionnaire consisted of demographics, a CRO knowledge test, social network, and predictions of plot owner characteristics. The network and prediction questions all related to the owners of the 15 selected plots in each leader's respective neighbourhood. For reference the leaders were given both official names and nicknames of each owner as well as a photo of the particular plot that was selected.

The survey concluded with price elicitation tasks. Leaders were asked to rank each of the 15 plot owners in their neighbourhood in terms of their willingness-to-pay for the title deed. After ranking, leaders had also to predict, for each plot owner, their exact willingness-to-pay. Each leader conducted the task under their randomly assigned treatment.

Leaders assigned to the *control* group were told that the research was conducted for academic purposes only. They were encouraged to be as truthful and accurate as possible to enable high quality research. Finally, they were ensured that their answers would not be used to change any procedure over the course of the study.

Leaders in the *stakes* group were informed that their responses to the price elicitation

[13] Practically, the respondents drew 1 of 75 plastic balls from an opaque jug. Each ball corresponded to a price between 0 and their full invoice value which was recorded on a reference sheet. The exact distribution depended on the size of their plot (thus, indirectly, the invoice value). In order to maintain goodwill the distribution of prices was shown upon request just before the price was drawn and none of the respondents asked to change their bid after seeing the distribution.

[14] Three of these cited unforeseen health issues with a family member that became a priority for the funds that were allocated to the title deed. One was unable to borrow the money that she had expected to. The last one went away on business and was unwilling to arrange a representative to make her payment.

tasks would be used to change procedures in the remaining study; particularly, to help decide which plot owners would have higher chances to win high discounts through the lottery. If leaders suggested a plot owner had a low willingness-to-pay, we would adjust the distribution of discounts available in the lottery to this plot owner so as to make it more likely that they win a high discount.[15]

Finally, leaders in the *incentives* group received the same instructions as the stakes group, but they had the opportunity to earn cash for their accuracy. We adopted an ex-post payment rule that would be implementable in a policy setting.[16] Leaders were given simple payment examples to work through. Each leader was reminded that being as truthful and accurate as possible was the best way to earn the cash. At the end, the leader with the most points was paid 30,000TSh and the four runner-ups were paid 20,000TSh each.

3.3 Data and Sample Descriptives

3.3.1 Summary Statistics and Balance

In Table 3.1 we present mean characteristics for the whole sample of both the plot owners (column 1) and leaders (column 2). Compared to leaders, plot owners tend to be younger and more highly educated but score worse on a short quiz about CROs and have lower household monthly incomes. While the majority of owners are male there is a significant share of female ownership (73% of plots have sole ownership, and 35% of these are owned by women). This is in line with previous findings that the cultural environment in Dar es Salaam is not particularly opposed to co-titling or female ownership (Ali et al. 2016). Leaders are also slightly more likely to be men, but 40% of them are women. Considering potential heirs, 92% of plots are owned by individuals with at least one child.

The average CRO invoice value is 526,000 TSh, or roughly two and a half times the median monthly income in our sample. Most plots are occupied by their owners, still 28% are owned by absentee landlords. A full 86% of plots were acquired by purchase, as opposed to inheritance or squatting, though only 24% of all plots have a certificate of sale (*hati ya*

[15] We adjusted the distribution of available discounts based on an average of leader predictions for the same plot owner, therefore mitigating concerns over the ethical aspects of this treatment.

[16] Before the price elicitation tasks, leaders were explained that, at the end of the study, we would pick one price level and observe which plot owners stated willingness-to-pay above that price. For each plot owner with stated willingness-to-pay above the threshold price, leaders would get a number of points corresponding to the assigned ranking position of that plot owner. This is implementable in a real world setting, since the policy maker will observe which of the plot owners do in fact uptake titles. If titles were purchased by owners for whom the leader ranked high, then the leader was accurate.

mauzo or *sale agreement*). Exactly half of the plots are owned by individuals with at least one other plot in Tanzania.

Leaders themselves tend to own their homes; 94% own the plot they live on while the remainder all live on a plot owned by a member of their household. Leaders typically have a long history of residence in Kilungule; while only 7% have settled in the past six years, a full 38% have been living there for over 19 years. Out of the fifteen owners sampled for each neighbourhood, leaders know on average 12 of these, though only 4.3 have ever come to the leader for official assistance. Leaders have few social connections among the plot owners in the study; on average 0.22 owners are family, 1.4 are friends, 1.8 meet together regularly for religious purposes, and 1.3 are considered highly esteemed by the community.

Also presented in Table 3.1 are differences in leaders' characteristics between *stakes* and *control* groups (column 4) and differences between *incentives* and *control* groups (column 5). There are only a few marginally significant differences, though standard errors are large. The *stakes* group has fewer women and more leaders with household income below 100,000TSh compared to the *control* group. The *incentives* group has more leaders with their home plot surveyed than the *control* group.

3.3.2 Demand for CROs

Figure 3.5 describes the demand for CROs elicited through the BDM. For the BDM demand curve we show, for each price, the share of plot owners whose bid was greater than or equal to that price. This is done by running successive logit regressions at each price point and correcting for heteroskedasticity in the calculation of the confidence intervals.

While the full sample of plot owners were not willing to purchase a title deed at their invoiced price there is still a significant amount of demand for CROs. Over 40 percent of plot owners would be willing to pay 200,000 TSh which is more than the monthly household income of half of our respondents. However, demand is still much below invoice fees that are currently being charged. The median invoice in our sample is 500,000 TSh, at such a price less than 10 percent of plot owners would be willing to pay. Even if all plots were charged 170,000 TSh, the minimum invoice value observed in our sample, roughly 50 percent of plot owners would not purchase a title deed.

In Figure 3.5b we compare the elasticity of demand calculated from smoothed version of the demand curve above. There is a wide range (200-600,000 TSh) where demand is relatively elastic, beyond which we have trouble estimating due to the sparsity of observations in the tail of the distribution.

3.3.3 Leader Predictions and Placebos

Because we are interested in knowing whether leader's have accurate knowledge of the willingness-to-pay for CROs in their neighbourhoods, we first check if they have knowledge on a related set of plot characteristics. In Table 3.2 we run regressions based on the model:

$$y_{ij} = \beta \hat{y}_{ij} + x_j' \gamma + \epsilon_{ij} \tag{3.1}$$

where y_{ij} is a characteristic of plot i related to leader j, $\hat{y}_{ij}$ is leader j's prediction of plot i's characteristic and x_j is a vector of leader controls for randomization strata, neighbourhood, and surveyor id.

Panel A shows that leaders predictions of plot and plot owner characteristics are positively associated with their true characteristics. In column 1 leaders are able to distinguish between owners with higher or lower income. Our estimate for income rank are very similar as those found by Rigol et al. (2020) for Indian entrepreneurs' predictions of their peer's income rank. In column 2 we show leader's predictions of CRO invoice value rank are positively associated with the true CRO invoice value rank of plots in our sample, and for column 3 this is also true across the full sample of plots.[17] Therefore column 3 signals that our plot owner survey sample is not selected towards plots that are easier to predict. In columns 4 and 5 we can see that leaders also have some ability to predict whether plot owners have paid their property tax or if they have a certificate of sale.

Panel B does placebo tests by comparing the relationship of leader predictions and actual characteristics across treatment groups. It is a placebo because all of these predictions were given by leaders before they were assigned their treatment. Although there are no significant differential coefficients of either the stakes or the incentives treatments, the stakes group does have a slightly higher differential between predicted and observed for each characteristic.

3.4 Results

3.4.1 Leader Predictions of Aggregate Willingness-to-Pay

Figure 3.6 compares demand for CROs elicited through the BDM mechanism with that elicited through local leaders. For the BDM demand curve we show, for each price, the

[17]The full sample of plots includes those plot owners that were not part of our plot owner survey.

share of plot owners whose bid was above that price. We follow a similar procedure for the leader predicted demand curve, but use the leader prediction of the plot owner's willingness-to-pay instead of the owner's bid. Since there are multiple leaders for any given owner, and so multiple predictions of their willingness-to-pay, we cluster standard errors at the plot owner level. The same 146 plots are used to construct both the BDM and leader predicted demand curves.

In Figure 3.6a, we only use leaders under the *control* group and compare the demand curve based on their responses with the BDM results. Whether demand is elicited from the BDM mechanism or predicted by leaders, the curves are strikingly similar. At least on an aggregate level, leaders seem to have knowledge of the distribution of willingness-to-pay in their neighbourhoods.

When leaders are told that their responses will be used to determine the likelihood that a plot owner receives a discount (*stakes*) they distort their responses. Figure 3.6b uses only leaders under the *stakes* treatment and compares the demand curve based on their responses with that based on the BDM. For most prices where demand is positive there is a large gap between the demand curve elicited from this group of leaders compared to the BDM. This suggests that, despite their predictive ability, eliciting aggregate demand from leaders may be difficult in a setting where their responses are used to price CROs in the community.

Offering a monetary incentive to leaders for their predictive accuracy (*incentives*) can mitigate the distortions created in the *stakes* environment. Figure 3.6c uses only leaders under the *incentives* treatment and compares the demand curve based on their responses with that based on the BDM. Whether demand is elicited from the BDM mechanism or predicted by leaders with incentives, the curves are statistically indistinguishable. This is not only due to wide confidence intervals. The largest gap between the point estimates of leader and BDM elicited demand curves is a 0.08 point difference, and for most prices the gap is less than a 0.03 point difference. The cash incentive has shrunk the gap that occurs when leaders are told that their responses will used to determine discounts.

3.4.2 Leader's ability to distinguish willingness-to-pay across owners

While leaders may be able to predict the aggregate distribution of demand fairly well, it remains to be seen if they can also distinguish between individuals with high and low willingness-to-pay. In this section we describe the ability of leaders to distinguish individuals with high and low willingness-to-pay by running regressions based on the model:

$$w_{ij} = \beta \hat{w}_{ij} + x_j' \gamma + \epsilon_{ij} \tag{3.2}$$

where w_{ij} is willingness-to-pay of plot i related to leader j, $\hat{w}_{ij}$ is leader j's prediction of plot i's willingness-to-pay and x_j is a vector of leader controls for randomization strata, neighbourhood, and surveyor id.

In Table 3.3 Panel A we show the coefficient on leader's predictions of different measures of owner's willingness-to-pay is always positively associated with the true measure of owners willingness to pay. Column 1 considers the within neighbourhood rank; an individual predicted to be one position higher in the ranking is on average 0.2 positions higher in the rank of plot owners' bids. Column 2 uses the actual level of willingness-to-pay; an individual predicted to bid 10,000TSh above another will on average bid 1,000TSh more. Column 3 takes the log of willingness-to-pay; a one percent increase in predicted willingess-to-pay translates to a 0.33 percent increase in actual willingness-to-pay on average. Column 4 takes the percentile rank of all owners in the sample (rather than within neighbourhood). Here moving from an individual at the median to one at the 60th percentile of predictions results in a 2.3 percentile increase in the true willingness-to-pay on average. Finally, columns 5 and 6 use the probability of being the top or bottom rank in the neighbourhood; an individual is 15 percentage points more likely to be the highest willingness-to-pay in the neighbourhood if predicted to be so, and 24 percentage points more likely to be the lowest willingness-to-pay if predicted so. On this last point it is of interest to note that one of the largest land surveying companies in Tanzania runs a 'free lunch' program, consulting local leaders before charging fees to determine plot owners in the neighbourhood who are in need of a discount. In one of their larger projects, they surveyed over 5,000 plots and used leader information to waive fees for about 2% of the plot owners.

Leaders may have knowledge of individual willingness-to-pay, and yet, they may distort their responses if it can help certain plot owners win or if they are paid incentives for accuracy. In Panel B we analyse the impact of the *stakes* and *incentives* environments. To do so we adjust model 3.2 to account for the differential coefficient for leaders in different environments:

$$w_{ij} = \beta \hat{w}_{ij} + \beta^S \hat{w}_{ij} \mathbb{1}(j \in stakes) + \beta^I \hat{w}_{ij} \mathbb{1}(j \in incentives) + \alpha_S + \alpha_I + x_j' \gamma + \epsilon_{ij} \tag{3.3}$$

where $\mathbb{1}(j \in stakes)$ is an indicator if leader j was assigned to the stakes environment,

$\mathbb{1}(j \in \textit{incentives})$ is an indicator if leader j was assigned to the incentives environment, and α_S and α_I are dummies for each treatment group. Returning to Table 3.3 Panel B, none of the differential coefficients of either environment is significantly different from zero at the five percent level. Therefore we find no evidence that the stakes or incentives environments create distortions in the predictions.

3.4.3 Property Characteristics to distinguish willingness-to-pay across owners

The government currently charges for CROs with a formula based on ward level land values, plot area and land use. In this section we examine the ability of this formula to target high and low willingness-to-pay individuals. In addition, we create a measure of property values based on photos of the plot, and local knowledge of the area.[18] We consider this measure of property value as another potential indicator to price discriminate on. Below we show how variation in property and invoice value relate to willingness-to-pay of plot owners.

In Table 3.4 we run regressions of the general form:

$$w_{ij} = \alpha z_{ij} + \beta \hat{w}_{ij} + x_j' \gamma + \epsilon_{ij} \tag{3.4}$$

where z_{ij} is either invoiced fee or property valuation of plot i related to leader j. When willingness-to-pay is transformed, we also transform the observable characteristic accordingly, e.g. in Panel A column 1 where the outcome is the rank of willingness-to-pay, we use the rank of invoice value as the explanatory variable.

In panel A we use invoice value unconditional of the leaders prediction of willingness-to-pay. Across columns 1-5 invoice values are positively associated with individual willingness-to-pay. Column 6 shows that the bottom rank willingness-to-pay is particularly difficult to predict with the invoice value. Otherwise the invoice value correlates strongly with willingness-to-pay, with coefficients that are typically closer to 1 than the leader predictions in Table 3.3 Panel A. Finally we note that, while variation in invoice value closely follows that of willingness-to-pay, the average invoice value is more than 2.7 times that of the average willingness-to-pay (Table 3.1 Column 1).

In Panel B we include the leader's prediction in addition to the invoice value. In columns 1-4

[18]This follows the procedure that is used for property valuation by local governments and the Ministry of Lands. The valuations are based on the subjective determination of three students from Ardhi University, a local university which specialises in surveying, planning, and valuation.

we show that, conditional on the invoice value, the leaders are still able to explain variation in the willingness-to-pay. This suggests that invoice formula and leader predictions could be applied complementary to one another. Finally, conditional on invoice value, leaders are not able to capture any variation when it comes to the the top rank of willingness-to-pay. Instead, when considering the bottom rank, leaders are effective while the invoice value is not.

Moving to Panel C we use property valuation unconditional of the leaders prediction of willingness-to-pay. In columns 1-4 the property valuation is positively associated with willingness-to-pay and the correlations are of similar magnitude than the leader predictions in Table 3.3 Panel A. However, in columns 5 and 6 the subjective valuation of the property has no ability to predict the top or bottom ranked willingness-to-pay. In Panel D columns 1-4 we show that subjective property value and leader prediction are both able to describe variation in willingness-to-pay conditional on one another. In columns 5 and 6 only the leader prediction is able to describe the variation in the top and bottom rank willingness-to-pay.

3.4.4 Can willingness-to-pay cover project costs

In this section we do back-of-the-envelope calculations to determine whether the willingness-to-pay is high enough to cover the cost of the project. Currently, 13% of invoices have been paid and their average fee was 616,000TSh. Therefore, the government raised about 80,000TSh on average. We sampled plots from the remainder of invoices, and here the average willingness-to-pay was 194,000TSh. Taking this figure as representative for the entire 87% of unpaid invoices, the maximum revenue that could be extracted from the remainder, averaged across the entire sample, is 0.87*194,000=168,780TSh. Together the average potential revenue is about 249,000TSh.[19] Considering that the average cost of surveying a plot is about 200,000TSh for large projects (quote from two private survey companies), and comparing this to the average willingness-to-pay, we realize that the costs of the project are covered and there is an average gain of 50,000TSh per plot. That means that the gains outweigh the costs, even only counting the perceived private gains to the title document. Furthermore, in a separate study we document that plot owners already perceive large benefits to the surveying and allocation of beacons regardless of the acquisition of the title document (Manara & Regan 2020). Thus the overall private gains are larger than those captured by the willingness-to-pay for the title alone. In addition, there

[19] Note that this is a conservative estimate since the willingness-to-pay of the 13% of property owners who hae aready paid, must have had a willingness-to-pay above their invoiced fee. Here we assume that their willingness-to-pay was equal to the fee.

are likely further gains to surveying and titling that are not internalised by the current plot owners, which will manifest in the long-run (Michaels et al. 2020). Together this suggests that the gains to formalisation can far outweigh the costs of surveying and planning.

3.5 Conclusion

Despite there being low uptake of property titles in much of urban Africa, we find that demand for property titles, as elicited by the BDM method, is substantial for a community in Dar es Salaam, Tanzania, where the uptake of land titles is currently low. We find that local leaders are able to predict this demand ex-ante, both on the aggregate and at the individual level. Under an environment where the responses of leaders are used to allocate subsidies, their predictions deteriorate. However, an incentive scheme of cash prizes can mitigate this.

African governments adopt land tenure reforms to contrast the socio-economic issues connected with unplanned and rapid urbanisation, essentially pushing for a transition from informal land tenure systems to more formalised ones. Our evidence suggests leveraging on the local knowledge of key actors in the left behind informal system as a means to make formalisation projects more affordable and inclusive.

3.A Figures

Figure 3.1: Project Timeline

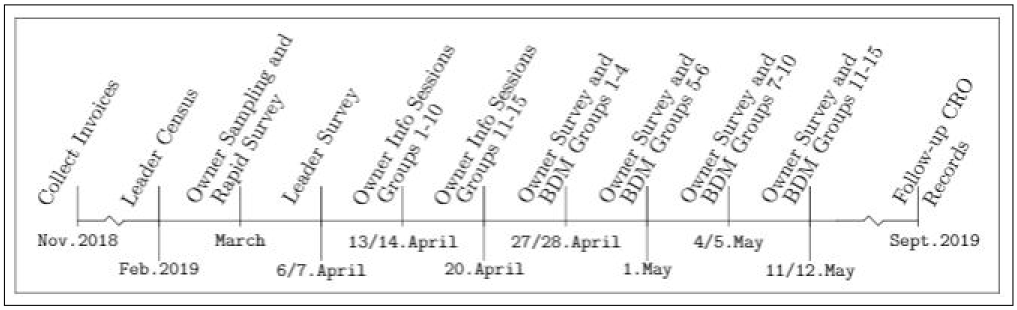

Figure 3.2: Example of an Invoice for a CRO

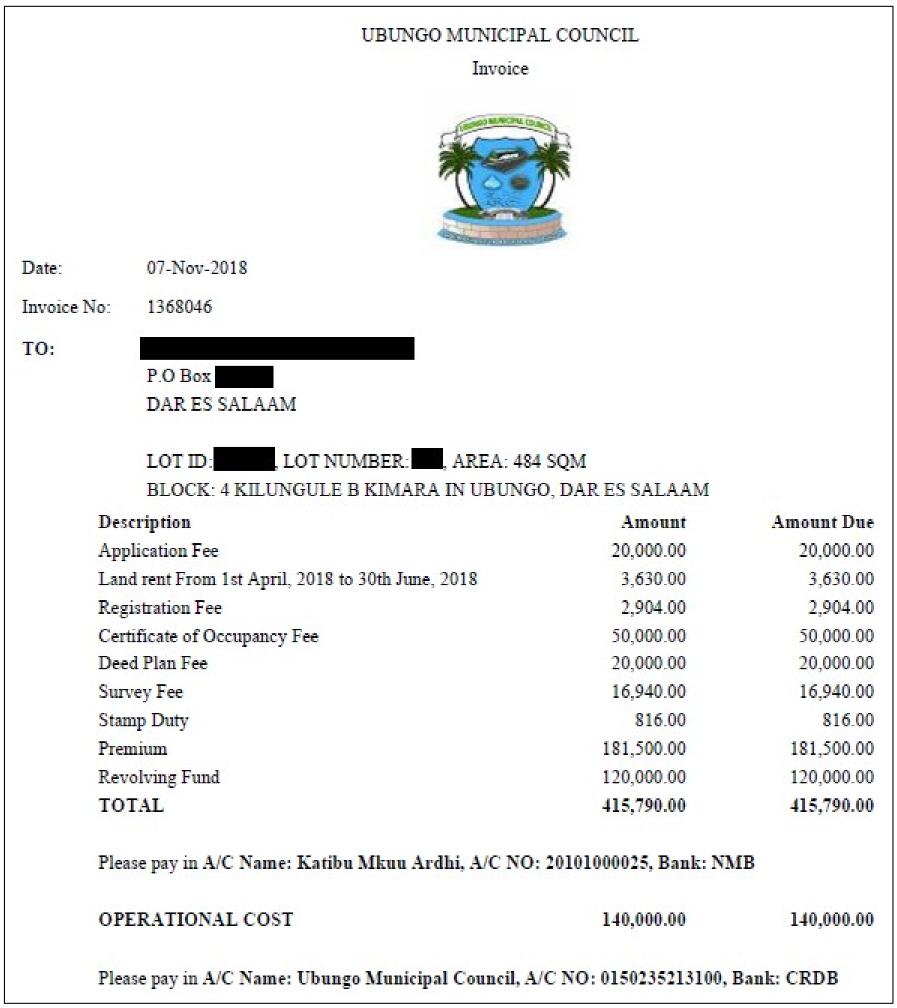

UBUNGO MUNICIPAL COUNCIL

Invoice

Date: 07-Nov-2018

Invoice No: 1368046

TO: [redacted]

P.O Box [redacted]

DAR ES SALAAM

LOT ID: [redacted], LOT NUMBER: [redacted], AREA: 484 SQM

BLOCK: 4 KILUNGULE B KIMARA IN UBUNGO, DAR ES SALAAM

Description	Amount	Amount Due
Application Fee	20,000.00	20,000.00
Land rent From 1st April, 2018 to 30th June, 2018	3,630.00	3,630.00
Registration Fee	2,904.00	2,904.00
Certificate of Occupancy Fee	50,000.00	50,000.00
Deed Plan Fee	20,000.00	20,000.00
Survey Fee	16,940.00	16,940.00
Stamp Duty	816.00	816.00
Premium	181,500.00	181,500.00
Revolving Fund	120,000.00	120,000.00
TOTAL	**415,790.00**	**415,790.00**

Please pay in **A/C Name: Katibu Mkuu Ardhi, A/C NO: 20101000025, Bank: NMB**

OPERATIONAL COST	**140,000.00**	**140,000.00**

Please pay in **A/C Name: Ubungo Municipal Council, A/C NO: 0150235213100, Bank: CRDB**

Figure 3.3: Location of Study Area in Dar es Salaam

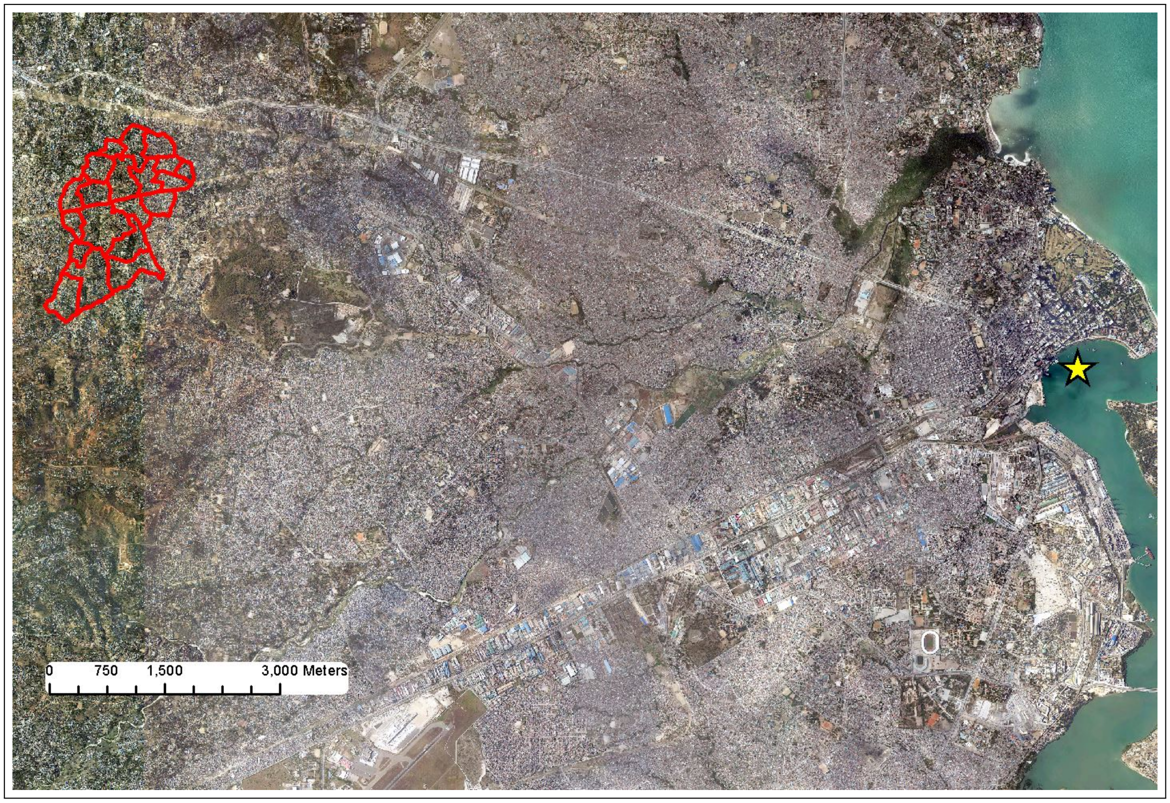

Figure 3.4: Example of a Plot Photo

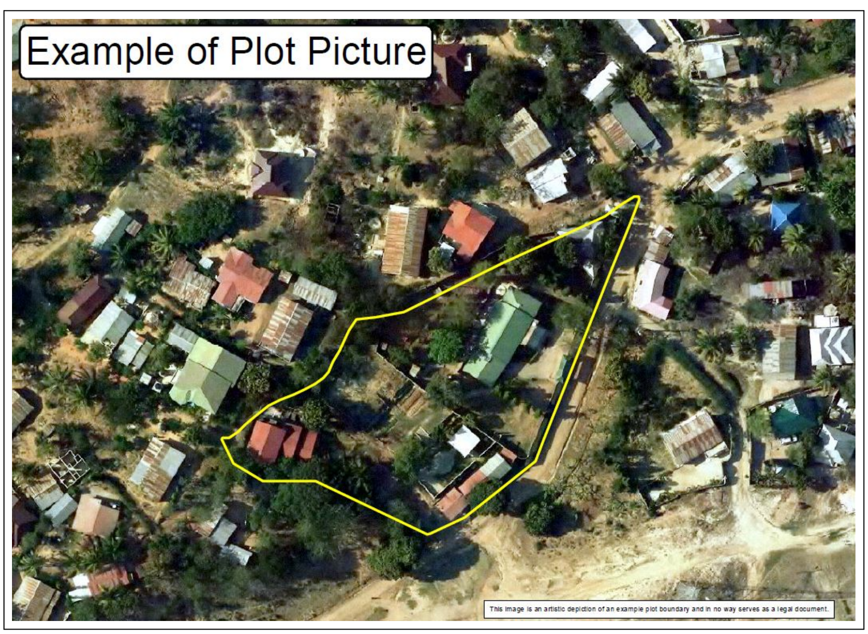

Figure 3.5: BDM Elicited CRO Demand and Elasticity

(a) Demand Curve

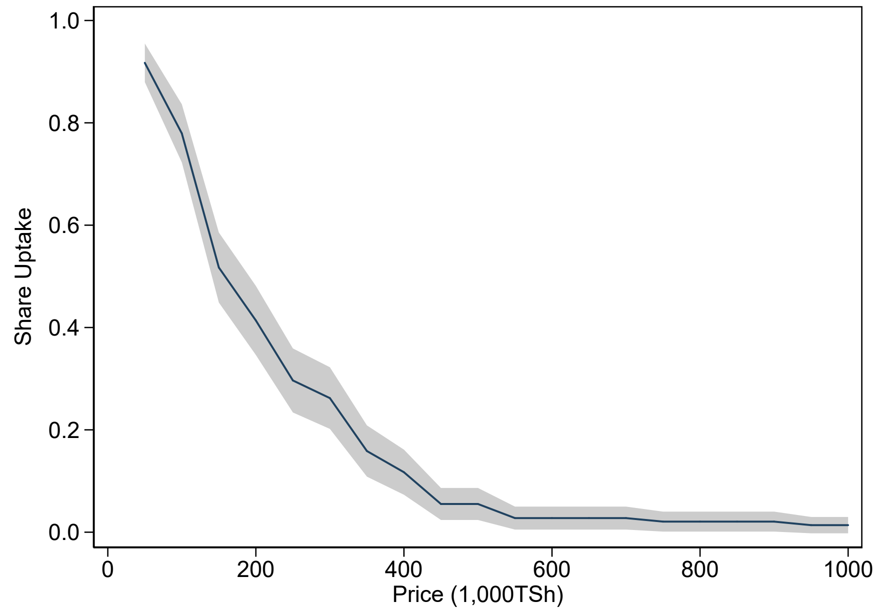

(b) Price Elasticity of Demand

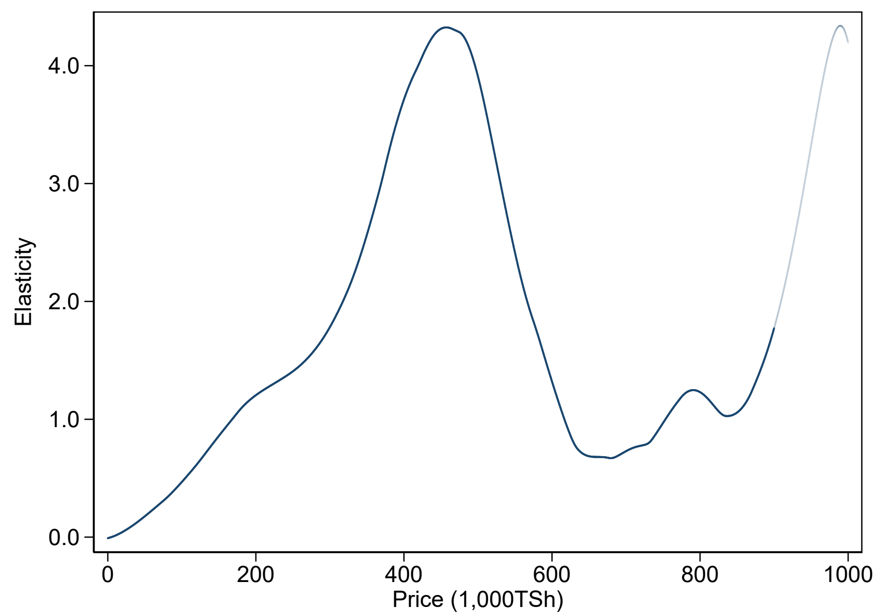

Notes: Figure 3.5a plots the BDM demand curve with 90% confidence bands. The demand curves indicate the share of respondents with a BDM bid greater than or equal to the indicated price. Confidence intervals are calculated using logit regressions (at prices TSh 50,000; 100,000; . . . ; 1,000,000) clustering standard errors at the plot level. The sample is 146 plots. Figure 3.5b shows demand elasticities using BDM predicted demand. The BDM elasticity is calculated by a local polynomial regression where, first demand is interpolated using a local polynomial regression with an Epanechnikov kernel, then the point elasticity is calculated and smoothed using a local polynomial regression. In Figure 3.5b to highlight the sparsity of data in the right tail of our data we lower the transparency over the range of the three largest observations used in the elasticity calculation.

Figure 3.6: Leader Elicited CRO Demand

(a) Control Group Leader Elicitation

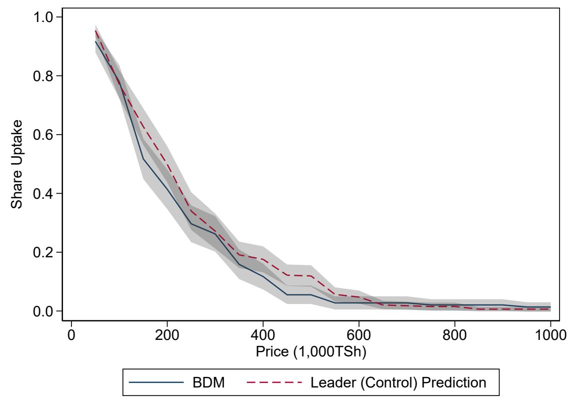

(b) Stakes Group Leader Elicitation

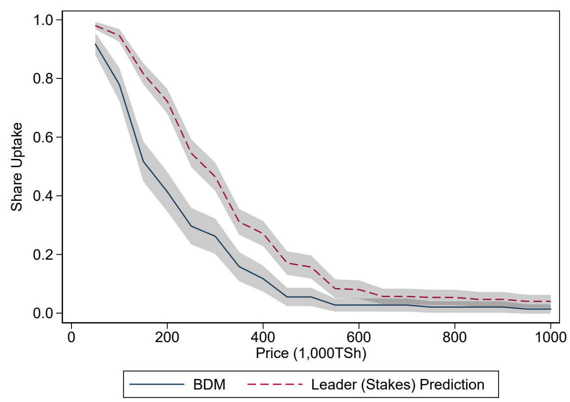

(c) Incentives Group Leader Elicitation

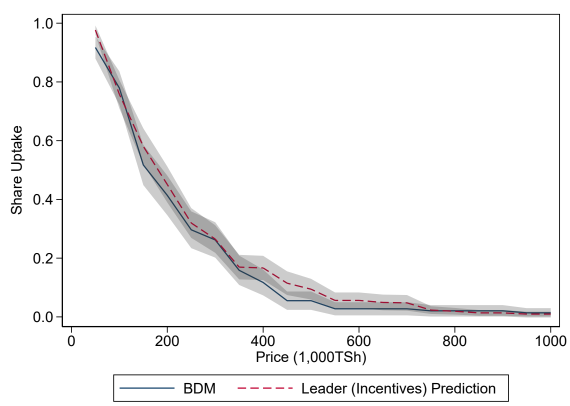

Notes: Figure 3.6 plots the BDM and Leader Predicted demand curves, with 90% confidence bands. The demand curves indicate the share of respondents with a BDM bid, or leader predicted WTP, greater than or equal to the indicated price. Confidence intervals are calculated using logit regressions (at prices TSh 50,000; 100,000; . . . ; 1,000,000) clustering standard errors at the plot level. The same sample of 146 plots are used for both, and predictions are frequency weighted by the number of leaders making predictions on that plot (i.e. each plot is equally weighted when calculating each leader predicted demand curve). Sub-figure 3.6a uses only leaders from the control group and compares the demand curve from their predictions with that of the BDM. Sub-figures 3.6b and 3.6c use leaders from the stakes and incentives groups respectively.

3.B Tables

Table 3.1: Owner and Leader Summary and Balance

Plots	(1) Mean	Leaders	(2) Mean	(3) Diff Stakes-Contr	(4) Diff Incent-Contr
Sole ownership	0.73 (0.037)				
Sole ownership and female	0.26 (0.036)	Female	0.40 (0.052)	-0.24* (0.124)	-0.15 (0.130)
Under 40 years old	0.23 (0.035)	Under 40 years old	0.07 (0.026)	0.06 (0.064)	0.04 (0.058)
Over 60 years old	0.20 (0.033)	Over 60 years old	0.33 (0.050)	0.09 (0.120)	0.11 (0.123)
Educ. primary or less	0.48 (0.041)	Educ. primary or less	0.57 (0.053)	0.05 (0.129)	0.05 (0.131)
Educ. above secondary	0.27 (0.037)	Educ. above secondary	0.12 (0.035)	-0.07 (0.087)	-0.06 (0.090)
Monthly income < 100,000TSh	0.34 (0.039)	Monthly income < 100,000TSh	0.19 (0.041)	0.25** (0.104)	0.00 (0.080)
Monthly income > 300,000TSh	0.35 (0.040)	Monthly income > 300,000TSh	0.42 (0.052)	-0.08 (0.127)	0.05 (0.132)
Avg. CRO quiz score	4.9 (0.114)	Avg. CRO quiz score	7.4 (0.124)	0.15 (0.309)	0.08 (0.318)
No children	0.08 (0.022)	Opposition party	0.14 (0.037)	0.03 (0.092)	0.00 (0.091)
Over 4 children	0.32 (0.039)	Assistant leader	0.39 (0.052)	-0.05 (0.126)	0.01 (0.130)
Absentee Owner	0.28 (0.037)	Owns their home plot	0.94 (0.024)	0.00 (0.064)	0.03 (0.058)
Acquired in last 6 years	0.11 (0.026)	Settled in last 6 years	0.07 (0.026)	-0.03 (0.056)	0.04 (0.074)
Acquired over 19 years	0.34 (0.039)	Settled over 19 years	0.38 (0.051)	-0.01 (0.127)	-0.06 (0.128)
Acquired by purchase	0.86 (0.029)	Home plot surveyed	0.91 (0.030)	0.07 (0.087)	0.17** (0.070)
Has sale certificate	0.25 (0.036)	Count of 15 owners known at all	12 (0.321)	0.54 (0.807)	1.1 (0.773)
Owns another plot	0.50 (0.042)	Count of 15 owners use services at all	4.3 (0.465)	1.7 (1.19)	0.33 (1.01)
Owns another surveyed plot	0.25 (0.036)	Count of 15 owners family members	0.22 (0.052)	0.09 (0.130)	-0.03 (0.114)
Owns another titled plot	0.10 (0.024)	Count of 15 owners close friends	1.4 (0.142)	-0.41 (0.358)	-0.08 (0.363)
Avg. invoice value (1000TSh)	526 (17.9)	Count of 15 owners religious affiliation	1.8 (0.267)	-0.48 (0.646)	0.07 (0.757)
Avg. plot area (sqm)	464 (32.8)	Count of 15 owners highly esteemed	1.3 (0.124)	0.02 (0.297)	0.08 (0.330)
Avg. BDM bid (1000TSh)	195 (14.5)				
N	146	N	90		

$^{*}p \leq 0.10$, $^{**}p \leq 0.05$, $^{***}p \leq 0.01$ for difference=0 t-test
Standard errors in parentheses

Table 3.2: Leader Predictions and Placebos

	(1) Income Rank	(2) Invoice Rank	(3) Invoice Rank Full	(4) Property Tax Paid	(5) Certificate of Sale
Panel A: Predictions					
Leader Prediction	0.20***	0.30***	0.34***	0.07**	0.09*
	(0.020)	(0.035)	(0.029)	(0.029)	(0.045)
N	871	871	1349	871	871
R^2	0.25	0.13	0.11	0.14	0.18
Panel B: Placebos					
Leader Prediction	0.19***	0.28***	0.32***	0.02	0.09*
	(0.035)	(0.058)	(0.046)	(0.050)	(0.047)
Stakes × Leader Prediction	0.05	0.06	0.08	0.06	0.03
	(0.048)	(0.086)	(0.068)	(0.065)	(0.057)
Incentives × Leader Prediction	-0.03	0.01	-0.02	0.09	-0.04
	(0.046)	(0.077)	(0.071)	(0.072)	(0.042)
N	871	871	1349	871	871
R^2	0.25	0.13	0.11	0.14	0.18

$^{*}p \leq 0.1$, $^{**}p \leq 0.05$, $^{***}p \leq 0.01$

Notes: Robust standard errors clustered at leader level in parentheses. Each observation is a leader-plot owner pair. Column 1 the dependent variable is the within neighbourhood rank of plot owner's income. The dependent variable in columns 2 and 3 is the within neighbourhood rank of invoice value. Column 2 restricts the sample to respondent owners, while column 3 includes all invoices. Column 4 the dependent variable is an indicator if the plot owner paid property tax in 2018. Column 5 the dependent variable is an indicator if the plot owner has a certificate of sale. The regressor is always the leader's prediction of the dependent variable. Fixed effects for leader strata, neighbourhood, and surveyor are included in all models.

Table 3.3: Leader's ability to distinguish variation in willingness-to-pay

	(1) WTP Rank	(2) WTP	(3) ln(WTP+1)	(4) WTP Percentile	(5) Top Rank	(6) Bottom Rank
Panel A: Predictions						
Leader Prediction	0.20***	0.10***	0.33***	0.23***	0.15***	0.24***
	(0.024)	(0.029)	(0.048)	(0.037)	(0.052)	(0.055)
N	871	871	871	871	871	871
R^2	0.19	0.09	0.10	0.09	0.03	0.12
Panel B: Distortions from real stakes						
Leader Prediction	0.175***	0.105	0.253***	0.159**	0.124	0.249***
	(0.041)	(0.081)	(0.095)	(0.069)	(0.085)	(0.093)
Leader Prediction	0.049	0.011	0.208	0.150*	0.043	-0.120
× Stakes	(0.058)	(0.091)	(0.133)	(0.089)	(0.125)	(0.130)
Leader Prediction	0.030	-0.025	0.089	0.098	0.019	0.103
× Incentives	(0.055)	(0.088)	(0.112)	(0.085)	(0.125)	(0.135)
N	871	871	871	871	871	871
R^2	0.19	0.09	0.10	0.10	0.03	0.12

$^{*}p \leq 0.1$, $^{**}p \leq 0.05$, $^{***}p \leq 0.01$

Notes: Robust standard errors clustered at leader level in parentheses. Each observation is a leader-plot owner pair. Column 1 the dependent variable is the within neighbourhood rank of plot owner's BDM bid. The dependent variable in column 2 is the value of the plot owner's BDM bid in Tanzanian shillings, and column 3 is the log value. Column 4 is the percentile rank across the entire distribution, rather than neighbourhood only. Column 5 the dependent variable is an indicator if the BDM bid is the highest in the neighbourhood, and column 6 indicates if the bid was the lowest in the neighbourhood. The regressor is always the leader's prediction of the dependent variable. Fixed effects for leader strata, neighbourhood, and surveyor are included in all models.

Table 3.4: Using Observable Characteristics to distinguish variation in willingness-to-pay

	(1)	(2)	(3)	(4)	(5)	(6)
	WTP Rank	WTP	ln(WTP+1)	WTP Percentile	Top Rank	Bottom Rank
Panel A: Invoice Formula						
Invoice	0.26^{***}	0.49^{***}	1.33^{***}	0.43^{***}	0.44^{***}	0.04
	(0.017)	(0.053)	(0.088)	(0.031)	(0.059)	(0.051)
N	871	871	871	871	871	871
R^2	0.25	0.37	0.15	0.20	0.19	0.09
Panel B: Invoice Formula and Leader Prediction						
Invoice	0.228^{***}	0.480^{***}	1.247^{***}	0.409^{***}	0.431^{***}	0.026
	(0.021)	(0.052)	(0.092)	(0.032)	(0.060)	(0.052)
Leader Prediction	0.151^{***}	0.053^{***}	0.253^{***}	0.179^{***}	0.062	0.239^{***}
	(0.023)	(0.015)	(0.043)	(0.035)	(0.044)	(0.055)
N	871	871	871	871	871	871
R^2	0.30	0.37	0.16	0.22	0.20	0.12
Panel C: Valuation						
Property Value	0.12^{***}	0.20^{***}	0.27^{***}	0.22^{***}	0.02	0.03
(1,000TSh)	(0.007)	(0.072)	(0.061)	(0.022)	(0.047)	(0.066)
N	870	871	871	871	871	871
R^2	0.15	0.07	0.09	0.09	0.01	0.08
Panel D: Valuation and Leader Prediction						
Property Value	0.075^{***}	0.171^{**}	0.229^{***}	0.195^{***}	-0.015	-0.011
(1,000TSh)	(0.012)	(0.073)	(0.059)	(0.023)	(0.054)	(0.064)
Leader Prediction	0.170^{***}	0.092^{***}	0.301^{***}	0.202^{***}	0.147^{***}	0.242^{***}
	(0.029)	(0.029)	(0.046)	(0.038)	(0.054)	(0.056)
N	870	871	871	871	871	871
R^2	0.21	0.09	0.11	0.12	0.03	0.12

$^{*}p \leq 0.1, ^{**}p \leq 0.05, ^{***}p \leq 0.01$

Notes: Robust standard errors clustered at leader level in parentheses. Each observation is a leader-plot owner pair. Column 1 the dependent variable is the within neighbourhood rank of plot owner's BDM bid. The dependent variable in column 2 is the value of the plot owner's BDM bid in Tanzanian shillings, and column 3 is the log value. Column 4 is the percentile rank across the entire distribution, rather than neighbourhood only. Column 5 the dependent variable is an indicator if the BDM bid is the highest in the neighbourhood, and column 6 indicates if the bid was the lowest in the neighbourhood. The regressor in Panels A and C are the invoice and property valuation equivalents of the dependent variable, respectively. While the regressor in Panels B and D are the leader's prediction of the dependent variable as well as the invoice and property valuation equivalents of the dependent variable, respectively. Fixed effects for leader strata, neighbourhood, and surveyor are included in all models.

Chapter 4

Life in a Slum: understanding living conditions in Nairobi's slums across time and space

4.1 Introduction

Countries across Africa are currently in a period of rapid urbanisation, with urban populations growing by an average 3.55 percent per year (UN Habitat 2016*b*). Kenya and its capital Nairobi are no exception; the rapid growth of the city, increasing from 2 million people to 3.1 million in the ten years between 1999 and 2009[1], is drawing attention to how the city can accommodate the millions of expected new residents. These new urban residents need housing, and if formal housing is underprovided many could end up living in slums. Already 56 percent of the urban population in Kenya are living in slums, an increase from 33 percent in Nairobi forty years ago (Lall et al. 2017).[2] A similar pattern is repeated across the continent, and while this issue is not unique to Africa, these numbers are far above the 34 percent of the urban population living in slums in Latin America and South Asia. As the process of urbanisation continues, both the role of slums in African cities and the conditions for those living in slums become increasingly important issues for policymakers.

In this paper we outline stylized facts about slums in Nairobi, revealing how their characteristics vary according to both location and time. In doing so, we are able to highlight the relative conditions for those living in slums compared to both formal and rural areas. We

[1] Authors' calculation

[2] This number is from page 38 and is comparable across others within the same report. However, as noted in Section 4.2 below, multiple definitions of slums exist. In our analysis of Nairobi, we observe a smaller share of the urban population living in slums, as we use a tighter definition.

provide descriptive analysis of who lives in slum areas, the living conditions they face, and the access they have to jobs, healthcare, and education. We analyse how these vary by location and how these have changed over time to highlight the dimensions along which slums are improving or deteriorating. Together, these provide a better understanding of living conditions in slums today, and the priorities for policymakers preparing for increasing slum populations over the coming decades.

4.1.1 How to conceptualize slums?

Slums are hard to define. While Programme (2003) and others have tried to establish a standard definition (an area where households lack at least one of the following: access to improved water, access to improved sanitation, sufficient living area, durability of housing, or security of tenure), the reality is that the local context matters and different countries adopt their own definitions. By their very nature, slum areas are also likely to change quickly, as the low-cost housing is simple to build, demolish, or improve. As a result, there are few studies that look at the dynamics of slum areas within cities, or at the characteristics of slums across cities. Data limitations typically prevent comparisons over both time and space. In this paper, due to the availability of detailed census data at two points in time, we are able to focus on the evolution of slums within one city.

Many of the productivity benefits that exist in cities result from their density: within a relatively short travel time, firms can access a wide set of potential employees and consumers, and households can buy goods and services from a large choice of firms. The average population density of slums in Nairobi was 28,200 people per km2 in 2009, 51 percent higher than just 10 years previously and far higher than in formal residential areas. These slums could therefore potentially be an efficient use of urban land providing high-density, low-cost accommodation near other households and near markets. However, there are also many reasons why slums may be an inefficient use of prime urban land. First, with low-quality one- or two-storey buildings and little public transport provision, their density leads to crowdedness. This both decreases liveability within slum areas and limits accessibility, reducing the potential benefits that density provides. Residents have little access to urban areas beyond the slum in which they live, leading to low mobility and preventing residents from accessing jobs and educational opportunities elsewhere within the urban area (Salon & Gulyani 2010). Second, dense areas are also subject to large externalities across households. Gollin et al. (2017) report higher crime rates in highly dense areas, and the risk of transmission of communicable diseases increases (Sclar & Carolini 2005). These externalities are worsened if there is underinvestment in services, with a lack of access to clean water and sanitation, in particular, having large negative health consequences; for example (Duflo et al. 2012) report that 88 percent of diarrhoea infections are attributable to poor sanitary infrastructure, unsafe hygiene procedures, and a lack of clean water. Finally, underinvestment in private living conditions also affects life quality. Cattaneo et al. (2009) show substantial impacts on child health from replacing dirt floors with cement floors and also find positive effects on adult well-being measures.

Most of these potential negative impacts of slums stem from underinvestment, both in private and public services. Low levels of investment in slum areas may arise for multiple reasons. It may simply be that households have a low willingness to pay for services. However, for many services such as water and sewage systems, installation costs are high and require coordination among users to pay for the investment or state provision. In addition, if land rights are weak and the owner or renter is uncertain that they will hold the rights to their building in the future, the incentive to invest in the local environment is reduced. Land market failures may also lead to the land being inefficiently allocated to those who may not value it the most. Finally, there may be political and economic reasons why the political elite have no incentive to improve living conditions within slums. As a result, local poverty traps could form, where residents are unable to benefit from the opportunities of the city (Marx et al. 2013).

4.1.2 Our investigation

Using detailed spatially disaggregated census data from Nairobi, we draw out stylized facts on who lives in slums, their wealth, the quality of buildings, the access to services, and child health and education outcomes. We examine how these may vary across both space and time within the city and compared to rural areas. The data are a repeated cross-section of individuals, so we cannot track people over time. We can, however, track the characteristics of neighbourhoods within the city, revealing how those living in slum areas differ, and how their living conditions are changing.

Slums are incredibly dense areas, with those near the city centre approximately ten times as dense as formal residential areas in the same part of the city. Visually, this density combined with poor-quality buildings gives an appearance of neighbourhoods stuck in poverty. Households are smaller and the number of rooms per person lower than in formal areas. Access to services including improved water supplies and sewage disposal are both very low, alongside ownership of assets. Along other dimensions, however, the picture is not as stark, particularly when we observe the changes over time. First, while housing quality and access to water and sanitation services is worse than in formal areas, it is significantly better than in rural areas. Second, there is some catch-up in housing quality between slum and formal areas over time—concentrated, however, in more distant slums. Third, in terms of access to electricity, as well as schooling and child mortality, the improvements seen over the last decade have resulted in slum households catching up with the levels seen for households in formal areas just 10 years previously. Slum areas are not static areas of poor living conditions, but instead are areas that are experiencing dramatic changes. Identifying both why these changes are occurring only in some dimensions and how these changes can be supported are key questions for research and policy to address.

Section 4.2 outlines the context of our analysis, including why Nairobi can inform policy-makers working on cities across the region, and the specifics in the history of Nairobi which may have influenced its slums. How we measure slum characteristics; adult demographics, private investment in building quality, public investment in services, and health and educational outcomes for children, and their changes over time and space is discussed in Section 4.3. We highlight the main empirical difficulties in a study of this kind. Section 4.4 then draws out stylized facts from Nairobi that address the issues outlined above. In section 4.5, we discuss how these issues are relevant to policy-makers, before concluding.

4.2 Nairobi context

Nairobi is home to over 3m people, including 180,000 in perhaps the most well known of African slums, Kibera, which is located just a few kilometres from the centre. Numerous other slums spread throughout the city, as shown in Figure 4.1, varying, among other things, in size, density, access to the centre, access to main arterial roads, and land rights.

Our analysis is conducted using repeated cross-sectional data on individuals using the 1999 and 2009 population censuses, giving a 10 percent sample of the city. Over this 10-year period, Kenya and Nairobi experienced major changes. First, national GDP grew quickly, at an average rate of 3.8 percent a year (Feenstra et al. 2015). Second, the urban population grew simultaneously, increasing by 50 percent in the same period. The rapid growth of Nairobi is not, however, the complete story of the city over this time period. It was also a decade of upheaval, some of which was focused specifically in slum areas. A wave of violence in Nairobi's slums following the 2007 elections may have had some influence on the outcomes that we study. This period of violence was hardly an unprecedented random shock, as Cheeseman (2008) notes, citing severe ethnic conflicts in the early and mid-1990s. We do not attempt to disentangle the effect of this violence from other channels that may have affected changes in slums. If slums are, for example, more prone to politically induced ethnic violence, the impact of this on slum outcomes is included within our results. We are able to document changes in slum conditions, but further work is required to reveal the mechanisms behind these.

Since change depends on the institutional and legal situation in slum areas, as well as the characteristics of these areas and their location relative to urban markets, the legal and institutional complexities of Nairobi are worth exploring. First, the legal ownership of land in slums is frequently under dispute. The example of Kibera illustrates how convoluted claims on property rights can be reinforced over time and potentially trap slums into areas with little investment. At the turn of the twentieth century, the British Colonial government enrolled Sudanese soldiers to serve in the King's African Rifles and fight for the British. These Nubian soldiers were granted permission to settle on the land that is now Kibera, and in spite of some pressure for them to leave, in 1950 they were granted official permission to stay. After Kenyan independence, however, the government revoked all Nubian claims to land in Kibera. Finally, a motion to provide the Nubians with land titles was passed in parliament, although the government continued to ignore their claim and the motion was never implemented (Joireman & Vanderpoel 2011). Currently the Nubians in Kibera continue to live on land that is formally considered property of the government, but over which they argue they have a legal claim. In addition to the Nubian settlers, since independence in 1964 migrants have flooded into Kibera, all settling without official government recognition. These settlements are facilitated by local chiefs who are government representatives but have no authority to grant land titles. The chiefs take payments for the construction of structures and allow the builders to act as landlords (Joireman & Vanderpoel 2011). As a result, the decisions today regarding who should be granted titles, or how to allocate compensation for evicted residents, are matters of dispute.

Second, in addition to legal complexities, misplaced incentives in the government can further distort the development process of slums. There is a long history in Kenya of government corruption pertaining to the use of land. Seminal work by Klopp (2000) argues that in the 1990s the Moi government was losing influence due to the decline of traditional forms of patronage, greater international scrutiny of corruption, and more political competition. In response, the government increasingly used the allocation of public land as a form of patronage. In 2004, a report headed by Paul Ndungu was published, researching the unlawful allocation of public lands under the previous administrations. Southall (2005) summarizes the findings in the Ndungu Commission's report, noting that, on urban land specifically, there was found to have been widespread abuse of presidential discretion in making land grants to individuals for political reasons, without consideration to the public interest and without proper pursuit of legal procedures, including outright illegal allocation. Given the context of public land allocation in Kenya, it is not hard to believe that slums on public land in Nairobi are being used as a means of patronage, and governmental power is abused for private gains.

Finally, in Nairobi, slum dwellers are not simply squatters, they pay rent to slum landlords even if the land is not owned privately. In a survey of 1,755 Nairobian slum households in 2004, Gulyani & Talukdar (2008) found that 92 percent were rent-paying tenants and only 6 percent owned both the land and structure. Furthermore, they estimated very short payback periods on housing investment in slums, highlighting the high costs to tenants of living in slums relative to the housing value. Such high returns to the landlords are often indicative of high entry barriers, which in Nairobi include the need for political connections, the payment of significant illegal fees, and a willingness to bear the risk of demolition and loss of capital. These slum landlords obtain security of their investment through political mechanisms, which are obtained despite the absence of formal legality, enabling a process that Amis (1984) calls the 'commercialization' of shanty towns. Specifically, a survey of landlords in Kibera by Syagga et al. (2002) reveals that many are government officials and politicians (as cited in Gulyani & Talukdar (2008)).

4.3 Empirical strategy

The analysis conducted in this paper uses individual- and household-level data from a repeated cross-section. Households and individuals in the 1999 and 2009 population censuses are grouped into comparable geographical units, enumeration areas (EAs), which are identified as slum or formal. We look at the association between various demographic characteristics and child health and education at the individual level, as well as housing, assets, and access to services at the household level, with whether or not that individual (household) is in a slum area. We conduct analysis using the individual level (household level) using OLS as follows:[3]

$$Y_i = \alpha + \beta D_i^{Slum} + \gamma D_i^{2009} + \delta D_i^{Slum\&2009} + \theta X_i + \epsilon_i \tag{4.1}$$

where Y_i is the outcome variable of interest for the individual (or household), typically a binary variable, such as whether or not a household has a solid stone or brick wall, although

[3] We also estimate the same analysis using probit estimations to check that we have similar results. For simplicity, and because our key explanatory variables are binary, we prefer the OLS estimations: we are in effect just estimating the sample means.

sometimes it can be a continuous variable, such as the child mortality rate. On the right-hand side, D_i^{Slum}, D_i^{2009}, and $D_i^{Slum\&2009}$, are dummy variables signifying respectively if the individual (household) lives in a slum or not, if the year of the observation is 2009, and if the individual (household) lives in a slum and the year is 2009. This allows us to separate the general trends in outcomes across the whole city over time, as well as the incremental changes felt only within slum areas. The marginal effects from these estimations are plotted, along with the 5 percent critical values to show the different means of the outcome variable concerned in slum and formal areas and in 1999 and 2009. Standard errors are clustered at the EA level. The average levels for these outcome variables in rural areas are also plotted for comparison.

We are also interested in how the correlation between living in a slum and our outcomes varies according to other characteristics. This can be captured using controls, Xi, which may include other characteristics of the individual, such as education level, and other characteristics of the area, such as distance to the city centre or land ownership of the plot. However, the inclusion of these controls leads to no substantial differences in the results. We allow therefore for a more general analysis, interacting the right-hand-side variables, D_i^{Slum}, D_i^{2009}, and $D_i^{Slum\&2009}$, with distance to the city centre, distance to the nearest industrial site, tenure status of the household, and individual slum identifiers. This captures heterogeneous relationships between slums and living conditions across space, tenure status, and slums. We discuss the results for all of these except distance to the industrial site, as there is little variation of interest along this dimension.

It is important to note here that the estimated coefficients show correlations: are people who live in slums faced with worse buildings, services, and social outcomes? They do not capture a causal impact of slums on these outcome variables, as the presence of sorting of households across areas within the city means that many, if not all, of our outcome variables are endogenous.

4.3.1 Issues and approach

There are inherently several major issues faced when analysing slums. First, there is no standard definition of slums, and different governments, non-governmental organizations (NGOs), and other bodies may all define slums differently. This paper uses a static definition of slum boundaries in Nairobi provided by the Kenya National Bureau of Statistics (KNBS) and created in 2009 as an indicator for informal settlements. From conversations with employees at the KNBS, EAs are defined as slum if they are 'unplanned'. Other slum definitions are available for Nairobi, including a mapping by the Center for Sustainable

Urban Development (CSUD) at Columbia University as part of a 2004 landuse map (see Williams & Klopp (2014) for their full methodology) and another by Limited & Consultants (2013) under the Kenya Informal Settlements programme. Although the criteria for these alternate definitions are similar, the lack of a coordinated effort to map slums through time has resulted in inconsistent mapping. For comparable measures across cities, greater conformity is required.

Second, very little quality data on slums in Africa exists. Surveys can provide very detailed measurement of outcomes in slums, but are both time-intensive and costly to run, which often limits their extent both across time and space. Another source of potential data is censuses, which is what we use in this paper. This has two main advantages. Census data capture the full population of an urban area, perhaps with 10 percent surveyed in greater detail, allowing comparisons across different areas within the city. Population censuses are also typically repeated regularly, allowing comparison of the same areas of cities over decades. However, while the full population is in theory recorded, it is often difficult to gain access to the data at a fine level of spatial disaggregation. When data are grouped into large administrative units, much of the variation in access and land ownership patterns is lost. The questions are also less exhaustive than in surveys, and less tailored to the economic questions concerned. There is also a growing literature that attempts to circumvent the difficulties posed by a lack of conventional data in slums. For instance, satellite imagery has been used in a variety of ways to measure otherwise intractable aspects of slums and other fine-scale areas in developing country cities. Marx et al. (2019) use a series of very high-resolution satellite images to measure the brightness of tin roofs in a slum as a proxy for household investment, and Henderson et al. (2016) use digital tracings from aerial imagery to measure built volume across the whole of Nairobi.

The third source of difficulties in assessing the development of slums is understanding the population being referred to. Over time, residents move in and out of any given area within a city. These movements may be in response to the conditions they are facing. First, the current residents of slums have selected into those areas. For example, a correlation between whether somebody lives in a slum area and the wage they earn may show that slum dwellers earn lower wages than other city residents. However, care must be taken to ensure that this is not attributed as a causal effect. The slum may provide lower access to jobs compared to other neighbourhoods, but the slum also attracts those who do not have a high paying job owing to cheaper rents. Any analysis of slums has to be aware of this sorting across neighbourhoods, and ensure that the differences between slum and non-slum populations are not attributed uniquely to the conditions in slums. Second, when using repeated cross-section data where individuals cannot be traced over time, the movement of households is lost. For example, if a slum area provides interim housing for new arrivals in a city, then when households successfully land a decent job and save up enough money, they may move on and settle in a different neighbourhood. If this were repeated, with a new arrival replacing them in the slum area, the characteristics of those who live in slums would appear to be static over time. Those who benefit from slums then move on and no longer appear in the cross-sectional data as slum residents. The slum conditions are such that there are positive benefits for these people that would not be captured by this analysis. The analysis in this paper is only comparing conditions within slums themselves over time, and against formal areas.

4.3.2 Data

This paper makes use of publicly available microdata from the Kenya Population and Housing Censuses for 1999 and 2009. The data are two repeated cross-sections at 10-year intervals, with the location of households recorded at the EA level, of which there are over 4,000 in the city, a level of precision rarely available to researchers. In general, the spatial extent that we examine is the administrative area of Nairobi County, which stays constant from 1999 to 2009. For our analysis along the distance to centre, we restrict focus to EAs within 2–15km from the city centre so that each kilometre bin contains at least 100 slum households and 100 formal households. We are limited by the questions the census records. For example, data on the type of job residents perform are limited, and information on water provision is coded differently in the two periods, preventing comparisons over time.

The KNBS slum definition assigns 2009 EAs as slums or not, and we are then able to classify 1999 EAs as slums or not depending on their spatial relation to the 2009 EAs. For our main set of results, we classify 1999 EAs as slums if at least half of their total area intersects a 2009 slum.[4] We run robustness checks using a range of different cut-offs and find similar results. Holding the slum area constant over time ensures we are comparing the same areas in the analysis. However, it does ignore the fact that some new slum areas may have emerged or, indeed, disappeared between the two time periods. New slums

[4] A complication that arose in mapping the census data spatially was in producing maps for the 1999 census. Previously, the highest spatial resolution available for mapping census data in Nairobi was to the 110 sub-locations (SL) in the city. Since many spatial characteristics, including slums, of urban areas manifest at a finer scale, much of the heterogeneity in the data would then be lost. In order to overcome this difficulty, we photographed the original hand-drawn census enumeration maps from the KNBS archives in Nairobi, a total of 194 maps which detail all EAs and some households in Nairobi from the 1999 Kenya Census. We digitized the boundaries and identifier numbers of all 4,636 EAs. After the digitization process, the census microdata were readily merged, allowing us to give highly precise locations of each surveyed household based on their EA identifier.

potentially have different characteristics, and are not captured in this analysis. Figure 4.2 shows the location of slums within Nairobi, alongside the 1999 EAs.

4.4 Stylized facts

4.4.1 Who lives in slums?

(a) Slums have population densities up to ten times higher than in nearby formal areas.

(b) Slum residents are more likely to be male and living in smaller households than those in formal areas. This is particularly true near the city centre.

(c) While slum residents are less well educated than the rest of the city's population, their level of education has increased over time.

(d) Slum residents are more likely to be in paid jobs, particularly if they live near the centre; this falls sharply in neighbourhoods further from the centre.

Slums are spread throughout the city of Nairobi, as shown in Figure 4.3 below. Most of the city's slum residents live within 3–12km of the centre, with few at the very heart of the city. One of the most striking features of slums is their dense living conditions and in Figure 4.4 we document the population density in slum and formal areas by distance to the CBD. The differences are large: between 3km and 6km from the centre, slums are approximately 10 times as dense as similarly located formal areas with a population density averaging over 60,000 people per km^2.

In Figure 4.5 we plot the sample means of population and household characteristics across slum and formal areas. Just 43 percent of those living in slums are female, an increase from 1999 but still below the natural average of 50 percent. In Figure 4.6, we observe how this varies by the location of the slum. Nearly 60 percent of adults living in slums near the city centre are male, a reduction from 1999 but still far above the natural average. An equal male/female ratio is only reached at 12km from the centre. Within formal areas, the male/female ratio does not vary by distance to the CBD. Households in slums are also smaller than in formal areas, with 2.9 people per household compared to 3.4. Again, Figure 4.6 shows how this differs by distance to the CBD. In formal areas, households are slightly larger near the urban core; however in slum areas, central households are approximately 20 percent smaller than their less central counterparts.

We also note that those living in slums have lower levels of education. The adult population has a 51 percent rate of some secondary education in 2009 compared to over 71 percent in formal areas. This share has, however, increased rapidly over the preceding decade, by 11 percentage points in slum areas and 14 percentage points in formal areas. This compares to an increase of just 3 percentage points in rural areas, where the secondary education rate is just 12 percent in 2009. This shows that the adult population in slum areas is dynamic; either new, higher-educated people are moving in over time, or the previously resident population is becoming increasingly educated. Figure 4.6 shows that this increase in education happened in slums throughout the city except on the fringes, indicating that the effects are not due to sorting of individuals across slums, for example through the higher-educated relocating into better slums.

We also show that slum residents are slightly more likely to work for pay than those in formal areas, although the difference is small, at 5 percentage points. This is particularly the case for those located near the city centre, where approximately 50 percent of slum residents work for pay, falling to 35 percent near the fringes of the city. This large decline highlights the importance of residential location; central slums offer better access to paid jobs, and those living in more distant slums have a decreased chance of working in paid employment. Unfortunately, more detailed data on the types of jobs performed are unavailable.

Together, these facts show slums to be dense areas with predominantly male populations living in small households. While education levels are lower than the citywide average, adults in slums are becoming increasingly well-educated over time and those nearest the city centre are more likely to be accessing paid jobs. These facts support a story of male family members being sent to the city to find work and earn money, living in small households near to urban labour markets. However, in Figure 4.5 we also show that, of those not born in the same district, those living in slums in 2009 have been there on average 8 years. While slightly less than in formal areas, this shows that the location decisions are not temporary. Men move into slums and remain there; they benefit from access to urban jobs and education, but we cannot rule out that they are stuck in these areas of urban poverty.

4.4.2 How wealthy are slum residents?

(a) Material wealth, as proxied by household goods ownership, is lower in slum households than in formal households, but still higher than in rural areas.

(b) Centrally located slum households are materially poorer than more distant neighbourhoods, suggesting a trade-off between rents and asset ownership. In contrast, centrally located formal households are materially richer than more distant neighbourhoods.

Figure 4.7 show the levels of asset ownership at a household level in slums and formal areas, estimated using OLS, and the mean levels of ownership in rural areas for comparison. These data were only collected in 2009 so we cannot observe how asset wealth has changed since 1999. Lower levels of asset ownership are observed across the board in slum areas versus formal areas, and particularly for televisions, fridges, and computers, all of which require a steady electricity supply. Compared to rural areas, however, households in slums have high rates of household good ownership. The only exception is bicycles which tend to be more common in rural areas where they may complement rural living or substitute for poor public transport. Despite the poverty in slums, those living in these urban areas can better afford these material goods compared to their rural counterparts.

In Figure 4.8 we examine good ownership across space. Within formal areas, households nearer the CBD have higher rates of asset ownership. Among households in slums, this trend is reversed; households that are further from the CBD have a higher probability of owning a radio, TV, computer, bike, or fridge. This suggests a trade-off between location and material wealth: living nearer to the centre is more costly and so, in choosing to do so, slum households may substitute away from buying other assets. The only good that does not follow this pattern is the mobile phone. While the probability of phone ownership varies across space, no clear pattern emerges. Together, the evidence suggests that those living in slums are materially wealthier than their rural counterparts, but that those in the central areas are materially poorer than those further out, which could either be due to lower incomes or to a trade-off between rents and asset ownership.

4.4.3 What does housing look like in slum areas?

(a) Slum housing is of poor quality, however some catch-up between slum and formal areas is observed over time.

(b) This catch-up varies spatially, with improvements in building structures more likely further from the CBD.

(c) There is little to suggest this varies according to ownership status. City residents are predominantly renters, both in slums and formal areas.

Figure 4.9 documents how living in a slum area correlates with building quality in Nairobi. First, slums have distinctively worse building structures, particularly with regard to whether they have a solid wall, with only 26 percent of households living in buildings with a solid wall versus 84 percent in formal areas in 2009, and in terms of solid roofs, with 9 percent of households in buildings with a tile or concrete roof compared to 45 percent in formal areas. It is interesting to note that the differences are less striking for floor material, with a much higher proportion of slum households having cement or tile floors. Flooring is a cheaper and more flexible investment than the main structure of a building, so it may be easier for slum dwellers to invest in.

Over the preceding decade, slum areas experienced a greater increase in the likelihood of having a solid wall or floor compared to formal areas. There was a 4 percentage point rise in the likelihood of a household having a tile or concrete roof in slum areas and a 9 percentage point rise in stone or brick walls in slums—both double that seen in formal areas, as well as a 11 percentage point rise in cement or tile floors, over three times the rise seen in formal areas.[5] When compared to levels in rural areas, in all three indicators housing in slums is of higher quality than the average housing in rural areas, and has improved at a faster rate over the last 10 years.

Second, we note that slums are crowded, with 0.6 rooms per person compared to 0.83 on average in formal areas. While in formal areas the number of rooms per person has remained constant over time, it has decreased in slum areas. Combined with the poor quality housing and high population density, this results in lower liveability in slum areas. On this indicator, slums perform worse than rural areas.

We also observe that slum households are less likely to own the house they live in, though urban formal households are also very likely to be renters. In 2009, just over 10 percent of slum households owned their own property, compared to over 15 percent in formal areas, far below the levels in rural areas. This proportion has decreased over the last decade. The measure unfortunately does not capture the underlying land rights—for example, whether the owner has secure tenure. However as renters may have lower incentives to invest in their property and surroundings compared to permanent owners, it may reveal part of the story behind underinvestment in slums.

We then look to see how these levels and changes may vary according to the location of the slum. Figure 4.10 shows that the quality of buildings in slums is inferior to that in formal areas at all distances from the CBD. However, beyond approximately 7km from the centre the gap between formal and informal areas narrows. In terms of the size of the building, slums have on average 0.6 habitable rooms per person up until around 10km from the centre, at which point the space per person increases and falls in line with that seen in formal areas. This highlights a trade-off between living conditions and city ac-

[5]For studies that rely on satellite imagery to detect changes in slum areas, this point is important to note. The evidence suggests that improvements in housing structures are taking place, without a corresponding improvement in roof quality. Using roof changes to estimate investment in buildings may underestimate the effects.

cess: poor-quality housing and a lack of space defines slums in prime locations near the CBD, and higher-quality housing and more spacious properties are available further out. Consequently central slums appear visually to have particularly low living conditions. Interestingly, for formal areas this pattern is somewhat reversed, with more spacious housing and a higher proportion of houses with solid roofs near the centre.

Over time, while houses throughout the city have seen similar increases in the share with a solid cement or tile floor, improvements in the wall and roof structures have only occurred at these more distant locations. This suggests that more-distant slums are improving faster, perhaps because there is more space for new buildings, or because the land rights are such that it is easier to re-develop areas. Unfortunately, we cannot disentangle the tenure status within slums; however, in Figure 4.11 we examine how each of these characteristics varies according to the ownership status of the household. No clear pattern emerges, except that privately owned properties are larger. This differs from our priors, as we expected areas where the occupier has less security over their property to be areas of lower investment. However, as we cannot distinguish the ownership rights of the landlord, we do not know if there are varying effects according to the strength of the landlord's tenure security.

We also repeat the analysis looking at individual slums and observe substantial variation. Some slums have nearly no housing made from stone or brick walls, whereas in others over 60 percent of housing is. No clear patterns arise: some areas, such as Mukuru, have a low share of households with solid roofs or walls, but a high share with cement floors, and others, such as Kayole, have a high share with solid walls, roofs, and floors, and larger properties than average, but have seen a large deterioration in each of these over the decade as the slums have grown and densified. This highlights the specificities of different slum areas. Various local characteristics including, but not limited to, location and tenure status will influence the development of slums and the investment in housing. Policy-makers both within Nairobi and beyond need to be aware of this local context.

4.4.4 How good is the access to services in slum areas?

(a) Slums have worse access to piped water, sewage systems, and electricity compared to other urban areas, but far higher rates than rural areas.

(b) Improvements are noticed over time; however, for sewer systems, the rate of improvement is not enough to dramatically alter the gap between formal and informal areas.

(c) Access to electricity in slums has improved, particularly electricity for lighting.

(d) Variation across slums is large, but does not systematically correlate with the distance to the city centre.

Figure 4.12 shows the average rate of access to services in slum and formal areas, highlighting the differences in level of provision. While 83 percent of formal households have access to piped water in 2009, and 78 percent to improved sewage disposal, just 63 and 25 percent, respectively, of slum households have such access. As outlined above, clean water supplies and proper sewage disposal are central to limiting the spread of contagious diseases and health outcomes. Small improvements in sewer access over time are present in slum areas, increasing by 7 percentage points compared to 4 percentage points in formal areas since 1999. Unfortunately, the data do not permit intertemporal comparisons on piped water access. In rural areas, just 15 percent of households have access to piped water and 1 percent to sewers, showing that slums, while areas of service deprivation within the city, provide far better access to water and sanitation than rural housing.

We also note gaps in the use of electricity, particularly the low levels of electricity use for fuel in slums. The use of other fuels for cooking may impact local air quality and pollution in slum areas. The rate of improvement in access to electricity in slums is particularly striking. While in formal areas the use of electricity or gas for fuel has risen by 10 percentage points, not far below the rate of increase for light of 15 percentage points, in slum areas the use of electricity or gas for fuel has risen just 4 percentage points in 10 years, compared to 28 percentage points for the use for light. This is consistent with a story of increasing availability of informal electricity supplies, with power sufficient for lighting but not for heating or cooking. Having access to lighting is beneficial to households by increasing the useful hours of the day for work and leisure time, but the limited increase in usage for fuel suggests these electricity networks are unsuitable to bring all the potential benefits of electricity. Anecdotal evidence suggests that local secondary electricity markets exist whereby a local seller connects to the grid and sells small amounts off to households through their own wiring system. Such a grid would have low capacity, consistent with our evidence showing higher electricity usage for lighting than for major household tasks such as cooking.

Figure 4.13 documents how this varies according to distance to the CBD. We note first that within slum areas access to piped water and sewers decreases with the distance. Access to electricity, however, does not. If the former require more government support to construct supply networks, this would suggest more attention historically has been placed on central slum areas. This could be due to higher population densities leading to increased cost-effectiveness for water provision, or because these areas are politically more sensitive, or because they are wealthier—unfortunately, a question we are unable to answer in this analysis. In terms of changes over time, we note that improvements occurred nearly uniformly throughout the distance bins. As these investments require some degree of coordination, one might expect that they are easier in denser, older slums near the CBD. Alternatively, one may also expect that it is easier to provide services in newer, less dense slums where there is space to construct new networks with less disruption. Neither of these two effects is dominating.

On a slum by slum basis, there is again substantial variation. Kayole, mentioned above to be the slum with the highest building quality but rapid decreases in this quality over time, has also seen some of the largest decreases in access to different services over the decade. The slum has grown and densified, and service provision has failed to keep up. Interestingly, slums, such as Uthuru, that have high levels of access to piped water do not always have good sanitation, and similarly slums with improvements in sanitation services are not the same slums that have seen improvements in electricity access. Again, this highlights the importance of the local circumstances in understanding the provision, or lack of provision, of key local services.

4.4.5 How well are children doing in slums?

(a) School attendance rates have increased and converged across formal and slum areas.
(b) These effects occurred throughout the city, ensuring all areas have similar access to primary schooling and higher access than in rural areas.
(c) Child mortality remains higher in slum areas, but again substantial gains have been made. The rate of child mortality varies substantially by slum.

Figure 4.14 shows that the levels of child school attendance are nearly identical across slum and formal areas in 2009, but that child mortality remains higher in informal settlements, at 76 children per thousand versus 54 children per thousand in formal areas.[6] Both of these measures show considerable catch-up over time and when compared to rural areas: 94 percent of children attend school in slum areas compared to 84 percent in rural areas today. The increase in access to schooling is likely to be at least in part attributable to the Free Primary Education Programme, brought in during the early 2000s, midway between our two census rounds, which offered all primary-age children free schooling. Previously,

[6] Note that child mortality is calculated using the number of child deaths per live births, without any age threshold. This makes it not comparable to standard measures that set, for example, a threshold at 5 years of age.

school fees had been a major barrier to entry to education (Glennerster et al. 2011), but the removal of this cost has led to near universal primary education. The lower levels in rural areas are likely due to access issues; within a dense urban environment, the nearest school is typically closer than it would be in a rural area. Unfortunately, nothing can be said about the quality of this education. From Figure 4.15 we observe that this impact was felt throughout the city; no matter where a child is living in the city there is near universal access.

In terms of child mortality, slum areas have also seen substantial catch-up. Slums have reached levels throughout the city in 2009 that are lower than those seen in formal areas just 10 years previously. However, a gap still remains. The improvements in child mortality vary wildly between slums, suggesting different local forces are at play here. In part, this may be explained by the lack of services, such as water and sewer systems, as discussed above. Figure 4.15 shows child mortality is higher in more central areas, areas which we have already shown to be dense and with poor-quality buildings, both factors in the rate of spread of disease.

4.5 Discussion and conclusion

The stylized facts above reveal slums to differ vastly from more formal areas within the city of Nairobi. Some of these characteristics are true across all slum areas, be it that slum residents are less well educated, live in smaller spaces, and that their housing is of lower quality and has poorer access to services such as water and electricity. However, looking more deeply at the data, we observe that there is variation across space, and across specific slums. The variation across space highlights the issue of using local surveys to draw general conclusions on slums, as the current conditions in slums and the pace and direction of change over time are both very localized. Some of these trends appear to vary in a consistent spatial way, such as that those living in more central slums are 15 percentage points more likely to be in paid employment than those on the urban fringes. Other variations are not explained by the data and therefore must be assumed to depend on other local characteristics, such as land tenure, security, political will, or proximity to certain urban services such as health clinics or NGOs. What we learn from one slum need not apply to another, and evidence from slums in Nairobi need not necessarily apply to slums throughout the East African region or beyond.

We can, however, draw out some conclusions from these Nairobi facts, and some questions that require further investigation. First, slums are difficult places to live in, with density in some slums ten times that in neighbouring formal areas, low levels of space per person, poor building quality, and poor access to services. Yet compared to rural areas, on many of these dimensions individuals living in slums are actually faced with better conditions. Second, slums are changing, both in terms of demographics and buildings and services. Between 1999 and 2009 individuals in slums became 11 percent more likely to have a secondary education, a trend observed throughout the city, thereby suggesting this is not just a story of the poorly educated being displaced from living in slums. Simultaneously, buildings in slums improved, particularly in terms of the wall and floor quality, investments which have substantial effects on living conditions. Access to electricity, and to a lesser extent sanitation, also rose. These improvements occurred faster than in rural Kenya. Third, the size of these changes vary by location of the slum and, in particular, central slums have different dynamics to more distant slums. In central slum areas, households have fewer assets, and rents are expected to be higher as the land provides greater access to services and jobs. Density is high and the number of rooms per person is low, and, in turn, improvements in building quality over time are limited. This gives an appearance of central slums being areas stuck in poverty. We cannot disentangle the reasons behind this slow progress, be it due to the lack of space or land tenure issues (though it is not correlated with a lower share of owner-occupiers) or political issues as discussed above. However, in central slums a higher share of residents are also in paid employment, indicating that there is a trade-off between better-quality living conditions and better access to urban opportunities. Fourth, despite these differences in dynamics, socio-economic indicators for children have improved throughout the city. Access to schooling in slum areas has caught up with the levels seen in other areas, and child mortality has fallen substantially. Together these facts show how slums can still be seen as places where people choose to live, despite appearances, at least compared to rural areas. In particular, living in slum areas provides better opportunities for children, and the potential for improving living conditions over time.

For policy-makers it is also important to underline what we cannot show, opening avenues into further research and allowing those invested in improving urban living conditions to be aware of the main caveats of this analysis. First, we cannot disprove from this the existence of any local poverty traps, but we can show that within slums there is the potential for development over time. For buildings and services, the reasons behind the slow improvements in central areas could be further investigated. Second, we also cannot rule out the idea that slums are inefficient use of land. They are providing accommodation for hundreds of thousands of urban residents in cheap housing; however, there are frictions in the allocation of this land and we do not know if this is preventing it from being put to a higher value use. Third, we show large variation across slums in the levels of investment in private buildings and public services. Some of this variation is explained by the location of the slum, but not all, and here the local context is of high importance. What is important to note is that for services that can be accessed through private investment, such as access to electricity for lighting, we have seen a large increase in provision. For services that require public investment or at least coordination between numerous households, such as sanitation, the improvements have been slower. Policy-makers should take note of the difficulties for local communities to set up such systems, and work to enable research and policy to support further improvements in this direction.

Finally, observing centrally located slums may give a sense of slums as areas of deprivation, with poor-quality buildings not changing over time while density rises. But the evidence suggests that these areas are more dynamic than first appearances. Internal investment in buildings is occurring, despite constraints on more substantial structural improvements; what can be done to support the latter as well as the former? Access to electricity is rising through likely private investments; can public services in slums be structured in a way that allows for similar improvements over the next decade? And child socio-economic outcomes are improving, with near universal school attendance following the introduction of free schooling; can similar strategies be used to bring child mortality down further?

4.A Figures

Figure 4.1: Location of major slums in Nairobi, relative to the central business district (CBD)

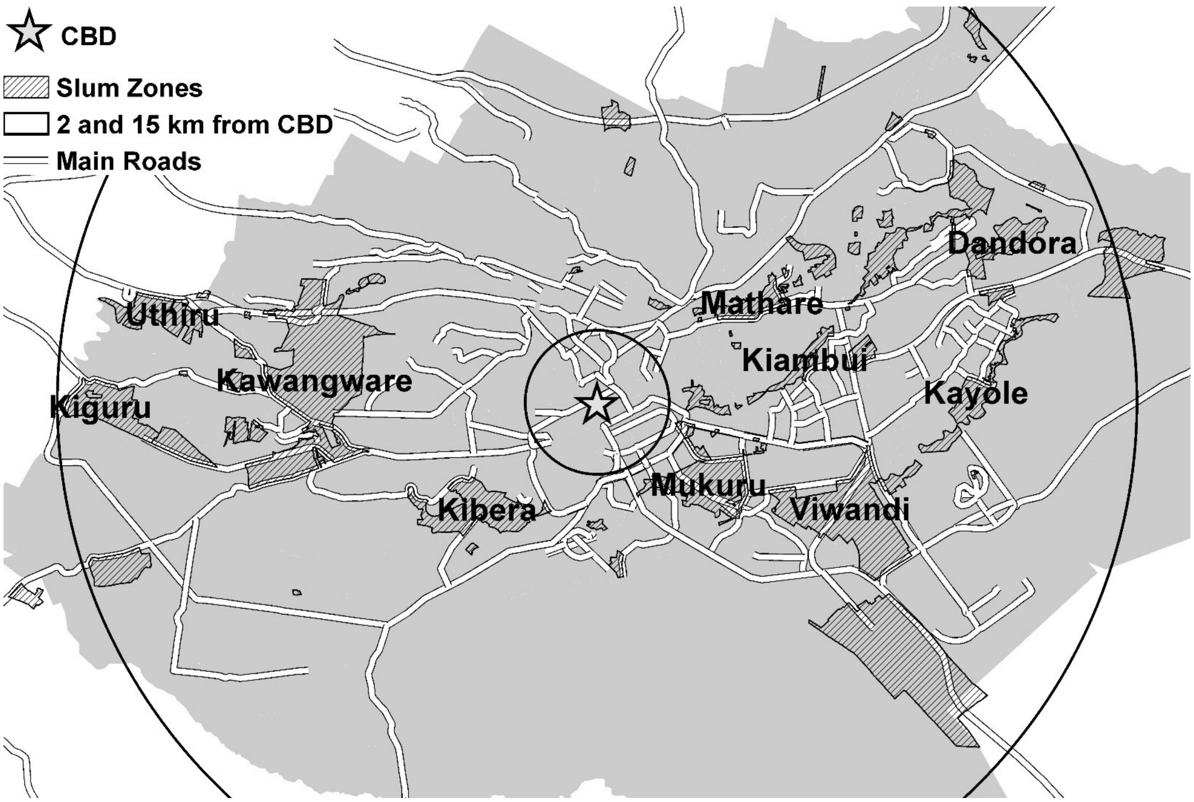

Figure 4.2: Slums and EAs in Nairobi

Figure 4.3: Location of residents relative to CBD

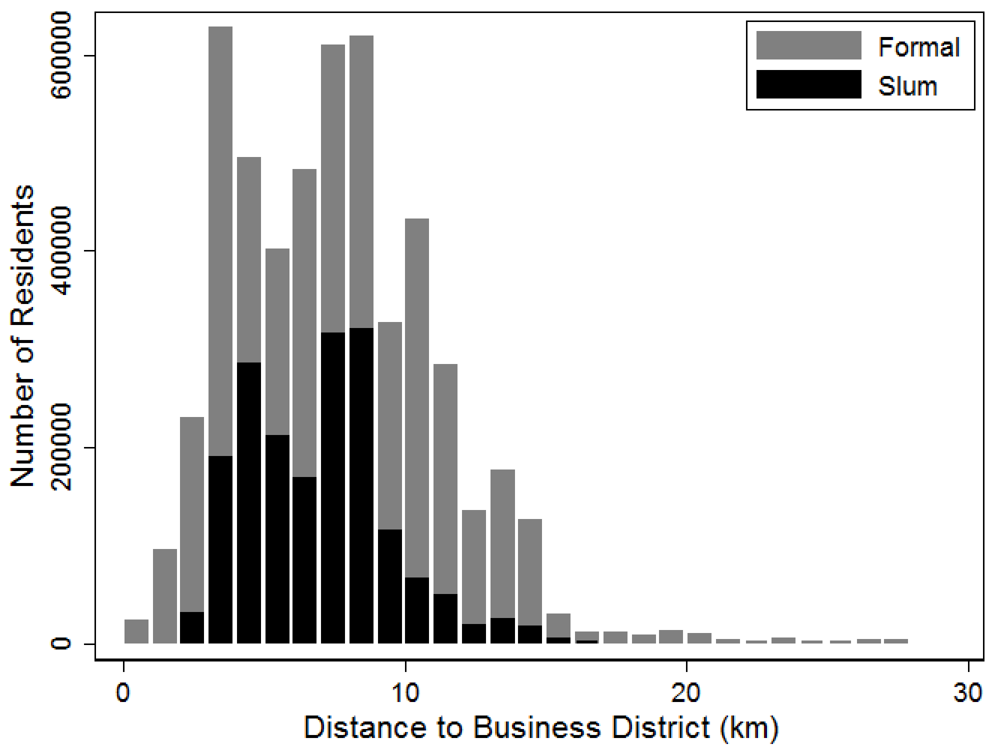

Figure 4.4: Population density across formal and informal areas, by distance to CBD

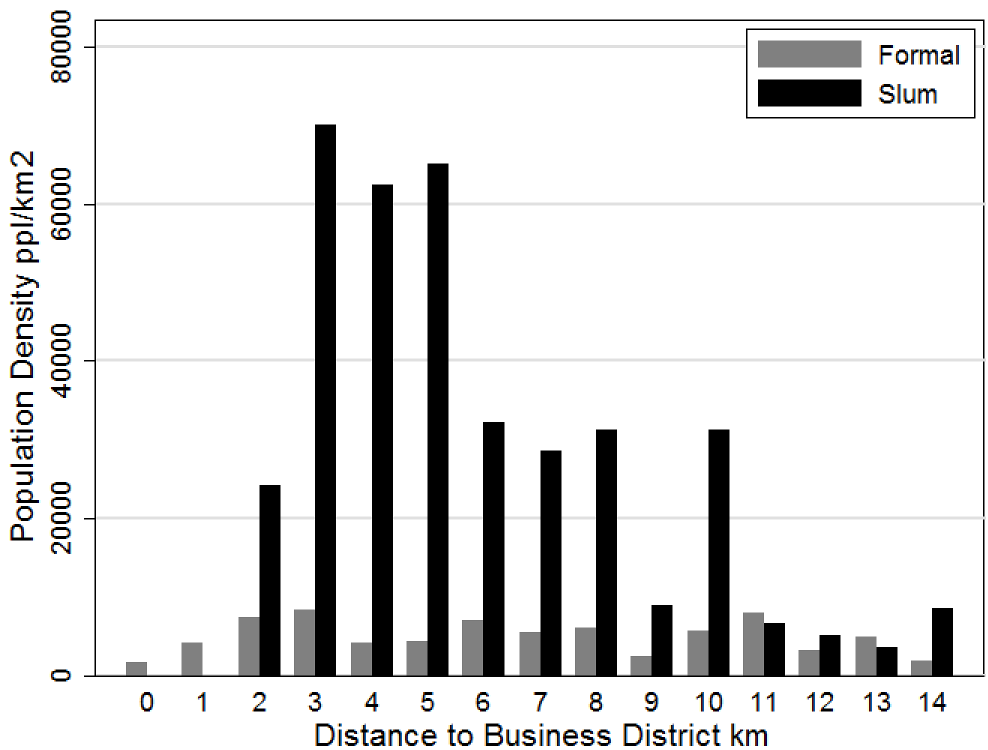

Figure 4.5: Adult population characteristics in slum and formal areas

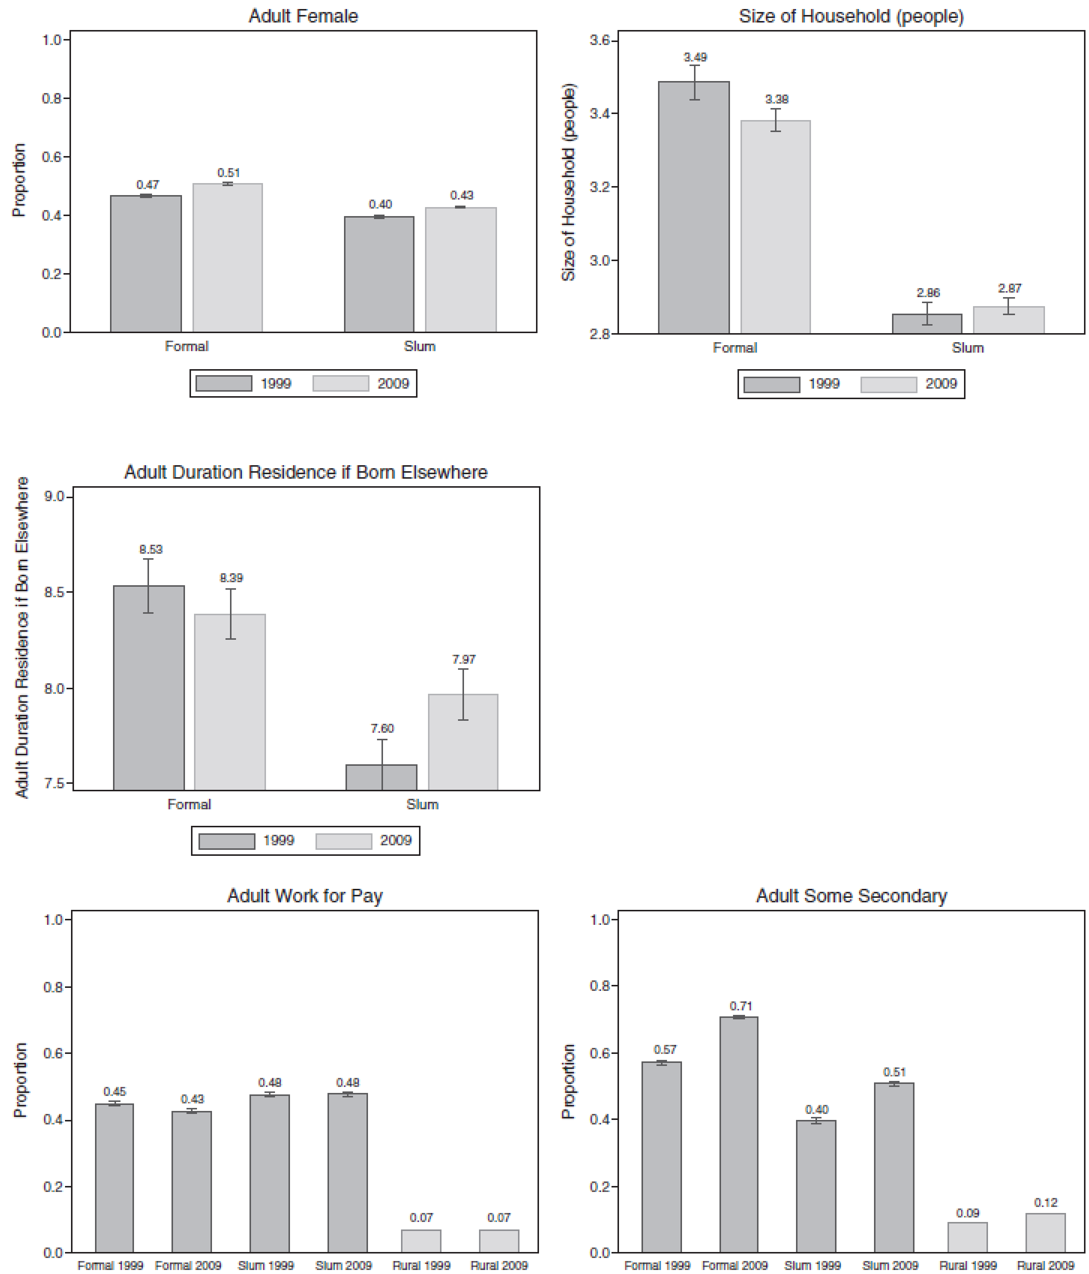

Figure 4.6: Population characteristics by distance to CBD, in slum and formal areas

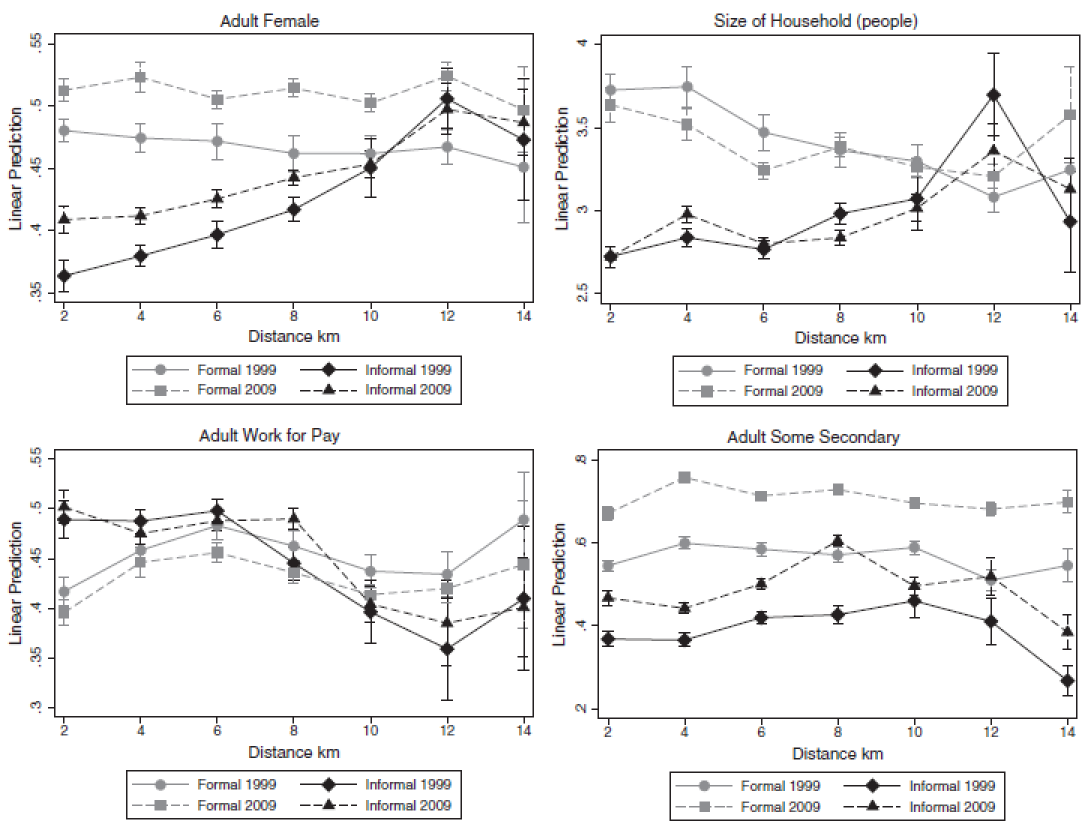

Notes: Marginal effects of being in a slum or formal area in 1999 or 2009, estimated using OLS over 2km bins, with 95 percent confidence intervals and errors clustered at the EA level.

Figure 4.7: Household asset ownership in slum and formal areas, 2009

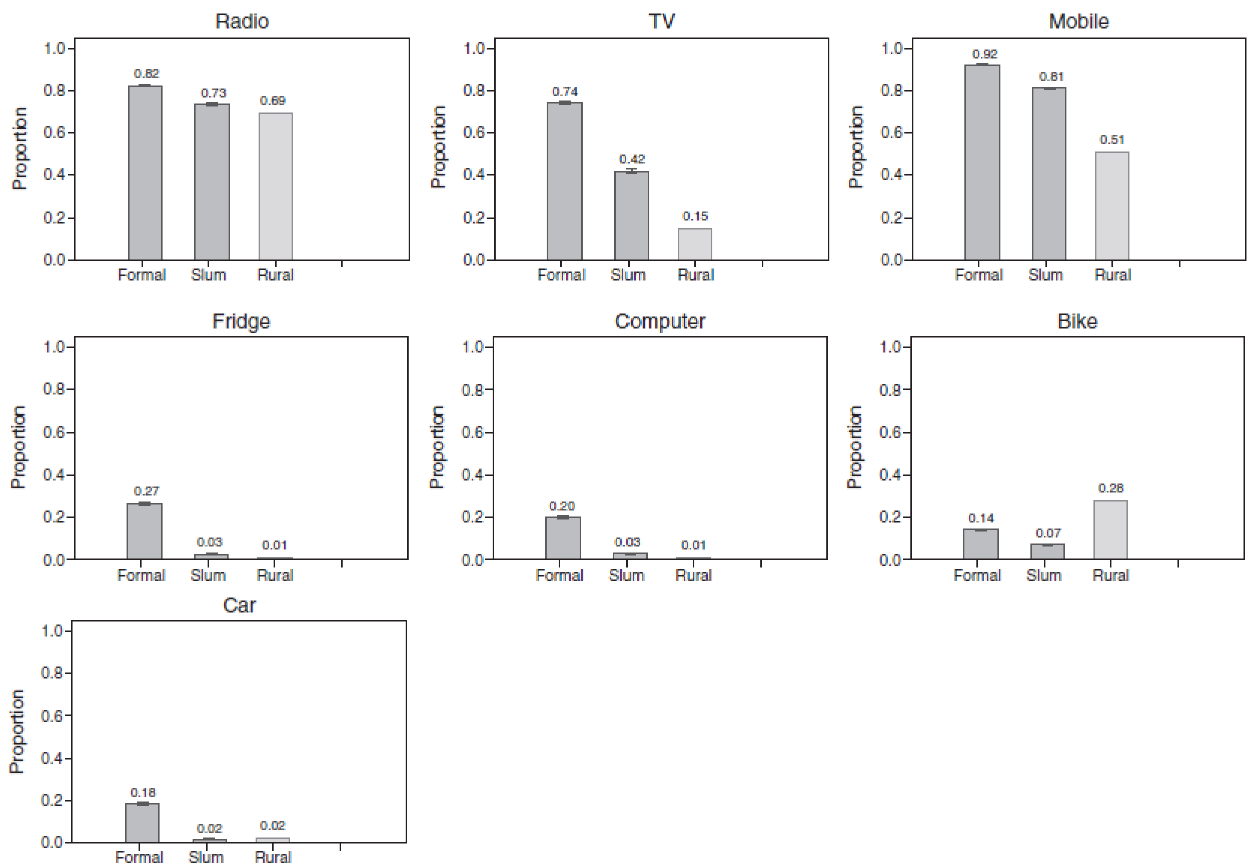

Figure 4.8: Variation in asset ownership according to distance from CBD

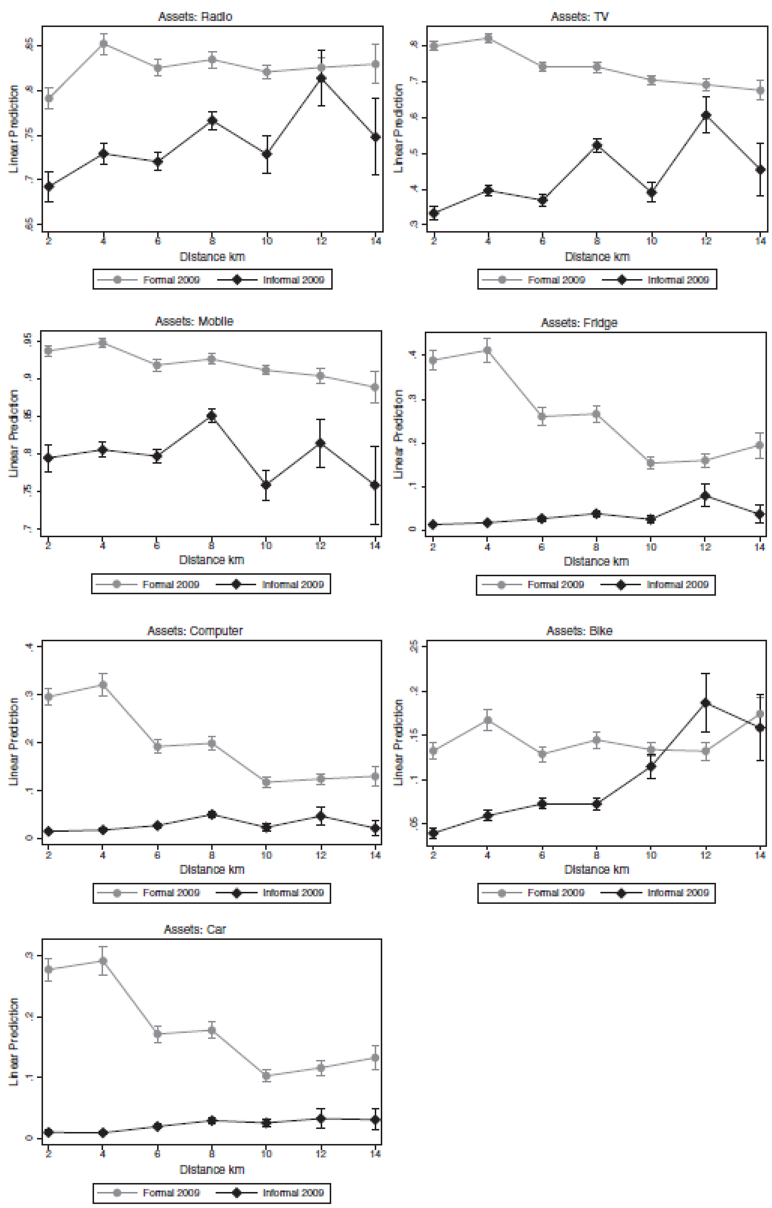

Notes: Marginal effects of being in a slum or formal area in 1999 or 2009, estimated using OLS with 95 percent confidence intervals and errors clustered at the EA level.

Figure 4.9: OLS sample means of building quality

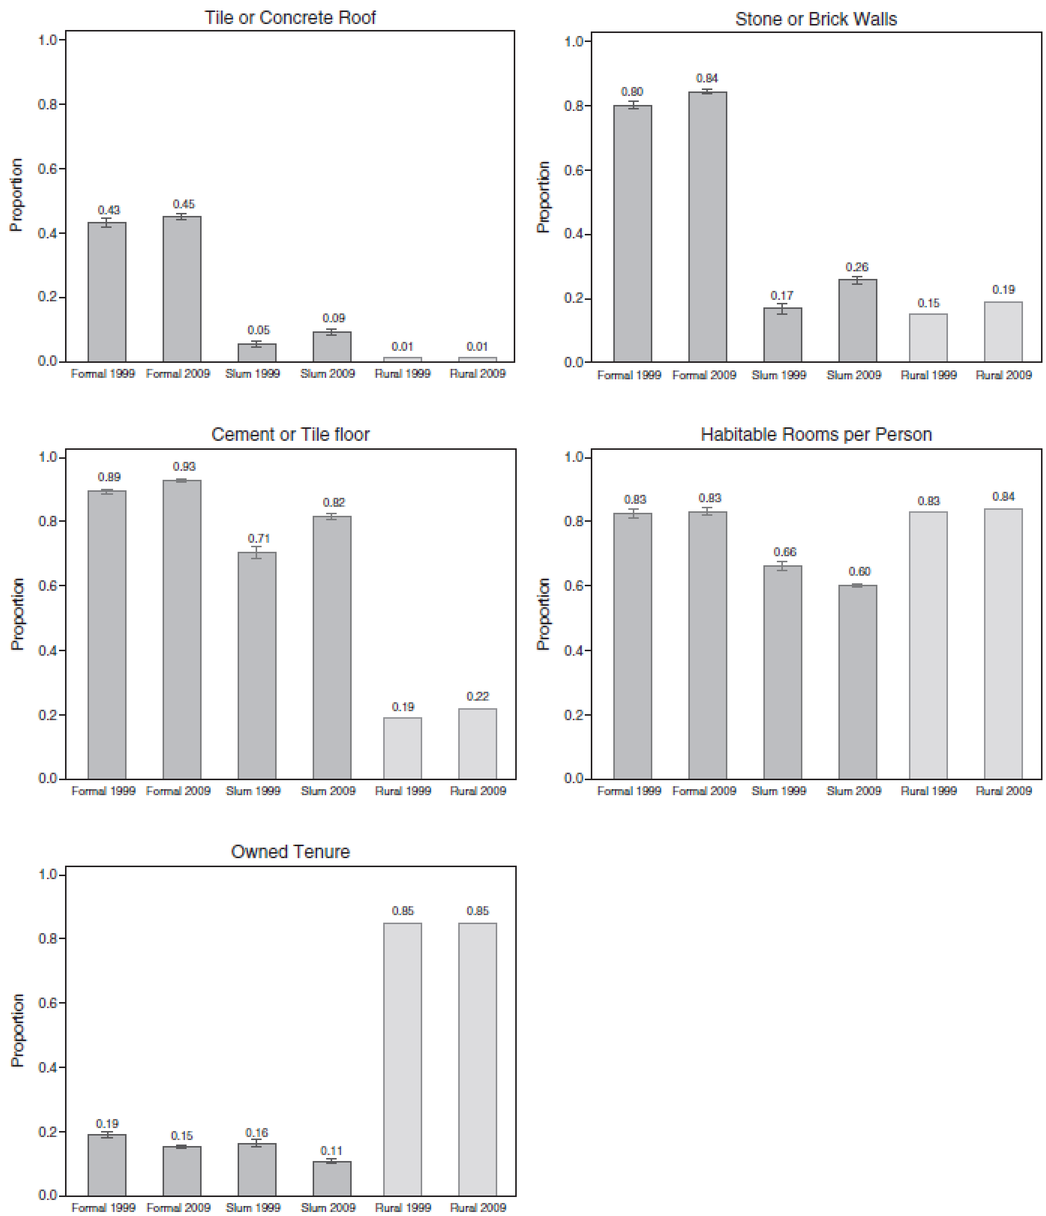

Figure 4.10: Variation in building quality by distance to CBD

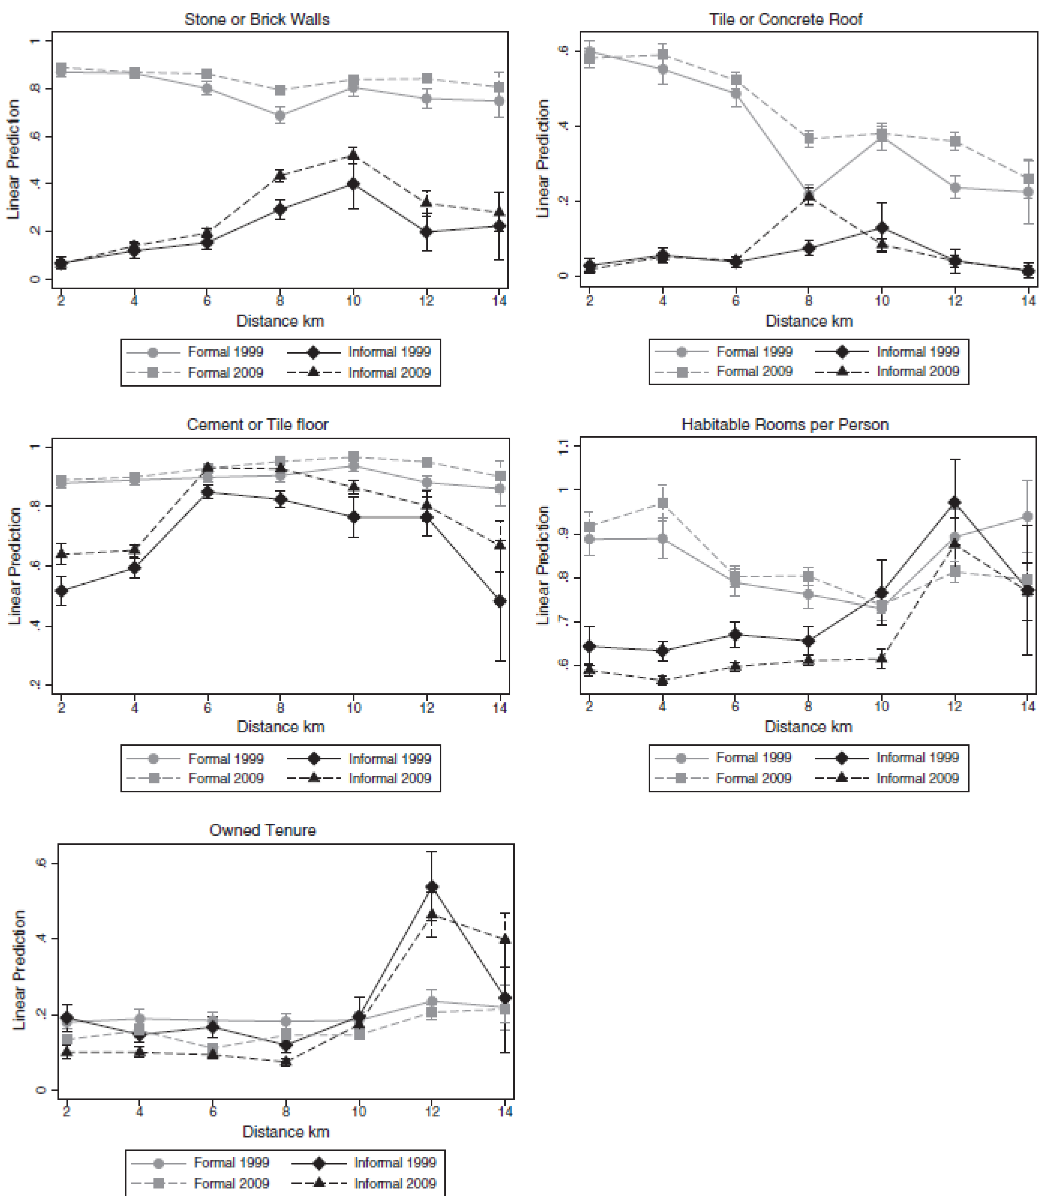

Notes: Marginal effects of being in a slum or formal area in 1999 or 2009, estimated using OLS with 95 percent confidence intervals and errors clustered at the EA level.

Figure 4.11: Variation in building quality by tenure status

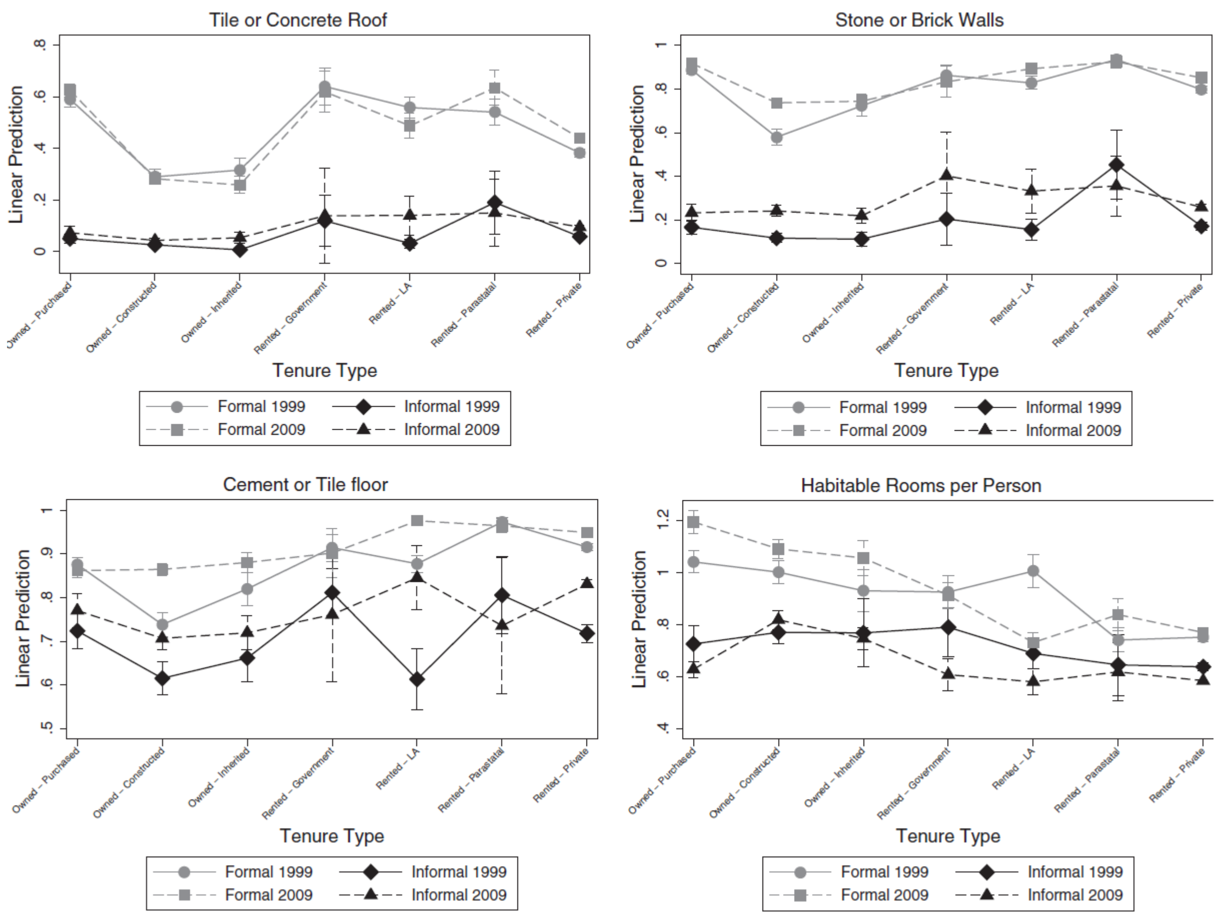

Notes: Marginal effects of being in a slum or formal area in 1999 or 2009, estimated using OLS with 95 percent confidence intervals and errors clustered at the EA level.

Figure 4.12: Mean access to services

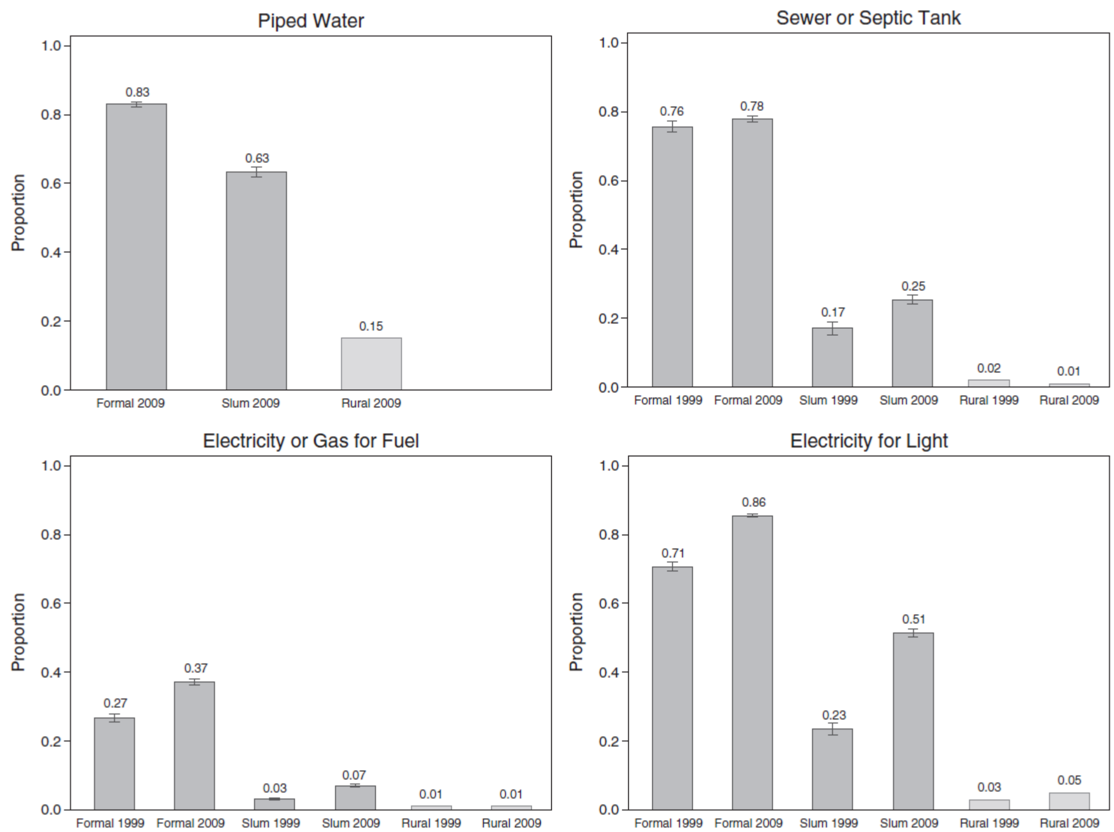

Figure 4.13: Variation in access to services by distance to centre

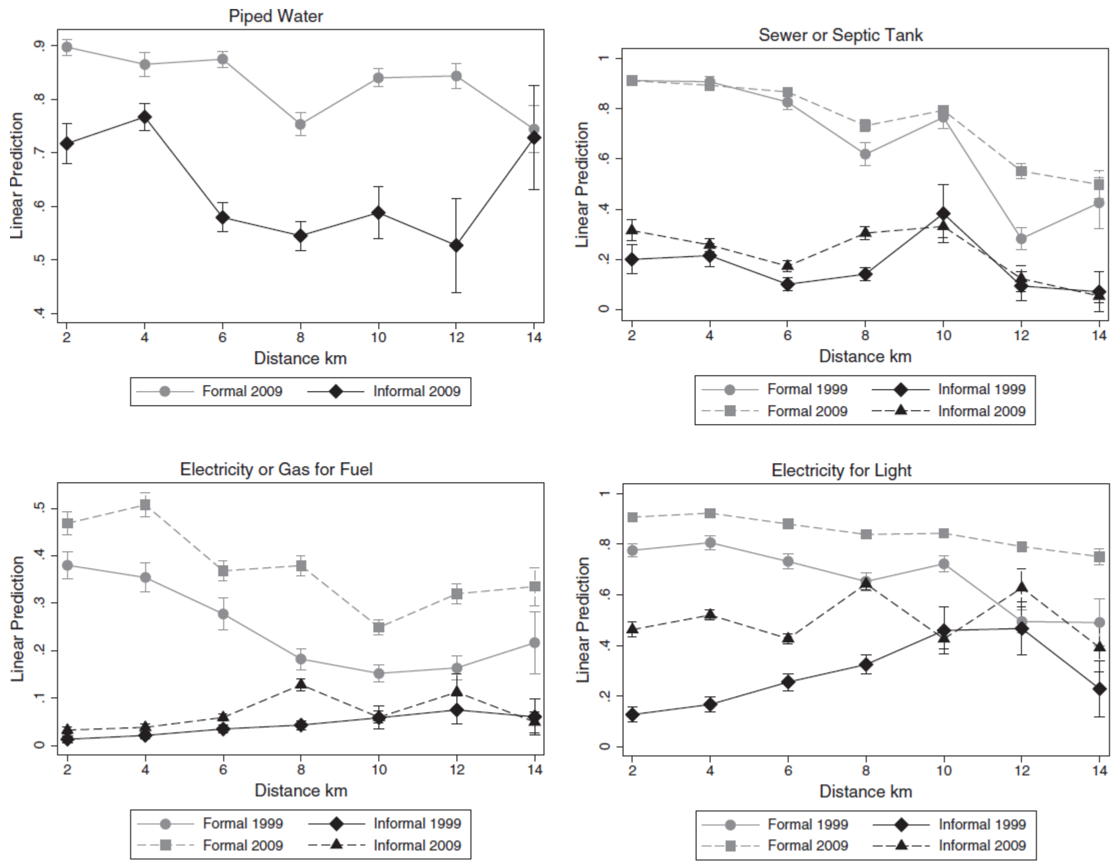

Notes: Marginal effects of being in a slum or formal area in 1999 or 2009, estimated using OLS with 95 percent confidence intervals and errors clustered at the EA level.

Figure 4.14: Mean child health and education outcomes

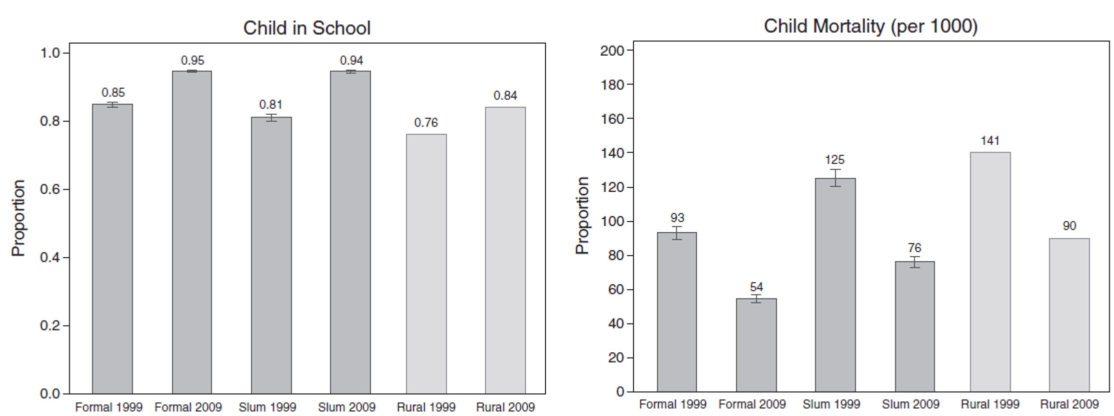

Notes: Note that child mortality is calculated as children lost from 1,000 live births, with no limit in age at which the child may have died, and is therefore not comparable to the standard definition with a 5-year threshold.

Figure 4.15: Variation in outcomes by distance to centre

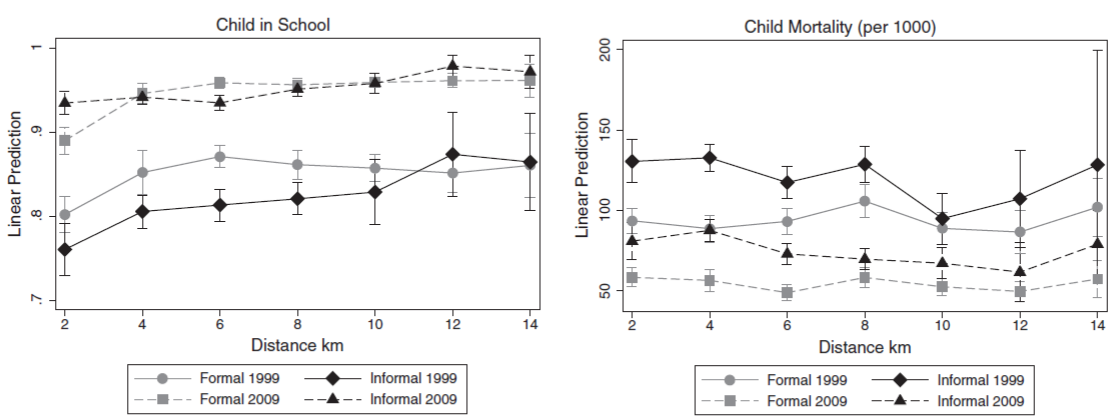

Notes: Marginal effects of being in a slum or formal area in 1999 or 2009, estimated using OLS with 95 percent confidence intervals and errors clustered at the EA level.

Chapter 5

Building the city: from slums to a modern metropolis

5.1 Introduction

The urban population of Sub-Saharan Africa is predicted to increase from 450 million to 1,250 million over the next 30 years (UN 2018). The populations of many large cities are increasing by about 50% every 10 years, implying a huge demand for increased building volumes. World Bank (2006) suggests that about two-thirds of a country's non-governmental capital stock are buildings, and urban construction and maintenance are a large and rising share of many countries' total investment. Yet, we know little about this investment process that reshapes cities and redeploys nations' capital stocks. African cities are subject to rapid redevelopment and expansion in the formal building sector; while at the same time, there is formalization of some slums, densification of others, and spread of new slums at city edges. In parts of cities, formal sector high-rise buildings on high price land are adjacent to low-level slums with limited returns to land, suggesting inefficient land-use and distorted investment decisions.

This paper models the process of a city's growth and applies the model to a unique data set we have developed on Nairobi, the capital of Kenya. The paper does four novel things. First, we develop a model of the built environment of a growing city. Given the difficulty of having both continuous time and space in the same model, especially with durable capital, virtually all models of cities are static. Fujita (1982) has an early attempt at a dynamic model and Braid (2001) put forth the first complete example of a dynamic monocentric city model with durable capital. To study the development of rapidly growing cities in middle- and low-income countries a full dynamic framework is essential. We extend and enrich the base model designed for developed countries by adding a slum sector and a role for geographical and institutional frictions, and we develop clear propositions about city growth.

Dynamics involve spread of a city's land area, intensified land use within the city, and demolition and reconstruction with increased density and building height. Our model captures these features, with a key distinction between formal and informal (slum) sectors. Formal buildings involve sunk capital costs, can be built tall, cannot be modified once constructed, and involve investment decisions based on expected future rents. As the city and housing prices grow formal sector buildings are periodically demolished and redeveloped to a greater height. In contrast, in the informal sector, buildings generally are much more malleable and built with non-weight-bearing materials. They are built low, with high land intensity. The building volume delivered by slums increases through time not by building taller, but by increases in crowding and already high cover-to-area ratios. New slums appear near the expanding city edge as the highest and best use of that relatively cheap edge city land. Within the city, the model predicts that as housing prices rise, there should be on-going conversion of older slums to formal sector use.

Second, we assemble a rich data set on the built fabric of Nairobi and present aspects of the 2015 cross-section and its evolution since 2003. While there is work on the USA using census data of building ages (Brueckner & Rosenthal 2009), we know of no work which utilizes citywide data on individual buildings, with demolition, redevelopment, and infill to detail the changes in the urban landscape. For the Nairobi data set, we have tracings of all buildings in the city from aerial photo images for 2003 and 2015, which give a precise delineation of the built environment. We develop an algorithm to overlay the two cross-sections of polygons and determine which building footprints are unchanged, which buildings were demolished and/or redeveloped and where and to what extent infill occurs. Building heights are derived from LiDAR data allowing us to measure building volumes. We also have formal sector vacant land prices for 2015, housing rent data for the formal and informal sectors in 2012, and data on land quality throughout the city, covering slope, ruggedness, water cover, and the like.

With these data, we construct a rich picture of the built environment of Nairobi, both in the spatial cross-section and its evolution through time. The (UN 2018) reports a population for the agglomeration of just over 4 million, a population that grew at 4.4% p.a. from the 1999 to 2009 census (Kenya 1999, 2009). We find that the built volume of the 2015 city area increased at 4.0% p.a., expanding by 60% between 2003 and 2015. Within just the 2003 spatial area of the city, overall volume growth was an astounding

47%, accounting for two-thirds of total volume growth within the 2015 city area. Much of this growth involves formal sector redevelopment. For example, at 3-4 km from the centre, redevelopment increased volume by 30%, achieved by tearing down one third of buildings and reconstructing more intensively. We find that newly redeveloped buildings are about triple the height of those torn down. Slower growing developed country cities generally have nowhere near this degree of change.[1]

Nairobi's slums deliver similar built volumes per unit land as the formal sector, and house just under 30% of the population in both 1999 and 2009, similar to UN Habitat (2016*a*) estimates for many developing country cities. Slum volume growth of 50% between 2003 and 2015 was slower than formal sector volume growth of 61%. As noted above, new slums appear on the city edge, often on land with private property rights, highlighting the ongoing role for slums to provide low cost housing in developing countries. However, apart from growth at the edge, a key feature is the persistence of several slum areas on high value land near the city centre, most notably Kibera, often referred to as Africa's largest slum.

Third, we estimate all the parameters of the dynamic model based on the characteristics of the 2015 cross-section of the built environment and on land sale and house-rent data, together with the discount rate that we take from World Bank data on Kenya. These parameters describe the intensity of land-use in the informal and formal sectors and gradients telling us how prices, building volumes and heights change with distance from the city centre in both sectors. We then use these parameters and the model to make predictions about the evolution of the built environment and population from 2003 to 2015 and compare these with the evolution we see in the data.

Finally, we calculate the welfare cost of inefficient land-use from delayed formalization of slums near the city centre. This highlights the role of policy for fast growing cities with major land market failures that deter investment. We argue that institutional frictions in Nairobi delay timely redevelopment of inner-city slums and quantify the losses from such delays. For older slums at 3-5 km from the centre, we estimate that the city could pay off illegal slumlords operating in these slums for the value of their land (reflecting their profits) in perpetual slum use. Even after that, formalization would still bring a gain of about \$18,000 per slum household, in a context where slum households spend on average about \$500-700 p.a. on housing. Our calculations do not capture the wider social costs

[1] To give perspective on the demolition rate of one-third, in the USA, the American Housing Survey data give demolition and removal by disaster (fire, hurricane, and the like). Annual rates of demolition and removal by disaster range from 0.5% to about 1.2%. For 12 years, this would involve 6–15% of building removal. Nairobi is thus 2–5 times that.

and benefits of formalization which we discuss, but they do quantify the direct land-use efficiency loss of delayed formalization.

5.1.1 Context and literature

Our modelling judgements are based on several key features of Nairobi, features that we think are applicable to many other developing cities. First, we treat Nairobi as a monocentric city. Monocentricity is the workhorse assumption of urban economics, including contemporary papers such as Combes et al. (2019) on modern French cities. As we will see, Nairobi appears to be primarily monocentric with building height, volume, and land prices all declining sharply with distance from the centre. This monocentricity appears to be a feature of many major African cities. Lall et al. (2017) show that African cities like Nairobi, Addis Ababa, Dar es Salaam, and Kigali are much more monocentric than Paris, London, Barcelona, or Atlanta. For our own sample of world cities similar to Nairobi in size, in Appendix Section 5.2.4, we show that developing country cities and Nairobi in particular display greater monocentricity than most developed country cities, with density dropping off three times more rapidly from the centre.

The second feature is that there are relatively well-functioning markets for the allocation of housing within both the formal and informal sectors, but a problem in the conversion of informal to formal sector land-use. While the two sectors offer very different qualities of housing, both are overwhelmingly rental markets; 87% of Nairobi residents overall and 89% of slum dwellers rent, proportions similar to those in many African cities.[2] The rental markets allocate land between occupants, but the individuals who control slum land and to whom rents are paid often have insecure land rights. This does not impede the functioning of rental markets, but insecure rights do deter those who control the land from sinking the capital required to construct formal sector buildings. Without security of ownership such investments are too risky to be undertaken. This is a problem particularly on government owned—but not managed—lands nearer the city centre in Nairobi. Politically well-connected slumlords 'illegally' operate vibrant rental markets, but they cannot switch these lands to formal sector use because that would open their control of the land to challenge. Their ability to operate the slum rental market involves corruption, and a complex political history, as we describe in Section 5.6. Marx et al. (2013) highlight the import-

[2]From UN Habitat (2011) on Ghana about 75% of residents in Accra and Kumasi rent. From authors' own calculations from IPUMS (2019) about 40% in Dar es Salaam rent. At a country level, among urban dwellers as defined in IPUMS for Ethiopia (2007), Nigeria (2010), and Uganda (2002), respectively, 51%, 60%, and 58% rent, where we expect higher amounts in primate cites. So while the urban Kenya rate in IPUMS is 56%, in Nairobi as noted in the text it is 87%. Some other counties including Mozambique and South Sudan have higher ownership rates.

ance of "investment inertia", where small groups of powerful individuals control large slum housing markets. Formalization and similar initiatives would challenge entrenched land control, so the status quo of informal status persists. These situations seem to be common in older slum lands near city centres in many developing country cities in Asia and Africa.[3] These cities exhibit the oft-photographed hotchpotch of land uses, with tall formal sector buildings bordering pockets of single storey corrugated iron sheet slum housing in the urban core as institutional frictions prevent conversion of these slums to formal sector development. It is this that motivates our modelling of distinct formal and slum sectors in this context, following the literature that generally makes the same sharp distinction, even if the line between informal and formal sector housing may be fuzzy at times.

Similar institutional barriers to formalization arise in slum areas around the edge of Nairobi. Here a higher proportion of land is privately owned. However, some owners may have only possessory rights which need to be converted to formal modern title prior to sinking capital in formal sector development. This again involves formalization costs, in this case in the form of fees, legal costs, delay, and possibly bribes. Frictions arise for natural as well as institutional reasons, including poor geographic conditions such as steep slope, ruggedness, marshes, and exposure to flooding. We examine these empirically to establish when and where each set of frictions matters. Our results suggest that institutional rather than natural frictions play the dominant role for older slums near the centre.

On other literature, we note that recent papers have varying perspectives in modelling the informal-formal sector distinction. For the Brazilian context, in Cavalcanti et al. (2019) slum owners make a choice between formalizing, which involves paying for secure rights and incurring taxes and regulation, versus staying in the unregulated informal sector with insecure land rights. We do not focus on taxes and regulation, because in Nairobi tax systems generally apply to both sectors and there is very low compliance in either sector. Similarly, land use regulations appear not to bind.[4] Two recent papers investigate whether

[3] There is no central source of information on slum locations and sizes. We explored sources including the UN Habitat (2016*a*), Davis (2006), and the repository of CIESEN at Columbia University. We complied a long list of slums near city centres. Well-known examples include Dhavari in Mumbai with population estimates of 7–800,000; Manshiyat Nassor (or "Garbage City") in Cairo with its edge within 3 km of the centre and population up to 800,000; and Ga Mashie and Old Fadama in Accra housing about 400,000. City centre slums in Jakarta and Dar es Salaam are analysed in Harari & Wong (2018) and Michaels et al. (2020).

[4] In Nairobi, land-use regulations seem to be largely ignored (Mwaura 2006) and property taxation of formal sector lands has very limited implementation and compliance (Kelly 2003). Narubi (2015) shows for property taxes compliance in Nairobi in 2014/15 was under 7%, and the collection ratio was about 16%, suggesting only some bigger commercial properties pay these taxes. Furthermore, from the 2010 cadastre for Nairobi county, most slum lands are assessed these same taxes. On regulations, the minimum lot size is generally 500m^2, and never less. In Figure 5.10, from a full cadastre of the formal sector, we show that modal lot size is at 180m^2, 54% of all buildings are in violation of the regulation, and that

slum upgrading programmes inhibit timely formalization. For Jakarta, Harari & Wong (2018) argue that slum improvement programmes serve to delay redevelopment of slums. Similarly, Michaels et al. (2020) argue that, compared to untreated neighbouring slums, improved slums either look no different than untreated ones 2–3 decades later, or have poorer characteristics, in Dar es Salaam. Finally, Gechter & Tsivanidis (2017) argue that in Mumbai very large formal sector developments can spur redevelopment of nearby slum areas. Our work complements these papers looking at the dynamic aspects involved in investing in land and using the dual structure of urban land markets that we think is relevant to many developing country situations.[5]

The basic model and core theoretical results are set out in Section 5.2 of the paper. Section 5.3 presents data and estimates key empirical relationships. Section 5.4 backs out the parameters of the model, based on empirical evidence, and validates key assumptions of the model. Section 5.5 analyses the evolution of the built environment and compares model predictions with the data. Section 5.6 develops the welfare analysis to derive costs of misallocation of land to older slums, and Section 5.7 concludes.

5.2 Theory

In this section, we develop the model of a growing city, focusing on investment decisions and consequent patterns of land-use and built structures. Sections 5.2.1 and 5.2.2 describe housing demand and supply, introducing preferences and technologies for informal (slum) and formal sector buildings. We derive the investment decisions and consequent land-rent that would be earned if land at a particular place and date were in informal or formal sector use (or left for agricultural use outside the city). In Section 5.2.3, we turn to landholders' dynamic optimization problems of choice of land-use, i.e. whether, at each place and date, land should be occupied by slums or formal sector buildings, and the timing of formal sector redevelopment. We look first at a particular point in the city and examine its evolution through time as it transitions from agricultural use to informal development, then formalizes and goes through successive waves of formal sector demolition and reconstruction. Section 5.2.4 draws out the way in which his path varies across points in the city, giving a description of both the cross-section of the city and its evolution through time. Section 5.2.5 adds frictions to the transition process, in particular spatial

there is a long right tail for commercial buildings. Most particularly, the curve for that distribution has no mass, step, blip, or kink around 500m^2.

[5] Additionally, Selod & Tobin (2018) model owners as choosing between different gradations of property rights at different costs, an issue relevant for some contexts. Ferreira et al. (2016) on Brazil focus on modelling how slums may be stepping-stones in skill formation and growth of cities.

heterogeneity in formalization and construction costs. Together, these analytical sections provide the relationships upon which empirics, estimation of key parameters and welfare analysis are based.

5.2.1 Bid-rents and housing demand

The value of a unit of housing varies across the city and through time, and we start by deriving the bid-rent for housing at places x and date t. A representative household has preferences between consumption of housing and other goods (the numeraire) described by $U(s(x,t)a, w(x,t)q(x,t)s(x,t))$. In this expression, $s(x,t)$ is the quantity of housing space consumed by an individual living at place x at date t, $q(x,t)$ is its observed rental price, and a is an amenity parameter which will be discussed further below. Household income available for housing and purchase of other goods is $w(x,t)$. We take the function $w(x,t)$ to be exogenous and assume that it is increasing in t and weakly decreasing in x. The former is based on the city using labour to produce output with constant returns to scale and experiencing productivity growth relative to the rest of the economy.[6] The latter can be a simple labelling of places. However, we generally interpret x as distance from the city centre, in which case $w(x,t)$ is interpreted as wages at the centre minus commuting costs to distance x.[7] We impose a particular functional form on $w(x,t)$ from Section 5.2.4 onwards.

The city is"open", meaning there is free migration between the city and outside, where utility takes exogenous value u_0. Equilibrium rent leaves all households indifferent between living at each occupied place x in the city or receiving outside utility u_0. Housing rent $q(x,t)$ therefore satisfies $U^*(\frac{q(x,t)}{a}, w(x,t)) = u_0$, where $U^*(.)$ is the indirect utility function (each household's location and $s(x,t)$ optimally chosen). Since the amenity factor enters the utility function multiplicatively with $s(x,t)$, it enters indirect utility in the form $\frac{q(x,t)}{a}$. Henceforth, we work with rent per unit amenity, defined as $p(x,t) = \frac{q(x,t)}{a}$, and referred to as space-rent. The net wage $w(x,t)$ therefore determines space-rent through the relationship

$$U^*(p(x,t), w(x,t)) = u_0, \text{ for example, } p(x,t) = \left(\frac{w(t,x)}{u_0}\right)^{\frac{1}{\beta}} \tag{5.1}$$

[6] Agglomeration economies are not included in this model. The price of output is exogenous to the city, and any variations in the output price have the same effect as changes in physical productivity.

[7] This is as in the standard urban model, where all households commute to the central business district (CBD). Newer approaches offer a richer model of commuting and the dependence of net wages on location (e.g. (Ahlfeldt et al. 2015)). We assume also that households have no other source of income, implying that land-rents are received by non-resident landowners or slumlords.

where the example has Cobb–Douglas preferences with housing share β and hence chosen quantity $s(x,t) = \frac{\beta w(x,t)}{ap(x,t)}$ (see Appendix Section 5.C.1). Outside utility is exogenous and will henceforth be supressed in notation.[8] We assume that housing-space is a normal good, implying that space-rent $p(x,t)$ is increasing in $w(t,x)$, and hence increasing with t and weakly decreasing with x. The dynamics of the model are therefore driven by urban productivity growth that increases $w(t,x)$, i.e. shifts this horizontal labour demand curve up through time. This increases space-rent, and the consequent growth of city population is determined by the supply response of urban housing.

We note that, while we have different types of building technology, we do not model different types of households. Were we to do so, land-use would be determined by the upper envelope of the land-rent functions generated by different households.[9] In our empirical work, land-rents in some parts of the city are systematically lower for land in informal use than for equivalent land in formal use. This is indicative of inefficient land-use, regardless of whether the land-rent function is that of a single household type or the upper envelope of many types. We will return to discussing the relevance of having different types of households in Section 5.6 when we conduct our welfare analysis.

5.2.2 Building technology and housing supply

There are two distinct building technologies, formal and informal, which supply building volume per unit land in different ways. The formal sector (F) can build tall, and the informal sector (I) can "crowd", increasing cover, i.e. the proportion of land covered by building footprints. The volume of building that a technology can supply on a unit of land at a particular place, x, at time t, is the product of height and cover, $v_i(x,t) = h_i(x,t)c_i(x,t),\ i = I, F$. Formal sector buildings are durable structures, which incur sunk costs of construction and increasing marginal cost of building tall. The informal sector is unable to build tall because the materials used are not sufficiently loadbearing, this also giving a "malleability", enabling structures to be rapidly modified and implying that construction costs are not sunk. Informal floor-space can be added at constant marginal cost, but crowding reduces amenity and hence housing rent.

This distinction captures the difference in sunk costs associated with each housing type, and its technological basis accords well with the difference between slums and formal sector areas in Nairobi. The 2009 Nairobi Census indicates that, in slums, the majority (about 52%) of households live in buildings where walls are corrugated iron sheets which can be

[8] u_0 is potentially time varying, but in the later analysis it is only the exogenous ratio $\frac{w(x,t)}{u_0}$ that matters.
[9] See for example (Duranton & Puga 2015)

easily reconfigured like Meccano parts; much other slum housing involves mud construction (about 19%) and other material with short duration (Kenya 2009). These materials are not sufficiently load-bearing to allow much height, and our data show that 85% of slum buildings are under 5m tall. In contrast, 84% of formal sector households live in a building that is made of stone or some type of brick or block, many with substantial height. The main missing category in both sectors is wood construction, where wood buildings in slums are mostly very different quality structures than those in the formal sector. While, in reality, the distinction between slums and formal sector developments can be ambiguous, the binary distinction is analytically tractable. We also think the sharp distinction captures the key features of Nairobi and is supported by the literature, including (Taubenbock et al. 2018) who analyse the morphology of slums for a large selection of cities worldwide.

Informal Sector

Informal sector construction materials are malleable and construction costs are a flow, occurring continuously through the life of the structure. This can be thought of as either the rental on 'Meccano parts' used in construction or as the cost of material whose life is one instant. This sector is unable to build tall, so height is fixed at $h_I = 1$. It can, however, change at any instant the proportion of each unit of land that is covered with buildings, $v_I(x,t) = c_I(x,t)$. Construction costs per unit volume in this sector are constant κ_I, so construction costs per unit land are $\kappa_I v_I(x,t)$. However, crowding more building onto land in slums reduces housing amenities, which we capture by assuming that the amenity is a diminishing and convex function of cover per unit land, so $a(v_I(x,t)), a' < 0, a'' > 0$. If land at x, t is in informal use, then land-rent is space-rent times amenity minus construction cost, all times volume per unit land,

$$r_I(x,t) = [p(x,t)a(v_I(x,t)) - \kappa_I]\, v_I(x,t) \tag{5.2}$$

The volume of housing supplied at any instant is chosen to maximize land-rent, taking space-rent $p(x,t)$ as exogenous and internalizing the effect of crowding on amenity.[10] The first-order condition equates marginal revenue to marginal cost,

$$\frac{\partial r_I(x,t)}{\partial v_I(x,t)} = p(x,t)a(v_I(x,t))\left[1 + v_I(x,t)\frac{a'(v_I(x,t))}{a(v_I(x,t))}\right] - \kappa_I = 0 \tag{5.3}$$

[10] A simplification is that we do not explicitly model this as an externality in which each developer's choice of crowding affects neighbours.

Generally, we assume that informal amenity is iso-elastic in cover, $a(v_I(x,t)) = a_I v_I(x,t)^{\frac{(1\alpha)}{\alpha}}$, with $\alpha > 1$. Then the chosen volume of housing and maximized value of land-rent are

$$v_I(x,t) = \left[\frac{a_I p(x,t)}{\kappa_I \alpha}\right]^{\frac{\alpha}{\alpha-1}}, \tag{5.4}$$

$$\begin{aligned} r_I(x,t) &= (\alpha-1)\kappa_I v_I(x,t) \\ &= \kappa_I(\alpha-1)\left[\frac{a_I p(x,t)}{\kappa_I \alpha}\right]^{\frac{\alpha}{\alpha-1}} \end{aligned} \tag{5.5}$$

The loss of amenity due to crowding, α, is a key parameter. It will be recovered in the empirics based on how informal sector volume in equation (5.4) varies with distance from the city centre. The parameter also determines the division of revenue between land-rent and construction costs. Equation (5.5) indicates that land-rent is multiple $(\alpha - 1)$ of construction costs, so total space-rent earned per unit land is divided with proportion $(1-\frac{1}{\alpha})$ going as land rent, and $\frac{1}{\alpha}$ to construction. Iso-elasticity implies further that informal sector space-rent adjusted for amenity is constant throughout the city at value

$$p(x,t)a_I v_I(x,t)^{\frac{1-\alpha}{\alpha}} = \kappa_I \alpha \tag{5.6}$$

Essentially, and as will be confirmed later in the space-rent data, increased crowding near the city centre offsets the advantage of improved access to the centre.

Formal sector

The formal sector differs in that buildings are durable, construction is a sunk cost, and volume is achieved by varying height, not cover. Height is chosen at date of construction, denoted τ_i, so $h_F(x,\tau_i)$ is fixed for the life of the structure, i.e. until demolition at date $\tau_i + 1$, where subscript $i = 1, 2, ...$ denotes successive redevelopments (or 'generations') of formal structures. For simplicity, but also based on the cross-section data, we assume that formal sector land cover is uniform at $c_F = 1$, so the volume of building at x developed (or redeveloped) at τ_I is constant at $v_F(x,t) = v_F(x,\tau_i) = h_F(x,\tau_i)$ for $t \in [\tau_i, \tau_{i+1}]$. Construction costs per unit land are one-off and sunk, and are an increasing and convex function of height (= building volume) on that land, $k(v_F(x,\tau_i))$, $k' > 0$, $k'' > 0$. Demolition incurs no costs.

We assume there is no amenity loss or gain from building tall, so we set $a_F = 1$. The space-rent of a unit of building volume is $p(x,t)$, exogenous to the developer. The present value of land-rent that accrues over the life of a structure, $t \in [\tau_i, \tau_{i+1}]$, discounted to construction date τ_i at interest rate ρ is denoted $R_F(x,\tau_i)$. With sunk costs $k(v_F(x,\tau_i))$ and volume fixed at the date of construction this is

$$R_F(x,\tau_i) = v_F(x,\tau_i)\int_{\tau_i}^{\tau_{i+1}} p(x,t)e^{-\rho(t-\tau_i)}dt - k(v_F(x,\tau_i)). \tag{5.7}$$

We define the ratio of the present value of space-rent per unit volume over its life relative to space-rent at date of construction as

$$\Phi(x,i) = \int_{\tau_i}^{\tau_{i+1}} \left[\frac{p(x,t)}{p(x,\tau_i)}\right] e^{-\rho(t-\tau_i)}dt, \tag{5.8}$$

so $R_F(x,\tau_i) = \rho(x,\tau_i)\Phi(x,i)v_F(x,\tau_i) - k(v_F(x,\tau_i))$. The integral $\Phi(x,i)$ is akin to the 'value-to-rent ratio' on a newly constructed property in the terminology of the real-estate literature (noting the time horizon in (5.8) is cut at τ_{i+1}).

If land at place x undergoes formal sector development (or re-development) at date τ_i, the volume chosen is intuitive and given by

$$\frac{\partial R_F(x,\tau_i)}{\partial v_F(x,\tau_i)} = p(x,\tau_i)\Phi(x,i) - k'(v_F(x,\tau_i)) = 0. \tag{5.9}$$

If the cost function is iso-elastic, $k(v_F) = \kappa_F v_F^\gamma$, with $\gamma > 1$, then chosen volume and the maximized present value of land-rent are[11]

$$v_F(x,\tau_i) = \left[\frac{p(x,\tau_i)\Phi(x,i)}{\kappa_F\gamma}\right]^{\frac{1}{\gamma-1}}, \tag{5.10}$$

$$R_F(x,\tau_i) = \kappa_F(\gamma-1)\left[\frac{p(x,\tau_i\Phi(x,i)}{\kappa_F\gamma}\right]^{\frac{\gamma}{\gamma-1}}. \tag{5.11}$$

The diseconomy in building taller, γ is another key parameter and will be estimated based on how the height of newly constructed buildings and hence volume in equation (5.10) varies with location as space-rent varies.

To define land-rent share, it is useful to have a continuous flow measure of land-rent, given by amortizing the one-off construction cost continuously over the life of the structure. If amortization is a constant proportion μ of revenue over the life of the building, i.e. $\mu p(x,t)v_F(x,\tau_i)$, then construction costs are covered by setting μ to satisfy

[11] The iso-elastic form implies an elasticity of substitution between land and capital (i.e. construction cost) of unity. This is at the centre of the range suggested in by Ahlfeldt & McMillen (2014)

$\mu p(x,\tau_i)\Phi(x,i)v_F(x,\tau_i) = k(v_F(x,\tau_i))$. With iso-elastic cost function the amortization rate is then $\mu = \frac{1}{\gamma}$.[12] Flow land-rent net of amortization is therefore fraction $(1-\frac{1}{\gamma})$ of gross revenue earned by land and structure together. This land-rent, net of amortization, is place, time, and date of development specific, and we denote it $r_F(x,t,\tau_i)$; it is given by

$$\begin{aligned} r_F(x,t,\tau_i) &= (1-\frac{1}{\gamma})p(x,t)v_F(x,\tau_i) \\ &= (1-\frac{1}{\gamma})p(x,t)\left[\frac{p(x,\tau_i\Phi(x,i)}{\kappa_F\gamma}\right]^{\frac{1}{\gamma-1}}. \end{aligned} \tag{5.12}$$

Land-rents are therefore fraction $(1-\frac{1}{\gamma})$ of space-rent in the formal sector, and fraction $(1-\frac{1}{\alpha})$ in the informal sector (equation 5.5), relationships that we will use later in the paper. Notice also that, comparing (5.11) and (5.12) at date $t=\tau_i$,

$$r_F(x,\tau_i,\tau_i) = \frac{R_F(x,\tau_i)}{\Phi(x,i)}. \tag{5.13}$$

5.2.3 Land development and construction phases

Equations (5.5) and (5.5) give the maximized returns to respectively informal or formal sector use of land at place x and date t. We now look at the choice of where and when land is in one or other of these uses. We pose this as the decisions of a landholder of when to develop informal structures and when to develop or redevelop formal structures on land at place x. At some date (say time 0) the present value of land-rent at location x that has not yet been developed is

$$PV(x) = \int_0^{\tau_0} r_0 e^{-\rho t}dt + \int_{\tau_0}^{\tau_1} r_I(x,t)e^{-\rho t}dt + [R_F(x,\tau_i) - D(x)]e^{-\rho\tau_1} + \sum_{i=2}^{\infty} R_F(x,\tau_i)e^{-\rho\tau_i}. \tag{5.14}$$

The first term is the present value of rent earned while the land is undeveloped land (flow rent r_0 which we take to be constant), discounted at rate ρ and accruing up to the date of first development, denoted τ_0. The second term is the present value of rent while in informal urban use during interval $[\tau_0,\tau_1]$. The first formal sector development, occurring at date τ_1, yields rent and incurs a potential one-time fixed cost $D(x)$ capturing frictions to formal development, such as unclear land rights or costs of land preparation; we discuss these further in Section 5.2.5. The final term in equation (5.14) gives the discounted value

[12] Using $k(v_F) = \kappa_F v_F^{\gamma}$ the condition $\mu p v_F \Phi = k(v_F)$ becomes $\Phi = \kappa_F v_F^{\gamma-1}$, and using equation (5.11) to substitute for v_F gives $\mu = \frac{1}{\gamma}$.

of land-rents earned over the lives of consecutive formal sector buildings, constructed at dates $\tau_2, \tau_3, ...$ Equation (5.14) assumes this sequence of development, and we will establish conditions under which this sequence is followed. Our key results emerge from solving this dynamic optimization problem and we proceed in stages, so as to explain results.

Dates of development and redevelopment are chosen to maximize $PV(x)$. For the first development the optimal τ_0 simply equates flow land-rents on undeveloped and informal land,

$$\frac{\partial PV(x)}{\partial \tau_0} = e^{-\rho t}[r_0 - r_I(x, \tau_0)] = 0, \tag{5.15}$$

which, with iso-elasticity and using equation (5.5) is,

$$r_0 = \kappa_I(\alpha - 1)\left[\frac{a_I p(x, \tau_0)}{\kappa_I \alpha}\right]^{\frac{\alpha}{\alpha-1}}. \tag{5.16}$$

Since space-rent $p(x, t)$ is increasing through time, for each place x there is a unique date τ_0 at which informal development commences.

The first formal development takes place at date τ_1 satisfying

$$\frac{\partial PV(x)}{\partial \tau_1} = e^{-\rho \tau_1}[r_I(x, \tau_1) - p(x, \tau_1)v_F(x, \tau_1) + \rho(k(v_F(x, \tau_1)) + D(x))] = 0. \tag{5.17}$$

Using the expression for land-rent (5.12) and formal volume (5.10), together with iso-elastic construction technology this can be written as

$$r_F(x, \tau_1, \tau_1)\frac{\gamma - \rho\Phi(x, i)}{\gamma - 1} - \rho D(x) = r_I(x, \tau_1). \tag{5.18}$$

This gives a unique switching point from informal to formal if $\alpha > \gamma$ (from equations 5.4 and 5.12), which we assume and which is a condition that our estimates of α and γ clearly satisfy.

The first redevelopment of formal land is at date τ_2 satisfying

$$\frac{\partial PV(x)}{\partial \tau_2} = e^{-\rho \tau_2}[p(x, \tau_2)v_F(x, \tau_1) - p(x, \tau_2)v_F(x, \tau_2) + \rho k(v_F(x, \tau_2))] = 0.$$

Generalising this for all redevelopments gives:

$$p(x, \tau_{i+1})[v_F(x, \tau_{i+1}) - v_F(x, \tau_i)] = \rho k(v_F(x, \tau_{i+1}))], \text{ for } i \geq 1. \tag{5.19}$$

Again using equations (5.12), (5.10), and iso-elastic construction technology this can be written as

$$r_F(x, \tau_{i+1}, \tau_{i+1}) \frac{\gamma - \rho\Phi(x, i+1)}{\gamma} = r_F(x, \tau_{i+1}, \tau_i). \tag{5.20}$$

Intuition on these redevelopment dates can be seen from inspection of equation (5.19).[13] This says that demolition and reconstruction occur at the date at which the instantaneous revenue gain from the change in volume equals the interest cost of the construction expenditure incurred. Equations (5.15)–(5.19) implicitly define the dates at which sites are (re-)developed. These equations, together with the definition of the value-to-rent ratio, $\Phi(x, i)$ in equation (5.6), form the basis of the analysis of the next sub-section.

5.2.4 Analysis

What do we learn from the characterization of development stages given above? Assuming space-rents grow at constant exponential rate $\widehat{p} > 0$ yields analytical results, where this growth is driven by the growth in urban wages relative to the outside option, derived from equation (5.1) (see Appendix Section 5.C.1). We look first at the time series development of a particular place x, and then at the urban cross-section.

Urban dynamics at any location

The process of redevelopment of land that has been formalized is summarized as follows:

Proposition 5.1 *If formal sector construction costs are iso-elastic in height (with elasticity γ), space-rents are growing at constant exponential rate $\widehat{p}$, $\frac{\rho(\gamma-1)}{\gamma} > \widehat{p} > 0$, and agents have perfect foresight then:*

[13]The switch dates in equations (5.17) and (5.19) are **not** when flow land-rents are equalized. This is because land-rents (net of amortization) jump as volume increases at date of construction. The terms in γ and $\rho\Phi$ make the appropriate adjustments. Notice that in equation (5.17) the term $\frac{\gamma-\rho\Phi(x,i)}{\gamma-1} = 1$ in a stationary model with $\widehat{p} = 0$ and the first-generation development lasting in perpetuity.

(i) The value-to-rent ratio takes constant value Φ and the time interval between successive formal redevelopments is constant $\Delta\tau$,

$$\begin{aligned}\Phi &= \int_0^{\Delta\tau} e^{(\widehat{p}-\rho)t} dt \\ &= \frac{1-e^{(\widehat{p}-\rho)\Delta\tau}}{\rho-\widehat{p}}, \\ \Delta\tau &= \frac{\gamma-1}{\widehat{p}} ln\left[\frac{\gamma}{\gamma-\rho\Phi}\right]. \end{aligned} \tag{5.21}$$

(ii) Successive rounds of formal sector building have greater volume (height) by a constant proportional factor

$$\begin{aligned}\frac{v_F(x,\tau_{i+1})}{v_F(x,\tau_i)} &= e^{\frac{\widehat{p}\Delta\tau}{\gamma-1}} \\ &= \frac{\gamma}{\gamma-\rho\Phi} > 1. \end{aligned} \tag{5.22}$$

(iii) Land value (capital value of a unit of land) in the formal sector at location x newly redeveloped time τ_i is given by

$$PV_F(x,\tau_i) = \frac{r_F(x,\tau_i,\tau_i)\Phi}{1-e^{-(\rho-\frac{\widehat{p\gamma}}{\gamma-1})\Delta\tau}}, \tag{5.23}$$

equivalent to constant growth rate per unit time of $\frac{\widehat{p\gamma}}{\gamma-1}$.

The first part of this proposition comes from integrating equation (5.8), using it in equation (5.20), and noting that there is a unique solution solving the two parts of equation (5.21) with Φ and $\Delta\tau$ constant over time and space. The second part follows from equation (5.20), noting that, from equation (5.12) the ratio of land-rents at date τ_{i+1} on vintages $i+1$ and i is simply the ratio of volumes. The thrid part gives the capital value of a unit of land in the formal sector at location x at newly redeveloped time τ_i derived by using $R_F(x,\tau_i) = r_F(x,\tau_i,\tau_i)\Phi$ in equation (5.12), and summing over infinitely repeated cycles of redevelopment; the sum is finite if $\rho > \frac{\widehat{p\gamma}}{\gamma-1}$. Flow land-rent $r_F(x,\tau_i,\tau_i)$ increases at rate $\widehat{p}$ during the life of a building, and jumps up by factor $e^{\frac{\widehat{p}\Delta\tau}{\gamma-1}} = \frac{\gamma}{\gamma-\rho\Phi} > 1$ at date of redevelopment, due to the greater building volume. The average growth of the land-value is greater than that of space-rent in anticipation of these jumps in volume. Relationship (5.23) between current land-rent and land-value (reflecting future rent increases) will be used in Section 5.3 to back out an estimate of $\widehat{p}$, the rate of increase in housing rents. Note also that the proposition tells us that, if the rate of growth of space-rent is the same in all formalized locations, then so too are Φ, $\Delta\tau$ and the growth rates of building volume and land-value.

Proposition 5.1 deals with redevelopment of land that has previously been formalized. The timing of informal and first formal development are given by equations (5.15) and (5.17), and hinge on comparisons between r_0, $r_I(x,t)$, and $r_F(x,t,t)$. With Φ now constant from proposition 5.1 we have:

Proposition 5.2 *If formal sector construction costs are iso-elastic in height (with elasticity γ), informal amenity is iso-elastic in volume (with elasticity α), space-rents are growing at constant exponential rate $\widehat{p}$, $\frac{\rho(\gamma-1)}{\gamma} > \widehat{p} > 0$, agents have perfect foresight and $D(x) \geq 0$ then:*

- *(i) Informal development exists for an interval of time $[\tau_0, \tau_1]$, positive if and only if τ^*, solving $r_0 = r_F(x, \tau^*, \tau^*)\frac{\gamma-\rho\Phi}{\gamma-1} - \rho D(x)$, is greater than τ_0.*
- *(ii) If informal development occurs, then it precedes formal development.*

The equation in the first part of the proposition implicitly defines τ^* as the date at which it would be profitable to switch undeveloped rural land directly to formal sector development. If $\tau^* > \tau_0$, then the switch date from rural to informal (equation 5.15) occurs prior to this date, so there exists a period of informal land-use prior to formalization. The inequality under which this holds, $\frac{r_0+\rho D}{\gamma-\Phi}^{\frac{\gamma-1}{\gamma}} \frac{\kappa_F}{\Phi}^{\frac{1}{\gamma}} > \frac{r_0}{\alpha-1}^{\frac{\alpha-1}{\alpha}} \kappa_I^{\frac{1}{\alpha}} \frac{\alpha}{a_I}$. As expected, an interval of informality is more likely the lower is κ_I, higher is a_I, the higher is $\frac{\kappa_F}{\Phi}$, and the lower is r_0.[14] The second part of Proposition 5.2 is proved by noting that r_I is increasing at rate $\frac{\widehat{p}\alpha}{\alpha-1}$ (equation 5.6) and r_F is increasing at rate $\frac{\widehat{p}\gamma}{\gamma-1}$ (equation 5.12). Given $\alpha > \gamma$, there is therefore a unique crossing, switching from informal to formal. The assumption $\alpha > \gamma$ means that there are sharper diseconomies to informal sector crowding than to formal building height, and that the share of land-rent in revenue is higher in informal development than in formal.

Figure 5.1 illustrates these results in a stylized benchmark city without frictions, using model parameters estimated in Section 5.3. Building volume is given on the vertical axis (log units), and on the horizontal plane axes are time t and location x. Location is distance from the centre, and we discuss the cross-section (variation across x at a given t) in the next sub-section. For the moment, look just at the development of a particular location through time, i.e. fix x and look along a line sloping up and to the right parallel to the t axis. Initially (at low t) this land is rural. Building volume becomes positive at date τ_0

[14] r_0 determines the date of edge slum development, but has no bearing on subsequent time intervals between this and formalization or subsequent redevelopments.

(specific to location x) when informal development takes place. The volume of informal development increases steadily, as increasing p causes Meccano pieces to be rearranged and building cover to increase. Formal development takes place at τ_1 and, as illustrated, leads to a small increase in volume, indicated by the second step. Subsequent redevelopments occur at fixed time intervals $\Delta\tau$ and bring the same proportionate increase in volume, achieved by building taller. The timing and volume of each of these formal investments is based on perfect foresight about the growth of prices and the date of subsequent redevelopments.

The urban cross-section and its evolution

We have so far concentrated on a single location, x, and now show how development varies across places in the city. We interpret x as distance from the city centre and assume the existence of commuting costs such that wages net of these costs, $w(x,t)$, decrease with distance from the centre at exponential rate θ. In Appendix Section 5.C.1, we give the underlying assumptions that yield this exponential form. Exponential decline with distance and exponential growth through time imply that space-rents are

$$p(x,t) = \overline{p}e^{\hat{p}t}e^{-\theta x}. \tag{5.24}$$

The switch points in equations (5.15–5.19) can now be interpreted either as giving the date at which place x develops, or the place that develops at date t. The latter interpretation gives the urban cross-section. Proposition 5.3 states how different stages of development vary across this stylized city as it grows, and Proposition 5.4 summarizes how this cross-section evolves through time.

Proposition 5.3 *If formal sector construction costs and informal sector quality are iso-elastic, space-rents are growing at constant exponential rates* $\widehat{p}$, *with* $\frac{\rho(\gamma-1)}{\gamma} > \widehat{p} > 0$, *and declining with distance at constant rate* $\theta > 0$, *and agents have perfect foresight then:*

(i) The distance from the city centre to the edge of new informal development increases through time according to $\frac{dx_0}{dt} = \frac{\widehat{p}}{\theta}$.

(ii) If $D = 0$, *the distance from the city centre to the outer edge of formal development increases through time according to* $\frac{dx_1}{dt} = \frac{\widehat{p}}{\theta}$.

(iii) The distance between successive formal sector redevelopments, Δx, *is constant,*

$$\Delta x = \frac{\gamma - 1}{\theta} ln \left[\frac{\gamma}{\gamma - \rho\Phi}\right].$$

These results are illustrated in Figure 5.1, where we now fix a date and move along a line parallel to the x-axis. At the city edge land is informal. Moving towards the centre, at the inner edge of informal land there is the most recent formalization; and, as we move in, further places that have been urban for longer have been through more stages of development, with a jump in housing volume and height at each stage. As the city grows, the first part of the proposition says new informal sector development is pushing continuously into rural land on the fringe. Similarly, the second part says the ring of most recent formal sector development is pushing continuously into the inner edge of the informal sector ring and that the width of the ring of informal area, $x_1 - x_0$ is constant through time. Hence, even in a circular city, the share of urban land area that is informal declines as the city gets larger. The third part says that the width of formal sector rings of development is constant. We also note that as we move from the city edge to the inner boundary of slums land, housing quality declines continuously as more investment is crowded onto each piece of land. Thus, iso-elasticity and exponential growth yield simple patterns in a benchmark city without frictions.

5.2.5 Spatial heterogeneity and formalization costs

Cities do not look like the neat pattern described above. Places have idiosyncratic features, and our focus here is on those that affect the ability to formalize land. These include institutional and geographic costs of moving from informal to formal settlement, all of which are summarized in our modelling by the parameter $D(x)$, the fixed cost of formalization at place x.[15]

The implications of this spatial heterogeneity are illustrated in Figure 5.2 in which $D(x)$ varies with distance from the centre, taking random non-negative values. Instead of the sharp edges of Figure 5.1, stages of development are now locally diffused because of the variation in $D(x)$. This yields "waves of development" in different greyscale shades in the figure. The leading edge of the formalization wave is places with low D and, within a particular distance band, this is followed a process of infill, with the fraction of the area covered by formal buildings increasing through time. The city cross-section therefore exhibits areas of informality next to formal structures, some of which may have gone

[15]While we have $D(x)$ as a conversion cost that occurs with formalization whether converting land from slums, farmland or vacant land, there may be an additional cost of clearing slums per se, including eviction costs even from private land. For properties where such costs may be high due to history and specific rights, landholders may skip the slum stage and go from farmland to formalization directly. This element is not added to the model here to limit complexity but is noted again in Section 5.5 in interpreting data.

through several phases of redevelopment. In the informal areas, volume and crowding increase. This is the pattern that we see in the data, as will be discussed in Section 5.5.

We know from Proposition 5.1 that the time interval between each redevelopment at a particular place is fixed, so it follows that if first formal development is later and taller, then so too are subsequent redevelopments.

Proposition 5.4 *If formalization costs $D(x)$ vary across places, space-rents are growing at constant exponential rates $\widehat{p}$, with $\frac{\rho(\gamma-1)}{\gamma} > \widehat{p} > 0$, and declining with distance at constant rate $\theta > 0$, formal sector construction costs are iso-elastic, and agents have perfect foresight then:*

(i) Formal development is delayed for places with high $D(x)$, but proposition 5.1 continues to hold for each x.

(ii) Variation in $D(x)$ means that first formal development and subsequent redevelopments may be undertaken in different places at the same date. For such developments, volume (height) and land-value gradients are

$$\begin{aligned} \frac{dv_F(x,\tau_i)}{dx}\frac{1}{v_F} &= \frac{-\theta}{\gamma-1}, \\ \frac{dR_F(x,\tau_i)}{dx}\frac{1}{R_F} &= \frac{-\theta\gamma}{\gamma-1}. \end{aligned} \tag{5.25}$$

The spatial gradients in the last part of this proposition are derivatives of equations (5.4) and (5.5) and will be used in the empirical section of the paper to back-out parameters of the model.

Spatial heterogeneity of formalization costs $D(x)$ supports the observed hotchpotch of adjacent formal and slum buildings. Some of the costs are incurred in securing private property rights (for either leasehold or freehold land) needed to avoid risk of expropriation, to obtain financing and insurance, and to clarify issues such as inheritance, compensation for takings and the like. Other costs arise from the institutional blockages in the long and tortured development of private property rights in Nairobi, as discussed in Section 5.1. Such frictions involving corruption and a dubious legislative history transferring land to the government have been a critical factor in the city's development. These frictions vary widely across areas of the city, as we discuss in Section 5.6.

Spatial heterogeneity in $D(x)$ may be due to geographical as well as institutional factors. Initial construction costs may vary with terrain, for example the costs of draining a swamp or levelling a rugged site for development. We offer empirical evidence for the presence of such effects as discussed in Section 5.5.1 and Appendix 5.E. We show that poorer geography delays initial development. High price land near that city centre undergoing first development has distinctly poorer terrain compared to the city edge, as all better terrain land near the city centre has already been developed. Evidence also suggests that poor terrain affects initial development costs, rather than costs of redevelopment, since it does not reduce building heights on redeveloped formal sector land.[16]

5.3 The built environment of Nairobi

Our empirical work provides evidence on the built fabric of Nairobi and its evolution through time.We use the data and analytical structure to estimate equations characterizing the 2015 cross-section, and these are used in Section 5.4 to solve for model parameters. In Section 5.5, we use the model and parameters to make predictions about the dynamics of the city and then in Section 5.6 to do welfare analysis.

5.3.1 Data and mapping

As noted in Section 5.1 and detailed in Appendix Section 5.D, we capture the characteristics of the built environment of Nairobi at a very fine spatial resolution.We work with 150m×150m grid cells, based on data aggregated from 40cm resolution to 3m×3m cells and to the grid squares that we work with. From aerial photographs, we have tracings of all building footprints in 2003 and 2015 and we have LiDAR height data for 2015 (Geosystems 2015, World Bank 2015). LiDAR records on average 1–3 points per square meter on the ground and is accurate in height to within a meter or less. Each building's height is the average of the LiDAR recordings that intersect with the building's 2D footprint and allows for sloping roofs and differential height in building segments. To infer 2003 heights, we assume that if building footprint is unchanged between dates then so is height. For demolished buildings, we assume 2003 heights equal the average height of unchanged buildings in the eight queen neighbouring grid squares, although that is likely to overstate their height, since demolished buildings are likely less tall than unchanged ones. Each gridcell is also assigned mean elevation and ruggedness based on USGS (2005).

In Figure 5.3, we show two mappings of slums and define the area of the city we work with. The star in Figure 5.3 marks the city centre, from which all gradients we estimate emanate. This centre is the traditional central business district and is defined by the brightest nightlights pixel in Nairobi in 1992. Figure 5.3 shows our economic boundaries

[16]Specifically, in Table 5.14 rebuilding heights are uncorrelated with elevation, presence of water in the neighbourhood, and do not decrease with greater ruggedness.

for the city, a dashed outline for 2003 and a solid one for 2015. These boundaries are based on built cover. For a (150m×150m) grid square to be in the city, the average roof cover in cells whose centroid is within a 900m radius of the cell must be above 10%; and we only keep those cells that contiguously connect to the centre. To focus on private sector development, we remove all grid squares entirely in permanent public uses and amounting to 11% of land in the 2003 city boundary; neighbourhood schools and roads remain (see Appendix 5.D). There is a large extensive margin expansion between 2003 and 2015 of about 60% in land area. Nairobi's sausage shape and the lack of expansion directly north and south of the centre arises because of permanent fixtures (Harari 2020): in the south an airport and a national park and in the north a state forest.

Our analysis requires a distinction between formal and informal settlements (slums). We base our slum mapping on a 2011 classification by IPE under the Kenya Informal Settlements Improvement Program, with details given in Appendix 5.D (Limited & Consultants 2013). The IPE mapping used satellite imagery, topographic maps, and on the ground surveying to define slums as having aspects of lack of planning and lay-out, low house quality, poor infrastructure, or insecure tenure. They also classify slums by land ownership, in particular private versus government, and some residual categories such as riparian and roadside slums, where riparian also includes large tracts of privately owned land and land under possessory rights. Figure 5.15 in Appendix 5.D shows how ownership and slum classification varies with distance from the centre. An earlier delineation of slums comes from a 2003 land use map prepared by the CSUD at Columbia University (Williams & Klopp 2014). While the effective definitions differ in precise detail across the two studies, we use the 2003 mapping to try to distinguish slum areas that underwent formal sector development between 2003 and 2011. While we rely on the 2011 IPE mapping to define slum areas in empirical work, the categorization of formal versus slum buildings and the drawing of slum borders can have fuzzy portions. While in Section 5.2.2, there was a clear general distinction in building materials between the sectors, there can be more of a continuum where, for example, some slum buildings are made of load bearing materials. They are just in crowded areas that are irregularly laid out.[17] Any categorization has issues, and we experiment in robustness checks with other divisions.

In Figure 5.3, slum areas are marked yellow if recorded in both studies, blue if only in the

[17]There are writings about semi-formal buildings, which are built tall of cheap materials and have a reputation in the press of collapse. This sector does not seem to be growing given the bad publicity, although semi-formal construction had a modest boom and the well-publicized collapses. There is no mapping of semi-formal areas, but based inspection of satellite images and on Huchzermeyer (2007) we think slum boundaries largely contain these structures, where the Huruma area Huchzermeyer studied is within the IPE slum mapping we use.

2011 IPE and red if only in 2003 CSUD. We will look at these red areas to show aspects of slum conversion to formal sector use. However, there seems to be little overall slum removal. Areas near the centre with no recorded slums are marked by the dashed and solid rings for respective periods. While the area with no slums expands considerably from about a 0.8km radius to about a 2km radius around the centre by 2011, the removed slum areas are tiny. The map suggests that most slum extensive margin expansion (blue) is near the 2003 fringe of the city and beyond. The well-known large slum of Kibera, which we will discuss later in some detail, is the large slum area directly south-west of the centre (ranging from about 3–5km of the centre).

Data on space-rent and housing characteristics are derived from a georeferenced household data set from the 2012 'Kenya: State of the Cities' survey by the National Opinion Research Center (NORC) and we will use these to infer base formal and slum rents (Gulyani et al. 2012). The NORC used the Census mapping of slums noted above but sampled within those areas to get what they thought were true slum versus formal sector properties; we accept their on-the-ground classification for the sample properties. Second, on prices, we have asking prices for vacant land listed in late 2015, obtained by scraping from Property24 (2015), a website that advertises property for sale in Kenya. All listings fall in the formal sector. In Figure 5.12 of the Appendix, we show a map of the locations of these properties (noting some dots overlap) and the NORC sites. Nairobi has a relatively sophisticated property market. We consider the listing services to give excellent coverage of formal sector land and there is a good scattering of all types of properties in the map.

We have made choices above as to what is the city centre, how slums are defined, and whether there is an important alternative centre. In Section 5.3.2, we perform robustness checks on our results, concerning these choices. Second, we rely on specific sources for land value and rental data. As such there is the issue of selection in the listings for land price data and in NORC surveying of slums, and the potential for omitted variables in regression analysis.We will discuss potential biases below. Third concerns what land value data reflect in markets where there can be bubbles. Appendix Section 5.D.4 analyses this issue for Nairobi, concluding that bubbles are not an issue in the data we use.

5.3.2 Characteristics of Nairobi's built environment in the cross-section

Our cross-section description of the city starts with how the built fabric varies with distance from the centre, from which we derive parameters of the model and evaluate key model assumptions. As well as fitting our model, looking at variation with respect to distance from the centre is standard in the empirical literature (e.g. (Ahlfeldt & McMillen 2018, Combes et al. 2019)). Such regressions take the form

$$\log(y_i) = b_1 + \beta_1 x_i + \beta_c Z_i + \epsilon_i.$$

Outcomes, y_i, include land prices and measures of the built fabric in formal and slum areas. Where the dependent variable is space-rent or land price observation i denotes the household or vacant lot; where it is a feature of the built fabric observation i denotes the 150m×150m grid square. x_i is distance to the city centre so that β_1 is the gradient. Following the model, but also the literature, regressions are semi-log, so a gradient coefficient can be interpreted as proportionate change per kilometre distance from the centre. b_1 is a constant term. Z_i are outcome-specific relevant control variables discussed below. For the error term, omitted variables will be a concern.

Results on estimated relationships are presented in figures and in tables. The figures for this section cover the 2003 bounded area of the city, cut at 10km from the centre, based on observations at the 150m×150m grid square level. Regressions cover all 6,470 grid squares of the 2003 city. For this section, regression results are presented in Tables 5.1 and 5.2 below where all regressions contain controls on grid square elevation and ruggedness. Regressions for prices have further controls, and tables report the gradients and an "intercept" at the city centre which is $b_1 + \beta_c \overline{Z}_i$ for $\overline{Z}_i$ evaluated usually at median values for the city. More details are in table notes and coefficients on all covariates of key regressions are in Appendix Table 5.10.

Formal sector

Intensity of land use is measured by volume of built space in cubic meters (m^3) per square meter (m^2) of land area, which we call the built volume to area ratio, BVAR and correspond to $v_i(x, t)$ in the theory. The gradient for BVAR is given in column 1 of Table 5.1 and illustrated for the raw data with a smoothed fit in Figure 5.4, which shows the confidence intervals and the 25^{th} and 75^{th} percentiles of observations.[18] Formal sector BVAR near the centre is very high, averaging around $8m^3$ of space per m^2 of ground area (where ground area includes roads and minor public spaces) in Figure 5.4; it declines at a rate of 4.9% per

[18] While we have the universe of observations, smoothing requires estimation to show a continuous gradient. As such, the displayed confidence intervals on the mean are extremely tight. To show the dispersion, we plot smoothed estimates of the 25^{th} and 75^{th} percentile of observations at each distance. Smoothing involves grid squares whose centroid is in a 300m moving window going out from the centre.

km from the centre from Table 5.1, column 1. The 25^{th} and 75^{th} percentile of observations. Note there is heterogeneity across grid squares in BVAR's at each distance, reflecting, in part, differential timing of formal sector development of grid squares, based on differential geography and the specific histories of properties' paths to formalization, as well features such as roads.

BVAR is the product of height and the cover-to-area (CAR, or $c_i(x,t)$ in the model), which have following gradients. In Table 5.1, column 3, average heights of buildings decline with distance from the centre at a rate of 7.6% per km. This compares with height gradient numbers for Chicago in the 1960s or 70s from Ahlfeldt & McMillen (2018).[19] As Figure 5.5 shows, there are very tall buildings at the centre in Nairobi; at 0–1km from the centre average grid square height (based on 3m$\times$3m pixel) is about 10 storeys (at 3m a storey); the data indicate that 6% of the grid squares are over 16 storeys. For CAR, we assumed in the theory that the formal sector has constant cover to area ratio throughout the city. The gradient in Table 5.1, column 2 has a zero slope; and the solid blue line in Figure 5.5 indicates that CAR does not change much across the city and is not decreasing with distance from the centre. The decline in formal sector BVAR is thus due to height rather than cover.

To recover formal sector parameters in line with the theory,we focus on the gradient for heights of newly redeveloped buildings, as given in equation (5.25).[20] We do not know building ages per se, just whether buildings are new since 2003. To pick the most recently redeveloped buildings from 2003 to 2015 interval, we first take the sample of redeveloped buildings from 2003 to 2015. Based on the theory where at any distance heights of new buildings will rise with time, we focus on the 80th percentile, balancing out wanting the most recent and hence taller buildings, against getting extreme outliers. Column 4 estimates the gradient for redeveloped buildings using a quantile regression. Fortunately, for redeveloped buildings, different quantiles, and OLS estimates of the gradient slope are very similar.[21] From Table 5.1, column 4 the slope is 10.1% per km, noticeably larger in absolute value than for the overall stock.

What are identification issues for these regressions on the built environment in cols. 1–4? The data are for the universe within the 2003 city, so there is no sample selection. The

[19] The authors use different functional forms and look just at tall buildings. But at say 5km from the centre their slope estimates are similar. However, for later years they have steeper slopes.

[20] Note we cannot do redeveloped BVAR, since we cannot properly assign land area to new vs. old buildings.

[21] Gradient coefficients on redeveloped buildings are stable, with OLS and quantile regressions up to the 90^{th} percentile all having gradients in the neighbourhood of 0.10. But intercepts rise, as we raise the percentile level.

concern would be that there are geographic conditions (affecting the D's in the model) which could be correlated with distance to the centre and affect gradients, which is why we control for elevation and ruggedness. That said, the two key equations we use to parameterise the model, the slum BVAR and the formal sector 80$^{\text{th}}$ percentile of height quantile regression, should not be affected by D magnitudes, according to the model. Correspondingly, whether we control for ruggedness and elevation or not, the estimated coefficients are stable.

Turning to price and rent gradients, column 6 of Table 5.1 uses scraped land price listings from the fall 2015 Property24 (2015) data to estimate the formal sector land price gradient with respect to distance from the centre, with asking prices (per square meter) in USA$. An exponential form captures the relationship well, and the gradient is steep, with price declining by 17.2% per km of distance from the centre. Such a rate of decline means land prices vary almost six-fold from the centre out to 10km. This slope is right in the middle of the range of land price gradient slopes across French cities in Combes et al. (2019).

For column 6 of Table 5.1, identification issues are a concern. There are two issues, sample selection, and bias in estimates within the selected sample due to unobserved property features affecting prices that could be correlated with distance to the centre. On selection, while we believe most lots in the formal sector are listed, we do not know the universe of properties for sale. Figure 5.12 in the Appendix shows a good scattering of listing points throughout the city. If there are unlisted properties, they may be different, on more challenged land or land with worse amenities, both of which could vary with distance from the centre. For example, if these omitted properties are first formal sector developments, we will show later that thesewould be onworse quality sites nearer the centre (last to be developed). In this case, selection could bias our estimate of the intercept upwards and our estimate of the slope downwards (more negative). On the issue of omitted variables, we control for plot area, ruggedness, elevation, and whether the listings have an exact street address (which many properties do not) as an indicator of a quality differential. However, unobserved amenities could be better or worse near the centre, and the direction of bias is ambiguous. In short, depending on the relative extent of selection and omitted variable biases, our estimate of the land rent gradient could be biased either way.[22]

The final column in Table 5.1 estimates the space-rent gradient with data from the NORC, using hedonic price regressions for space-rents. Column 7 reports gradients and intercepts

[22] Using reports from HassConsult, an alternative set of vacant land listing prices, the estimated gradient from regressing ln price per square metre on distance gives a similar slope 0.198 (0.0309), within one standard error of that estimated here. See Appendix Section 5.D.4 for details.

in 2015 USA$ for m^3 of volume.[23] We will utilize the intercept only in work below; this space-rent gradient can be sensitive to specification unlike other columns. These regressions control for house characteristics in order to define rents per unit volume for a typical formal sector house; full hedonic regressions are reported in Table 5.10. Table 5.1, column 7 gives the gradient of formal sector space-rents, which is flatter than that of land-rents as suggested by the theory. In column 7, its slope of 8.6% per km is close to what we back out of the model (7.1%) in Section 5.4.1 below.

Slum areas

Table 5.2 gives gradients for the slum sector. Slum areas should exhibit a decline in volume with distance from the centre, with the decline driven by diminishing cover per unit area, not by a change in height. Table 5.2 and Figures 5.4 and 5.5 give results. Column 1 of Table 5.2 indicates BVAR declines significantly with distance from the city centre at a rate of 9.5% per km. In column 3, slum heights are flat in terms of distance from the city centre as assumed in the model and as pictured in Figure 5.5 by the red-dotted line. With constant slum height, gradients for slum BVAR in column 1 and slum cover to area ratio, CAR in column 2 should be the same; the estimates of 0.0948 and 0.103 in Table 5.2 cols. 1 and 2 confirm this.[24] Slum CAR near the city centre in Figure 5.5 is very high at 50% or so, more than twice the 25% number for the formal sector. This means that slums have little green/open space around houses, with attendant loss of amenity.[25]

The final column of Table 5.2 gives the slum space-rent gradient. Note that in the model and in the data, typical formal and slum sector houses are not comparable; and we have separate gradients for each. The slum space-rent gradient is flat, in contrast to that in the formal sector. In equation (5.4) of the model, observed space-rents are constant throughout the city; the insignificant coefficient of 0.0094 validates the modelling that gives this result. The intercepts of the space-rent equations in the formal and slum sectors in Tables 5.1 and 5.2 are, respectively, 3.148 and 1.886. These will be used later in welfare analysis to infer land-rent differentials by sector. The formal-slum differential in space-rents reflects different amenity values, incorporated in the model by the parameters a_i.

Again, there are issues of identification for this NORC sample. On selection, Figure 5.12 in the Appendix shows a good scattering of points throughout the city for formal and

[23] Calculations assume each m2 of floor space yields $3m^3$ of volume (given typical ceiling height) inflates 2012 nominal rents at 8% a year based on reports from www.hassconsult.co.ke, converts monthly to annual rents and Kenyan shillings at the 2015 exchange rate of 100KS to a USA$.

[24] Note the slopes for slum CAR and height (0.103 + 0.0107) sum to almost that for BVAR, 0.0948.

[25] We also have paved road surface from high-resolution SPOT satellite data for 2015. Formal sector roads are about 15% of area near the centre while in slums they maybe hit 5%.

slum housing, covering newer and older slums within the 2003 city. Households were randomly sampled and stratified based on the 2009 Census definition of slums. On bias and unobserved attributes that could vary systematically with distance from the centre for developed formal sector and slum housing, we simply do not know the likely direction of bias. We also note that features like ruggedness and elevation on average are the same across sectors with insignificant differences.[26]

A further point of interest in Figure 5.4 is that, at 2km and beyond, slums and the formal sector deliver generally the same BVAR. For the opposing views of whether formal sector height trumps slum coverage in providing volume of built space per unit land, in Nairobi, they do equally well on average, albeit at different quality levels. However, while each sector has similar BVAR, slums occupy a smaller fraction of the city's surface area and hence provide just 10% of total building volume.

Robustness

We perform robustness checks on various definitions or classifications, we imposed on the data. The first concerns the city centre being chosen as the brightest nightlights' pixel in Nairobi in 1992. Alternatively, using the current nightlights from NASA's Visible Infrared Imaging Radiometer Suite to identify the top 1% of grid-cells readings in Nairobi, we can define the centre as the brightest pixel in the largest cluster of these top 1% pixels (Appendix Section 5.D.4). That moves the star in Figure 5.3 by about 400m. The second check concerns the presence of a second major centre in Nairobi around which gradients might orient. The most likely candidate is a large industrial area to the east and somewhat south of the centre. We experiment with specifying this as a second centre influencing or even potentially dominating the gradient from the city centre that we estimate. Finally, there is the definition of slums. Alternative to IPE in 2011 is a 2009 Census map of slums classifying census tracts as slums if they are "unplanned"; this results in a looser definition of slums and larger tracts, as can in seen in Figure 5.11 in the Appendix (Kenya 2009). We will experiment with using this definition of slums.

We report robustness results for the five equations that we use below to derive parameters of the model: the land, slum rent and formal rent price equations and the equations for slum BVAR and formal sector redeveloped height. For the slum BVAR and formal sector height equations, these involve adding a control for distance to the main industrial centre to the base equations, using the alternative centre point based on recent (not 1992) nightlights in

[26] Mean [SD] log elevation in the slum and formal sector are 7.43 [0.045] and 7.40 [0.028] and log ruggedness are 1.58 [0.54] and 1.43 [0.57], respectively.

the equations, and using the Kenya (2009) Census definition of slum areas (as unplanned settlements and having density over 3000 per sq. km). For the price gradients, we examine the impact of having a differently placed city centre and having a second centre. We do not change slum area definitions for the price equations. Land listings cover only the formal sector; and we accept that, in their sampling, NORC have correctly defined on the ground allocations of households to slum and formal sector environments. The results are reported in Table 5.12. The new city centre definition has a negligible impact on all estimated gradients and intercepts for the five equations. Adding distance to the industrial centre to the base specifications adds a covariate that is always insignificant, has a small impact on all intercepts, and only has a noticeable effect on the CBD distance gradient in one case, formal sector rents. However, we do not use the gradient for that outcome to solve the model parameters below, only the intercept which is minimally affected. Finally, using the census definition of slums similarly has a negligible impact on gradients and intercepts for slum BVAR and redeveloped heights.

5.4 The data and the model

We now derive parameters of the model, using 2015 cross-section information from Tables 5.1 and 5.2 and relationships in the model. The parameter estimates are of intrinsic interest, and we compare them with other estimates in the literature, where available. In Section 5.5, we use the parameterized model to predict dynamic changes between 2003 and 2015 and compare these with what we find in the data.

5.4.1 Solving for the cross-section built environment

The model, in cross-section, is fully determined by parameters describing technology and amenity, γ, α, κ_I, a_I, and $\frac{\kappa_F}{\Phi}$, together with knowledge of space-rent at the centre and its gradient, $p(0,0)$ and θ. (Formal sector choices depend on $\frac{\kappa_F}{\Phi}$; here we calculate the ratio, and in Section 5.4.3 separate the components.) Relevant values and model equations are set out in Table 5.3 which divides into two parts, the upper block giving gradients of variables with respect to distance from the centre and the lower the levels of parameters and variables at the centre, $x = 0$ (i.e. intercepts reported in Tables 5.1 and 5.2).

Looking first at Table 5.3 panel a; for the formal sector, we use gradients on newly developed buildings, taking the distance coefficient of new building height (Table 5.1, column 4) as the measure of how volume varies with distance, and similarly, land values as the measure of $R_F(x, \tau_i)$, the present value of land-rent that accrues over the life of a structure. For slums, the volume measure is BVAR; there is no land value measure available since slum land is not widely transacted. The three equations in this part of the table give values for parameters γ, α, and θ.

Table 5.3 panel b equates intercepts from the regressions with expressions for the level of variables from the theory. These parameters play a role in model simulations where we need levels. For formal sector volume of new buildings, we use height times the median cover of 0.247 throughout the city. This is in preference to the intercept in column 2 of Table 5.1 (which gives a CAR of 0.188) but has low explanatory power given the undulating but basically flat relationship for CAR in Figure 5.6 below. For slums with their malleable capital all that matters is BVAR. The final row of the table uses the space-rent data and together with values of γ and α, they give values of κ_I, a_I, $\frac{\kappa_F}{\Phi}$, and $p(0, 2015)$.

5.4.2 Discussion

Technology and amenity parameters, γ and α, come from cross-section gradients and give diminishing returns in formal sector construction costs and slum crowding, respectively. They also imply shares of land-rent and construction cost in spending on housing. The value of $\gamma = 1.70$ implies that the share of land-rent in formal sector revenue, $1 - \frac{1}{\gamma}$, is 0.41, while $\alpha = 3.98$ implies a corresponding share in the informal sector, $1 - \frac{1}{\alpha}$, of 0.75. The formal sector share of 0.41 is similar to that in Case (2019) for the USA but higher than the 0.30 for Paris in Combes et al. (2017). For slums, we know of no data to make comparisons, since typically slum lands are not transacted privately. However, given the low construction costs of slum housing, such a high land-rent share seems reasonable. Note that $\alpha > \gamma$ is the condition for the model to predict that slum land-use at the city edge precedes formal development.

Land-rent is the product of space-rent per unit volume (\$ per m^3), space per unit land (BVAR m^3 per m^2) and the share of space-rent that is attributable to land, $1 - \frac{1}{\gamma}$ in the formal sector and $1 - \frac{1}{\alpha}$ in slums. We find that, at the centre in 2015, $x = 0$, and $t = 2015$; $r_F(x, t, t) = \$65.38$ and $r_I(x, t) = \$17.67$ (using equations (5.5) and (5.12) with numbers from Table 5.3). Moving away from the centre these imputed land-rents drop at respective proportional rates $\frac{\theta\gamma}{\gamma - 1} = 0.17$ per km and $\frac{\theta\alpha}{\alpha - 1} = 0.095$ per km, nearly twice as fast in the formal sector as the informal. The wide gap in land-rents near the centre is indicative of inefficient land-use (an argument that we will develop precisely in Section 5.6), although we note that from equations (5.15) and (5.19), efficient switch points are generally not at the point of equality of r_F and r_I.

What are the components of these changing land-rents? In the formal sector, the distance gradient is driven principally by commuting costs reducing net wages and hence space-rents (as given by equation (5.1)); land-rents fall faster than space-rents because of declining volume per unit land. For slums, as we saw above, observed space-rent is constant as lower commuting costs near the city centre are offset by greater crowding; the decline in land-rents is therefore due to lower volumes.

There are two other predictions of interest. The amenity derived from a unit volume of slum housing at the centre is 28% of a unit of formal volume; slum crowding decreases with distance, so this fraction rises, reaching 58% at 10km.[27] This change reflects the quality continuum aspect of slums. The construction cost of a unit volume of slum housing is 12% that of formal sector construction cost per unit volume at the centre (cost expressed on an annual amortization basis); formal construction costs per unit volume decrease with distance as buildings get less tall, so at 10km this fraction reaches 25%. As implied, formal sector unit construction costs per unit volume at 10km are 50% of what they are at the centre.

5.4.3 Solving for the dynamic environment

The city is characterized by changing land-use and periodic redevelopment, as is modelled in the theory. Quantifying this requires two further pieces of information. One is the rate of interest, ρ, and we use the real interest rate $\rho = 0.057$, which is the average of the World Bank real interest rate (lending rate adjusted for inflation) for Kenya for the 14 years post 2002.[28] The other is the expected rate of increase of space-rent, $\widehat{p}$. We have estimates of formal sector land-rents, $r_F(x,t,t) = \$65.38$ ($x = 0$, $t = 2015$, from previous sub-section) and the capital value of land at date of redevelopment, $PV_F(x,\tau_i) = \$1414$ (=exp(7.254), Table 5.1, column 6). Equation (5.23) links this capital value to the level and rate of increase of flow rents, dependent on price growth, and hence enables us to back-out the expected price increase. Using these numbers and the value for γ from Table 5.3, we solve forΦ, $\Delta\tau$, and $\widehat{p}$ in equations (5.21–5.23), restated here as:

[27] The amenity parameters $a_I = 0.734$ and $a_F = 1$ are constant.
[28] data.worldbank.org/indicator/FR.INR.RINR?locations=K

$$
\begin{aligned}
\Phi = \int_0^{\Delta\tau} e^{(\hat{p}-\rho)t} dt = \frac{1 - \Delta\tau e^{(\hat{p}-\rho)t}}{\rho - \hat{p}}, \\
\Delta\tau = \frac{\gamma - 1}{\hat{p}} ln \left[\frac{\gamma}{\gamma - \rho\Phi} \right], \\
PV_F(x, \tau_i) = \frac{r_F(x, \tau_i, \tau_i)\Phi}{1 - e^{-(\rho - \Delta\tau \frac{\hat{p}\gamma}{\gamma - 1})}}
\end{aligned}
$$

Solving these equations gives $\hat{p} = 0.0092$, $\Phi = 20.64$, and $\Delta\tau = 89.29$ years. Thus, real space-rent appreciation is estimated to be about 1% a year. This is broadly consistent with the space-rent series produced by HassConsult as described in Appendix Section 5.D.4. The value-to-rent ratio, $\Phi = 20.64$, is close to that implied by an infinite stream discounted at 5.7%, and is in the middle of range of ratios reported on the internet by realtors for US cities.[29] The length of life of new buildings, 89 years, is novel and an estimate for which we know of no easy comparison.

5.5 City dynamics: the data and the model

As reported in Section 5.3.1, we have tracings of all building footprints for 2003 and 2015, and LiDAR height data for 2015. In this section, we use these data to describe the patterns of change that occurred in the city, and then compare these actual changes with predictions of the model. As we have seen, the model is calibrated on 2015 cross-section data, so using the model to predict the timing and extent of changes that occurred in the city is a demanding comparison.

5.5.1 Development and redevelopment in the formal sector

We look first at the formal sector, where the model predicts that there is a wave of development, followed by a wave of redevelopment (Section 5.2.5 and Figure 5.2). Figure 5.6 is the empirical counterpart, giving formal sector volume change (2003–15) per unit area at each distance and its decomposition into two parts, infill and redevelopment out to 12km of the 2015 city. Two distinct waves of change are apparent. The wave further out, peaking at around 9.5–10km and significantly higher than other rates, is largely infill, defined as new buildings not intersecting any 2003 buildings. This is the empirical counterpart of formalization; it is construction of formal structures at the edge of the city, some of them on land temporarily occupied by low-density slums, and some on land previously

[29] For example, smartasset.com/mortgage/price-to-rent-ratio-in-us-cities

in non-built use.[30] The wave nearer the city centre peaks at 3–4km from the centre and is principally redevelopment, defined as new buildings that overlap the footprint of 2003 buildings, where redevelopment volume is net of what was demolished.

These two waves are as predicted by the model. How do they line up in terms of exact timing and location? Table 5.4 gives timings of formalization and first redevelopment predicted by the model using values of parameters that are reported in the text above. The first row gives date of formalization, and the model predicts from equation (5.18) that, if $D(x) = 0$, the 2003–15 peak would occur at around 12–12.5km; the data in Figure 5.6 have peak at 9.5–10km, i.e. the model is predicting formalization slightly early. This is consistent with generally positive values of $D(x)$ that delay formalization. The predicted date of first redevelopment in Table 5.4 at 3.5km is about 2030, while the distance band predicted to redevelop during the period 2003–15 is around 0.5–1km from the CBD (Table 5.4, second row). The model is therefore predicting first formal development later than the wave that occurs in the data. The discrepancy arises as the predicted distance between places experiencing formalization and first development is around 11.5km (89 years times $\frac{\hat{\rho}}{\theta}$), whereas the peaks in Figure 5.6 are just 7km apart. The "late" first redevelopment predicted by the model could—amongst other things—be attributable to using the World Bank discount rate of 5.7%; a higher rate reduces the time interval between developments, $\Delta\tau$, and hence also the distance between successive building peaks. Notwithstanding these remarks, we note that Table 5.4 gives first formalization in the city at 1917, surprisingly close to its date of founding, 1899. Of course, such long-run projections are way out of sample and highly speculative, given historical changes through that time period.

A critical aspect of redevelopment in the data and in the model is the change in building height that it brings about. Figure 5.7 splits out the formal sector height gradient in Figure 5.5 to show how the mean heights of redeveloped buildings compare with those that did not change. In the formal sector beyond 1.5km redeveloped buildings are significantly higher, nearly double the height of unchanged buildings in the interval 2–4km. To get a number conceptually comparable to the model, we compare the height of the 80$^{\text{th}}$ percentile of redeveloped buildings, mimicking heights of newly completed buildings, to the 20$^{\text{th}}$

[30]This formal sector development near the city edge occurs as the fraction of land covered by formal buildings increases between 2003 and 2015, as discussed in Section 5.2.5. Not all of this land may have been occupied by slums prior to formal development. However, since we rely on one slum map in 2011, we do not know the formal status of all of this land in 2003. Costs of evicting slum tenants may also inhibit landowners from having land in slum usage close to the time of first formal sector development. Note also that "infill" can be on developer assembled lots where vacant land is built on and formally covered land is left as green space, so it is not truly infill but a rearranging of how land is used on assembled lots. There is also a small residual category (not shown) of torn down buildings with no new footprint appearing (yet) over the old.

percentile of unchanged buildings, mimicking the height of those just about to be torn down. This is graphed in the figure and regressions on these slopes are reported in Table 5.1, cols. 4 and 5. From 1 to 4.5km, the regression (using intercepts and slopes) and the figure both suggest that building height increases around 2.75–3 times with redevelopment. This compares with a model prediction given by equation (5.22) where volume increases by a factor of 3.23, at each redevelopment.

Role of terrain

An element of $D(x)$ which can be observed and that bears on the timing of development (Section 5.2.5) is the terrain. In Appendix 5.E (Land Quality), we report the results of investigating this with water (river, wetland, pond) and ruggedness as independent variables at various spatial scales.[31] There are two main findings. Figure 5.6 indicates infill occurring at all distances, including close to the centre. Table 5.13 in the Appendix shows that formal sector infill near the centre is on much worse quality land than already developed buildings, indicating delay in the date in which poor terrain land is first developed.

The effect falls rapidly with distance so, by 10km out, there is no clear difference between infill and already developed areas. Second, Table 5.14 shows that the height of redeveloped buildings is unaffected by underlying terrain fundamentals. Overcoming terrain involves a sunk cost of land improvement that does not affect the height of structures that are then constructed (i.e. κ_F in equation (5.10) is not affected).

5.5.2 Slum dynamics

The model predicts that, if the condition given in the first part of proposition 5.2 is satisfied, new slums will form on the edge of the city. That near the city edge there will be conversion of existing slums into formal sector use. That existing slums will become increasingly crowded. And that there will be no slums remaining near the city centre unless there are substantial formalization costs, $D(x) > 0$.

The data are broadly consistent with these predictions. First, there are new slums. We have a proxy measure, which are areas identified as slums in the 2011 mapping but not in the 2003 mapping. As noted above, some of the differences are due to different methodologies employed by researchers in defining slums. Nevertheless, in the blue areas in Figure 5.3, we see large tracts of new slum development at and beyond the 2003 city edge. Referring back to dates of first formalization in Table 5.4, we see that in the early 2000s these should

[31] Ruggedness is defined as the standard deviation of elevation across the 30m×30m grid squares of the 90m×90m square neighbourhood which is queen neighbours plus the own square.

be at some distance beyond 10km (i.e. beyond places formalizing). In Figure 5.8, we show the change in slum volume for slums as mapped in 2011 and the figure indicates that at about 8km there is one peak of infill (new slum cover) and then another at 10.5km. The difference between these two peaks is interesting. The peak at 8km is on government owned slum lands where we presume D's are too high to allow formal use yet and slum usage is intensified as prices rise. The peak at 10.5km is disproportionately on private land, so it is on 'market driven' slum expansion. Figure 5.15 in the Appendix displays these ownership patterns.[32]

Second, there are some slum conversions we might infer from the 2003 and 2011 mapping of slums. In Appendix Figure 5.13, we show rates of building teardown for 2011 slums (non-converted) and slums that disappear between the 2003 and 2011 mappings (converted). Teardown rates for these converted slums average nearly 60%, three times greater than in non-converted areas. Appendix Figure 5.14 also shows that buildings on converted slum lands are generally built significantly taller than those in current slums.

Third, there is densification of existing 2011 slums within the 2003 city-boundary. In Figure 5.8, we show total slum volume changes and their decomposition into "redevelopment" and infill. Up to about 6km, "redevelopment" is akin to rearrangement of the Meccano parts and is conceptually not really distinguishable from infill. Both are part of a process of increasing crowding in slums, as CAR increases. Figures 5.5 and 5.7 indicate, as assumed in the model, that there is essentially no increase in slum height, which remains low and uniform across the city. Volume changes are driven entirely by CAR changes.

Finally, no slums get created near the centre, but there is persistence of some large slum areas, as we discuss in detail in Section 5.6.

5.5.3 Total volume and population

What are the implications of these changes for aggregate building stock and population? We have data on each of these for 2003 to 2015. As noted in Section 5.1, total built volume within the 2003 boundary increased by 47% between 2003 and 2015, and that in the 2015 boundary by 60%. Based on the Kenya (1999) and Kenya (2009) censuses for the 2003 area of the city, population grew at 3.8 p.a. which, for the period 2003–15, would project to 57% for 2003 boundary and 68% for the 2015 city boundary.

In the model, volume comes from integrating over all x at a particular date, weighting

[32] In addition to purely private, we are including some riparian lands as graphed in Appendix Figure 5.15. From Kenya (2018), we know these areas contain private developments, as well as areas with possessory rights.

volume per unit area by area at the place.[33] For population, volume at each space is divided by each household's demand for space, $s_i(x,t)$, $i = I, F$, so takes the form

$$L(t) = \sum_{i=1}^{imax(t)} \int_{x_{i+1}(t)}^{x_i(t)} \frac{n(x)v_F(x,\tau_i)}{s_F(x,t)}dx + \int_{x_1(t)}^{x_0(t)} \frac{n(x)v_I(x,t)}{s_I(x,t)}dx. \tag{5.26}$$

In this expression, $n(x)$ is land area (in a circular city, proportional to distance x from the centre). The first term integrates over land in its i-th generation of development at date t (i.e. land in interval $(x_{i+1}(t),\ x_i(t))$, and sums over generations up to that which has been redeveloped the most times, denoted $imax(t)$. Applying this equation to volume (i.e. setting $s_i(x,t) = 1$) gives the volume increase along a ray $(n(x) = 1)$ over the 12-year period of 38%. Population increase exceeds volume as rising income and space-rents create income and substitution effects in household demand for space, with the latter dominating. This gives a population increase along a ray of 56%. In a circular city each of the rates of increase are lower, at 27% and 34%; this is because the model predicts that redevelopment during this period is close to the centre than it is as noted above, i.e. covering a small land area.

To conclude this section, we remind the reader that parameters of the model were estimated entirely on cross-section data, together with an imported value of the discount rate. The model does well in predicting many aspects of Nairobi's development through time, albeit with the proviso that the long time period between developments has the effect of making the model predict redevelopment later than is in the data.

5.6 The cost of delayed formal sector development

Finally, we turn to quantifying some of the costs of persistent slums in central parts of Nairobi. Chief amongst these is Kibera, often described as the largest slum in Africa. The costs and benefits of such slums, and policy towards them, is a complex and contentious issue. We do not seek to quantify all elements (such as community dislocation), and focus on a single element of the equation, namely the real income loss due to inefficient land-use. This is measured by the loss in land values associated with informal settlement on high opportunity cost land.

[33] This is integrating under the surface of Figure 5.2 (anti-logged) at a particular place, and weighting by the area of each place x.

5.6.1 Formalization costs

Nairobi was founded in 1899 and the British colonial government, as was typical throughout Africa, housed the African population in informal settlements without land title (Olima 2001). After independence in 1963, a series of reforms resulted in over 85% of land in Nairobi becoming privately owned under charges of widespread corruption (Southall 2005). However, older slums mostly within 6km of the centre remain 'government owned'. These areas are not managed by the government but rather by slumlords who operate 'illegally' and make high profits. Gulyani & Talukdar (2008) estimate payback periods on an investment in a single room of just 20.4 months. This is consistent with the fact that land is "free" to slumlords; and, by our calibration, land's share in revenue in the informal sector is 75%. Moreover, slumlords have characteristics that are problematic for formalization. In Kibera, for example, of 120 slumlords surveyed, 41% were government officials, 16% (often the biggest holders) were politicians, and 42% were other absentee owners ((Syagga et al. 2002) as cited in Gulyani & Talukdar (2008)). The political economy issue is that if the government were to auction the land it "owns" (or give it to the tenants), the slumlords would lose their claim to the revenue. Having well-connected bureaucrats and political figures opposed to formalization presents a political problem.

This problem is magnified when there are historical private claims to the land, as is typically the case. Kibera gives a nice example. The 1000 acres in Kibera were awarded to Nubian soldiers for service in 1912 by the British. They occupied a portion of the land, but at independence their claims (but not tenancy) were revoked and land reverted to the government. The majority area of present day Kibera which is not occupied by Nubians and their descendants was settled by others, on the basis of claims illegally allocated by local chiefs and bureaucrats. The moral right of the Nubian descendants to at least the land they occupy is well recognized but the unwillingness to grant them title has historically been a road-block to redevelopment (Etherton 1971, Joireman & Vanderpoel 2011). The literature has similar stories for other major slums in Nairobi.[34]

[34] For example, Mathare 3km northeast of the centre was originally a stone quarry. When the quarry closed the land went to the Department of Defence and the area was occupied by squatters. There then followed a long history of squatters attempting to set up collectives to "buy" the land in competition with land-buying companies, dissolution of the cooperatives, allegations of corruption, and competing claims on the land (Medard 2006). Today the majority of the slum part of Mathare (about 1 km^2) is private, but significant portions remain under government ownership of some sort (police, central government, Nairobi City Council). Syagga et al. (2002) analyses how Kenyan tenure legalization can take decades to implement due to the need to reconcile interests of stakeholders and offers more examples such as the Korogocho slum.

5.6.2 The loss of land value due to slum persistence

In our open city model, the welfare costs of inefficient land use are given by the present value of land-rents foregone. To assess the cost of delayed formalization, we use the model and estimated parameters to calculate the present value of land-rents earned by land in different uses. The gap between the values of land in formal versus slum use measures the real income loss due to inefficient land-use, although it does not capture other costs and benefits outside the model, such as social costs of disruption involved in slum redevelopment and community dislocation, or possible productivity benefits of spatial reorganization of the city.

For a unit of land at place x, we calculate the present value, in 2015, of being held in slum use until conversion to formal use (with $D(x) = 0$) at date z, and denote this value $PV(x, z)$. For perpetual delay versus conversion in 2015 the expressions are, respectively

$$
\begin{aligned}
PV(x, z = \infty) &= \int_{2015}^{\infty} r_I(x, t) e^{-\rho(t-2015)} dt \\
&= \int_{2015}^{\infty} r_I(x, 2015) e^{[\frac{\hat{p}\alpha}{\alpha-1} - \rho](t-2015)} dt
\end{aligned}
\tag{5.27}
$$

$$
\begin{aligned}
PV(x, z = 2015) &= \sum_{i=0}^{\infty} R_F(x, z + i\Delta\tau) e^{-\rho i \Delta\tau} \\
&= \frac{r_F(x, 2015, 2015)\Phi}{1 - e^{-\Delta\tau(\rho - \frac{\hat{p}\gamma}{\gamma-1})}}.
\end{aligned}
\tag{5.28}
$$

These present values are given in Table 5.5 as a function of distance from the centre, x. The first row gives the present value if formalization occurs at date $z = 2015$, and the second row gives the present value if it never occurs. For each distance reported in the table, 2015 is beyond the efficient date for formalization given in Table 5.4. We compare the cost of never formalizing with formalizing in 2015 for each distance, $PV(x, z = 2015) PV(x, z = \infty)$. These range from \$920 per m^2 in the 0–1km distance band to \$249 per m^2 at 6–7km from the centre.

We illustrate the gains from conversion by looking at lands 3–4km from the centre, which includes some parts of Kibera. At 3–4km the cost of perpetual informality as compared to switching to formal sector use in 2015 is row 1 minus row 2, or \$774\$284 = \$490 per m^2. There are 1.07mn m^2 of slum land in that distance band. There is thus an aggregate gain from converting at 2015 compared to perpetual delay of \$525mn. We give three perspectives on this.

First, with so much money on the table, why is land at 3–4km not converting from slum to formal use at a faster rate? We have argued that there are financial costs and political barriers to formalization, and a lower bound on the monetary value of these costs is given the bottom row of Table 5.5. This number is per m^2, and is derived by solving equation (5.18) as an equality. At 3–4km this is \$248, equal to 32% of the present (market) value of land in formal use and 87% of the value in slum use, clearly non-trivial amounts consistent with our political story on Kibera. If there were slums in the centre the lower bound on D would be \$493 but, since there are none, we know that the pressure to redevelop was so intense as to overcome formalization barriers.

Second in terms of the one-off gain in land values from conversion at 3–4km, the surplus of \$525 mn is about 10% of Nairobi's GDP in 2015.[35] If we add all the net gains out to 7km, that is about \$2.5bn or 45% of Nairobi's GDP.

The third perspective hints at a political solution. Suppose slumlords were fully compensated for conversion by \$284 per m^2, as if they had the right to utilize the land forever. That still leaves a remaining surplus of \$490 per m^2 at 3–4km. For the 29,070 households affected, the gain is about \$18,000 per household. This is a very large sum, for households paying on average less than \$700 a year in rent (from the NORC data and consistent with the rent calculations we use in Section 5.3.2). At 4–5km, a distance including much more of Kibera, the same type of calculation gives a surplus of about \$19,000 each for the 46,000 households.[36]

In these calculations, we have ignored moving costs for slum residents, broadly defined. For those forced to move there could be losses in terms of job location and social networks, although proper relocation programmes could mitigate those. Nevertheless, a key point of the calculations is that, relative to income and space-rents paid, there is an enormous surplus to play with to compensate residents. In one version of a "just" non-political world this would be solved by giving tenants the land titles and allowing them to sell themselves when and if they are ready.

These calculations are subject to bias. First, to be valid, slum lands near the centre must be the same quality as available formal sector lands. That is, it is slum history and formalization costs driving the current delay in development of slum lands, not geography. The basis of comparison is the vacant land for which we have sales price data. In Appendix

[35] GDP per capita in Kenya was \$1,090 in 2015. We set it at \$1,300 in Nairobi for a population of 4.2mn.

[36] These model-based numbers are substantially higher than preliminary back-of-the-envelope type calculations based on raw data reported in CEPR DP 11211 which give a gain of \$13,000 US per slum household for the core city.

Table 5.15, we perform a border experiment at slum boundaries comparing quality of slum lands with formal sector infill (vacant) lands that have just been built upon. We show that at distances out to 6km (where these old slums are), for elevation, ruggedness in the small, ruggedness in the large, presence of water, and being lower than mean neighbourhood elevation, of the twenty cases for slums compared to formal sector, seventeen show no differential, two are better, and one worse. It seems differentials in geography do not drive non-conversion of slums near the city centre.

5.7 Conclusions

This paper develops a framework for analysing a growing monocentric city in which residents make housing and location choices and developers sink capital in built structures, acting on the expectation of future rental income. We tailor the model to a developing country, so include an informal or slum sector in which structures are cheap and malleable, so capital is not sunk and construction can take place even if land-rights are insecure.We show how the model predicts the conversion of land from informal to formal as the city grows, and how areas go through successive waves of tear-down and redevelopment. The model provides a platform for further work, such as exploring the role of expectations in urban growth, endogenizing urban productivity, and moving from a monocentric to polycentric urban structure.

The application is to Nairobi, a fast-growing city that, we argue, has many features in common with other developing cities. Based on a unique data set constructed from high-resolution aerial photographs, we study the city in cross-section (2015) and change (2003–15). We characterize how building height and ground cover, in both the formal and informal sectors, vary with distance from the city centre and use this cross-section to calibrate key parameters of the model. We find—consistent with assumptions in our model—that, closer to the centre, both sectors offer more built volume per unit area; the formal sector by building tall and the informal by crowding (i.e. more cover per unit area).We derive estimates of parameters of the two building technologies, indicating that land-rent (as opposed to construction) cost, is a much higher share of total rental income in the informal sector than in the formal. Turning to changes through time, the model does well in predicting the pattern of slum growth and two waves of formal sector development that occurred in the period 2003–15. One was redevelopment of formal structures relatively close to the city centre, with redevelopment typically trebling building height as predicted by the model. The other was new formal sector development further out, on the fringe between formal and informal areas.

A feature of Nairobi and other developing cities is the persistence of slums in high land value areas near the city centre. This is rationalized in the model as being due to a package of barriers to formalization. In Nairobi, many of these barriers arise as "slumlords", who rent out property at market rates, do not have sound claims on the land, so resist a process of formalization that might deprive them of their flow of rental income. We use estimates from the model to calculate the losses due to these barriers, and find they are very substantial. Even if slumlords were to be paid-off (compensating for a loss of perpetual control), conversion would yield a surplus of about $18,000 per slum household in a context where they pay about $500–700 a year for their housing.

5.A Figures

Figure 5.1: Urban development with perfect foresight

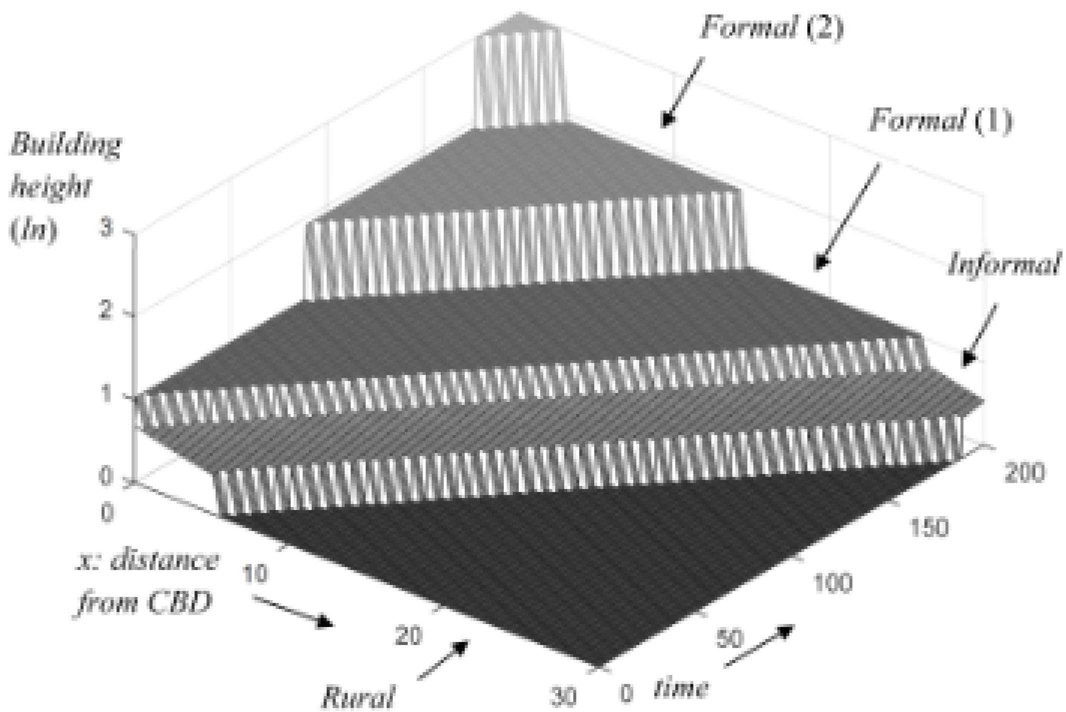

Figure 5.2: The hotchpotch—random variation in fixed costs of development, D(x)

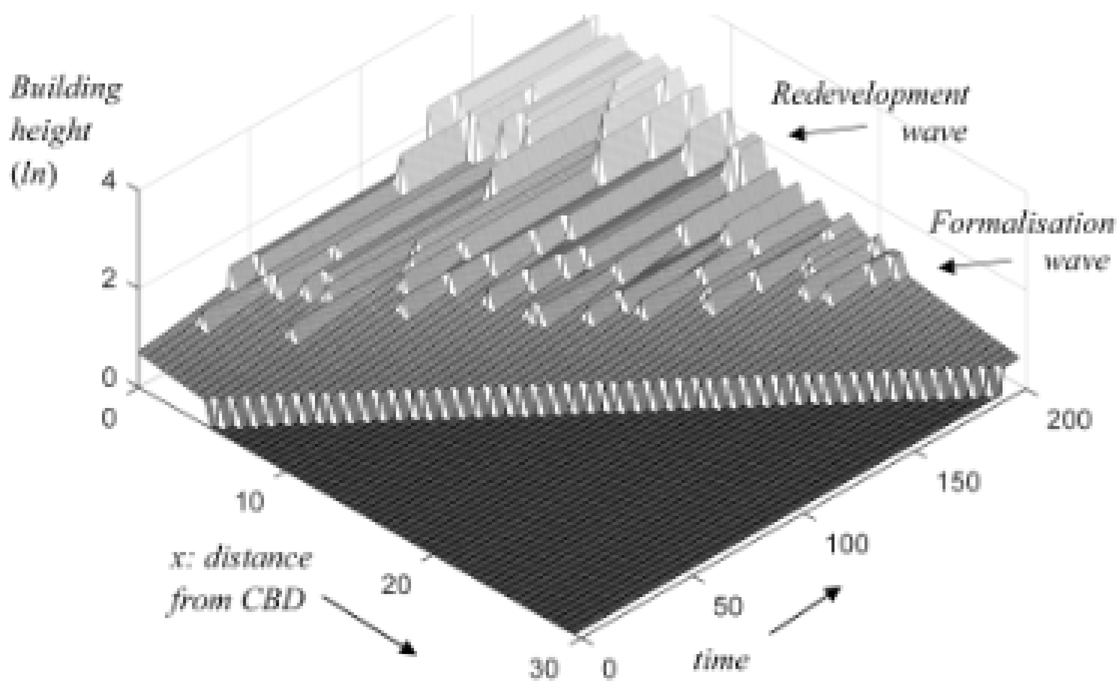

Figure 5.3: City shape and slums

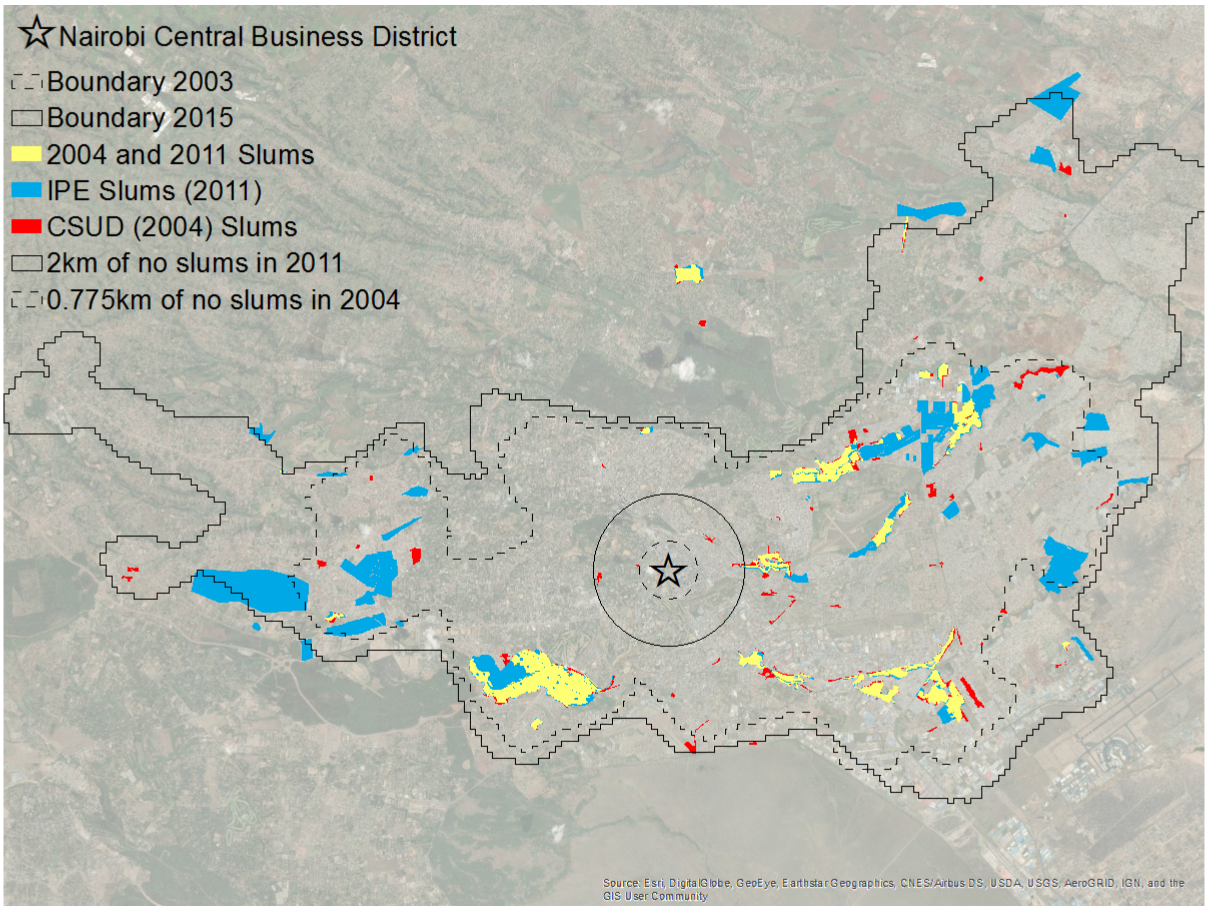

Notes: This figure maps two slum definitions; IPE from 2011 (blue), CSUD from 2003 (red), and the overlap of the two (yellow). The 2015 boundary of the city is traced with a solid black line, and the 2003 boundary is traced with a dashed black line. The central business district is denoted with a star and is circled for areas without slums in 2004 (dashed circle) and in 2011 (solid circle). Satellite imagery is faded in the background to give a sense of context.

Figure 5.4: Built volume per unit area (BVAR)

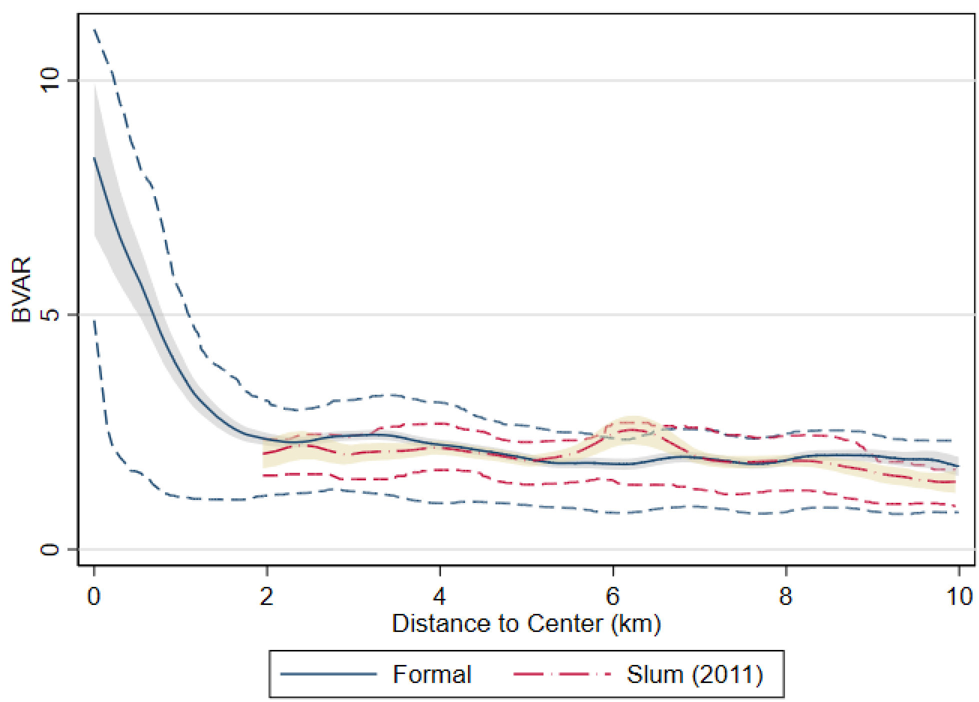

Notes: Built volume to area ratio (BVAR) by distance from the centre for formal and slum sectors. Lines show the local average BVAR, shaded areas show local 95% confidence intervals, and dashed lines show local 25th and 75th percentiles. Local statistics calculated using an Epanechnikov kernel with bandwidth of 300m.

Figure 5.5: Built cover per unit area (CAR) and height

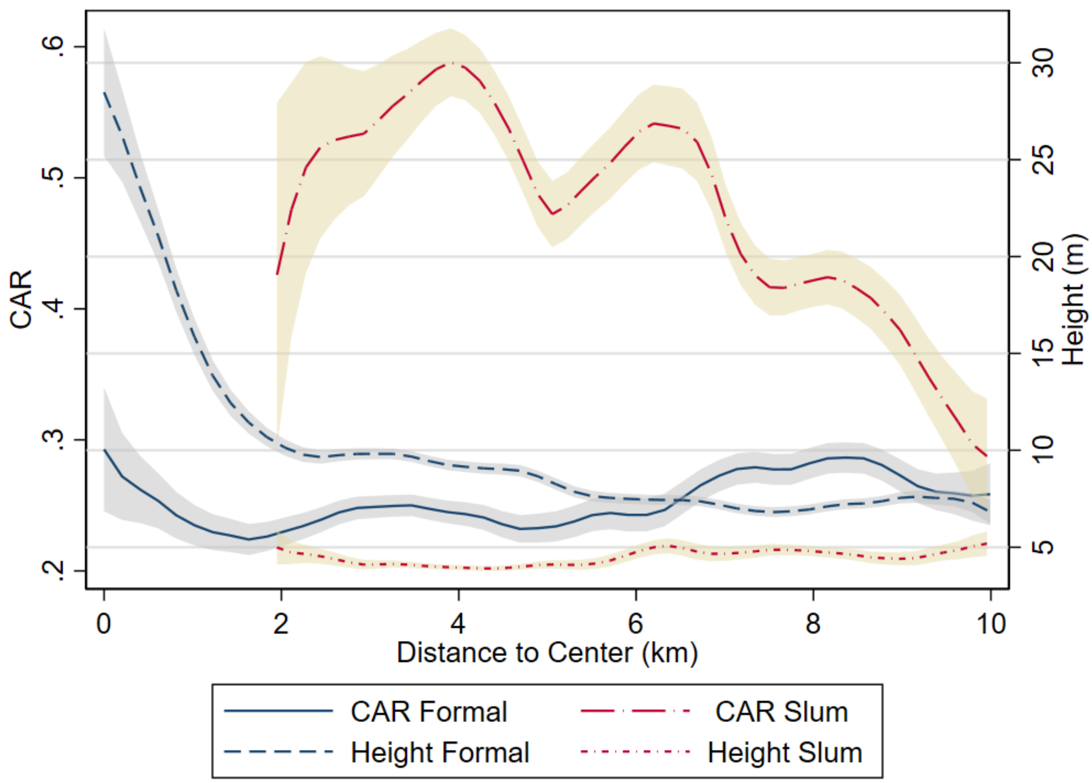

Notes: Cover to area ratio (CAR) and building height by distance from the centre for formal and slum sectors. Lines show the local average CAR and height, and shaded areas show local 95% confidence intervals. Local statistics calculated using an Epanechnikov kernel with bandwidth of 300m.

Figure 5.6: Formal sector changes in volume per unit area (2015 city to 12 km)

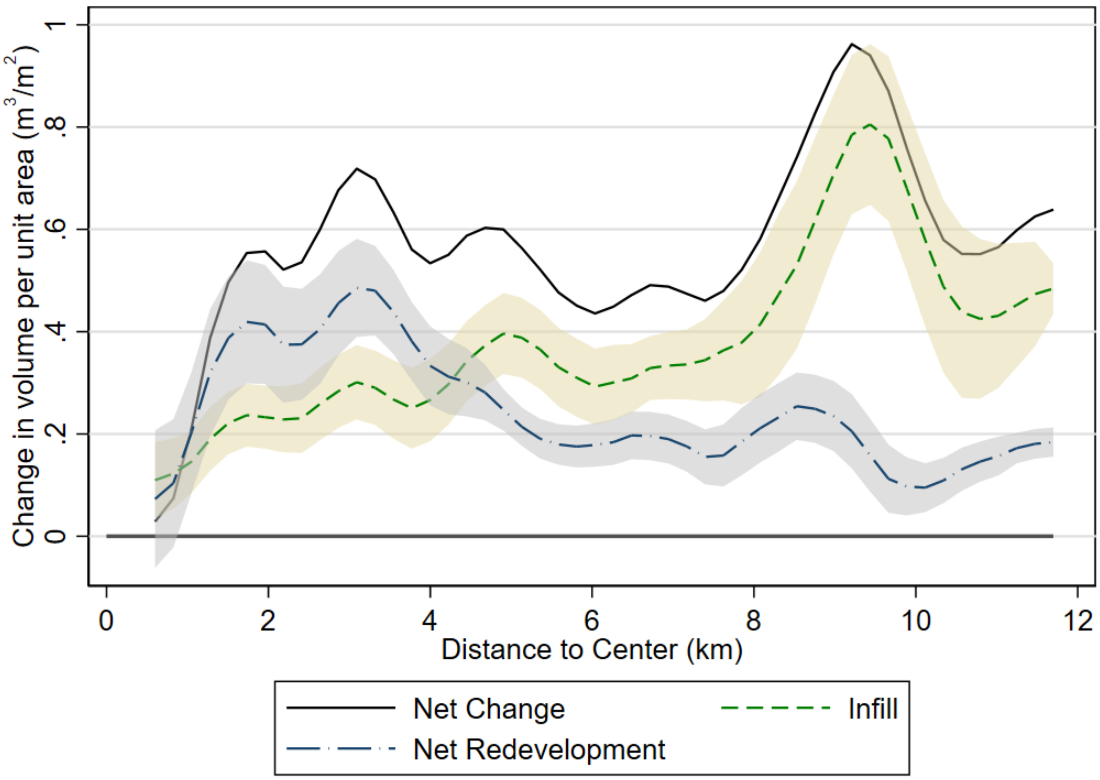

Notes: The change in formal sector volume, by redevelopment, infill and total net change, calculated in bins of 300m inside the 2015 city boundary and divided by the area of each bin. Shaded areas show local 95% confidence intervals for infill and redevelopment only. Local statistics smoothed using an Epanechnikov kernel with bandwidth of 300m.

Figure 5.7: Mean height—unchanged and redeveloped buildings

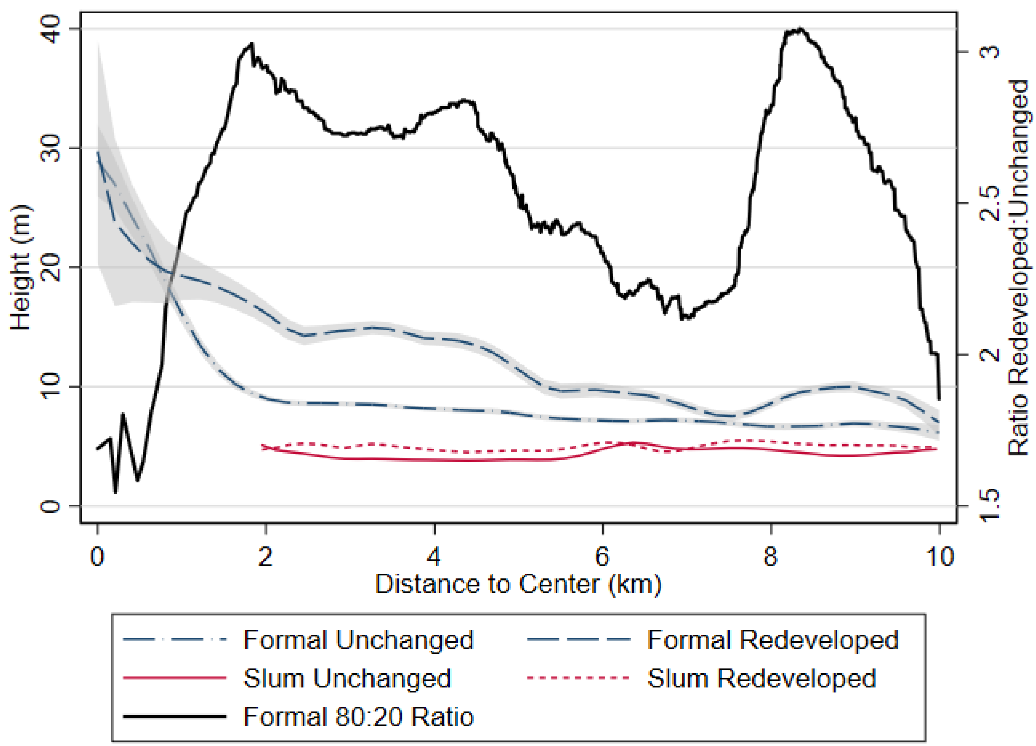

Notes: Height of unchanged and redeveloped buildings by distance from the centre for formal and slum sectors. Lines show local average height and shaded areas show local 95% confidence intervals for the formal sector only. The right axis is the ratio of the local 80th percentile of redeveloped height to the local 20th percentile of unchanged height. Local statistics calculated using an Epanechnikov kernel with bandwidth of 300m.

Figure 5.8: Slum sector changes in volume per unit area (2015 city out to 12 km)

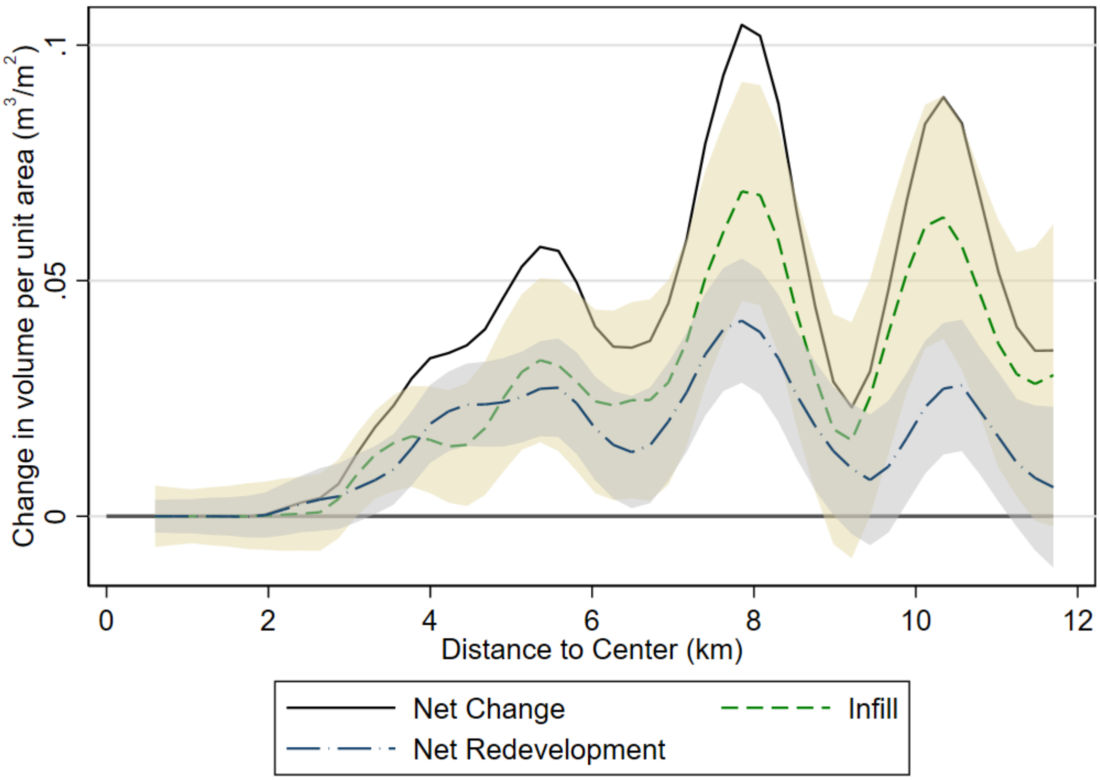

Notes: The change in slum sector volume, by redevelopment, infill and total net change, calculated in bins of 300m inside the 2015 city boundary and divided by the area of each bin. Shaded areas show local 95% confidence intervals for infill and redevelopment only. Local statistics smoothed using an Epanechnikov kernel with bandwidth of 300m.

Figure 5.9: F1 Scores under ICP Routines

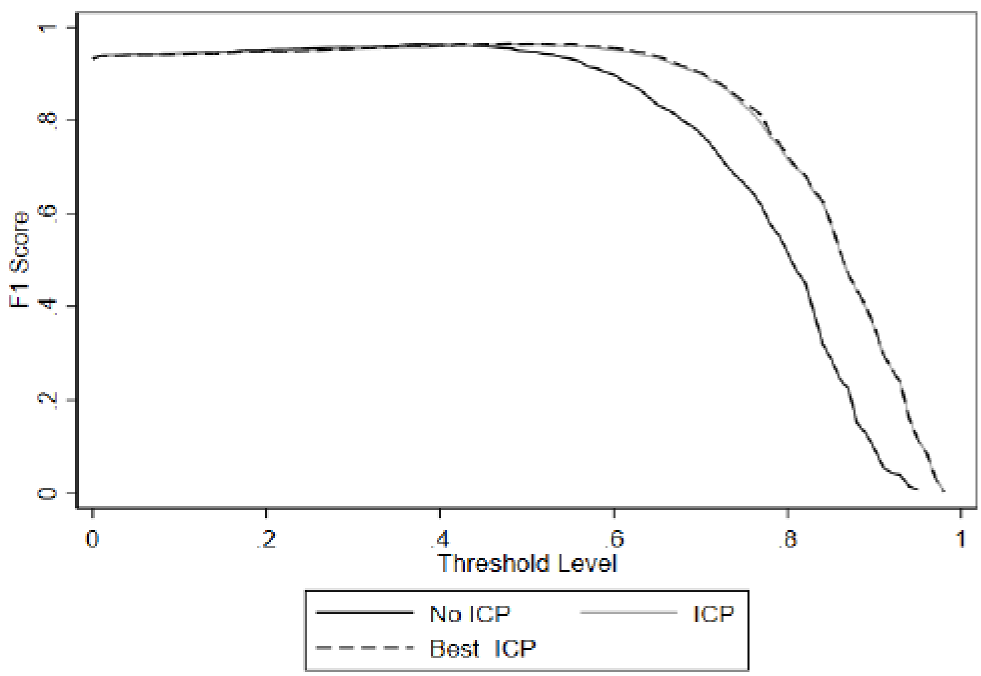

Notes: This figure shows average area-weighted F1 scores by threshold level of the overlay ratio, and by ICP translation routine. The solid dark line shows results using the overlay ratio determined without the ICP translation adjustment, the solid gray line uses the overlay ratio after ICP translation, and the dark dashed line uses the larger of the two overlay ratios (essentially in-line with that using ICP only).

Figure 5.10: Distribution of plot sizes: formal sector

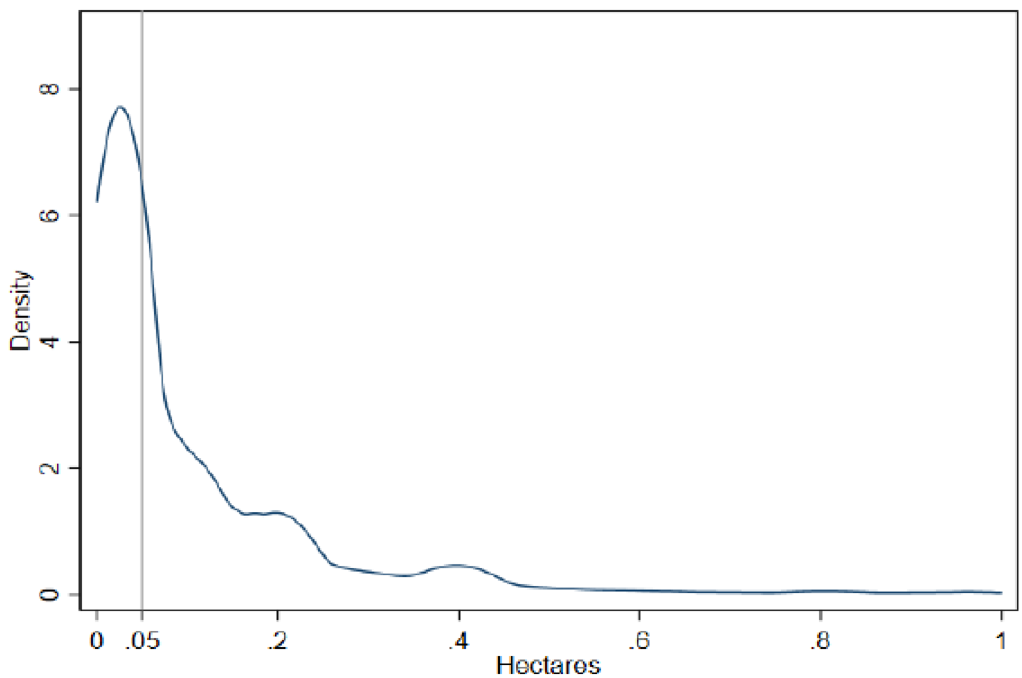

Notes: This figure shows the distribution of formal plot sizes from the cadastre. Plots above one hectare are truncated from the graph. The vertical line denotes 0.05 hectares. Of all recorded plots, 54% are smaller than 0.05 hectares, and modal density is at 0.018 hectares. Density is calculated using an Epanechnikov kernel with bandwidth of 0.025 hectares.

Figure 5.11: Census 2009 slums versus IPE slums

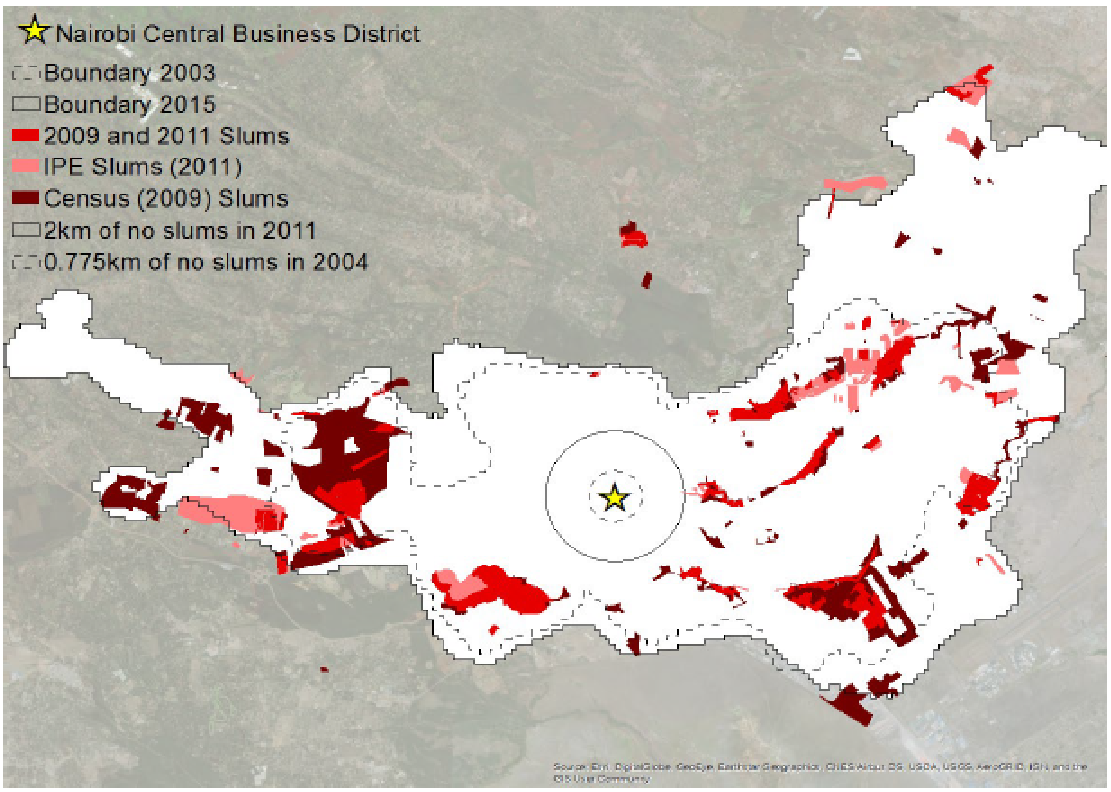

Notes: This figure maps two slum definitions; IPE from 2011 (light red), Census from 2009 (dark red), and the overlap of the two (red). The Kenyan National Cenusus Bureau defines all Enumeration areas as either planned or unplanned, even if the are unpopulated. Our 'Census' definition here maps all Enumeration areas that are unplanned and have at least 3,000 people per km^2.

Figure 5.12: Location of vacant land listings and of NORC surveying on house rents

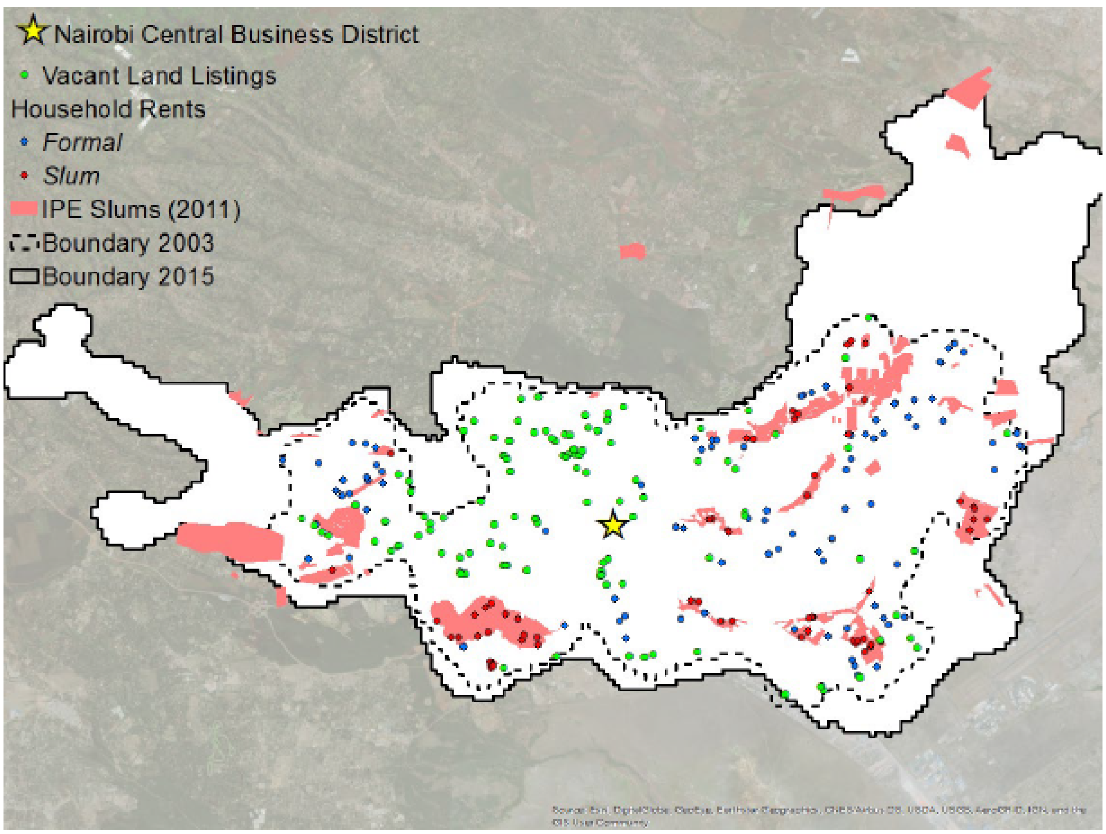

Notes: This figure maps formal (blue) and slum (red) household rents from the NORC survey and vacant land listings (green).

Figure 5.13: Teardown converted vs non-converted slums

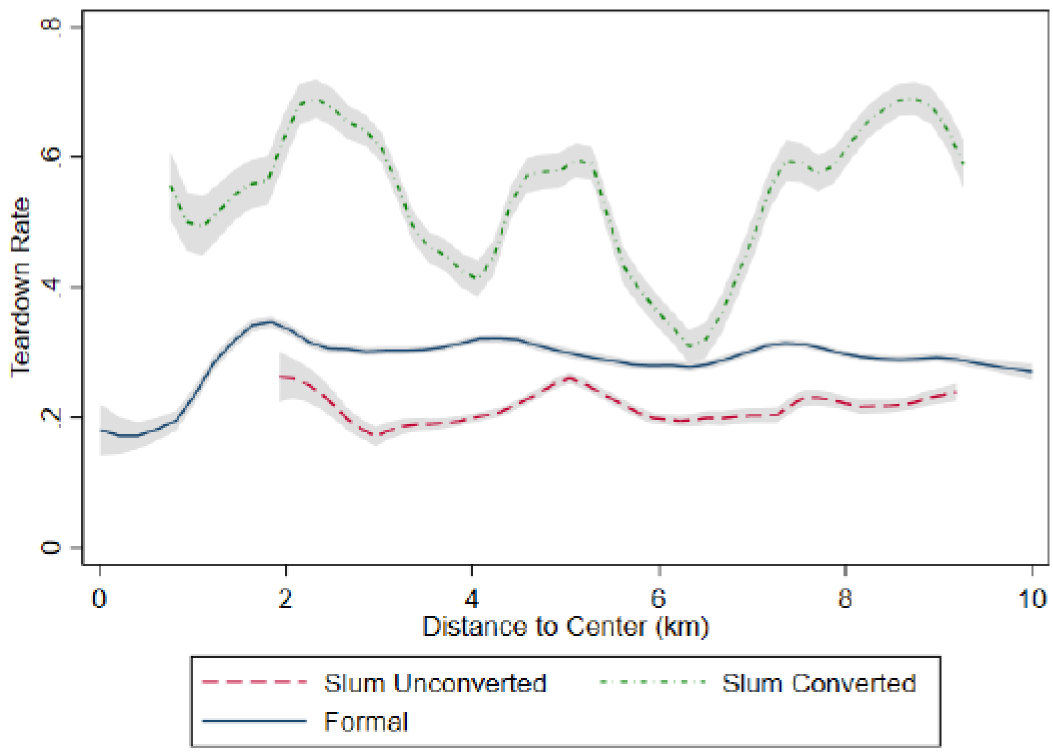

Notes: This figure shows teardown rates of 2003 buildings along distance to the city centre for the formal sector, and the 2003 slum sector. The 2003 slum sector is further broken down into unconverted slums (areas that remained classified as slum in 2011 by the IPE) and converted slums (areas that were not classified as slum in 2011 by the IPE). Teardowns are buildings that were either redeveloped or left unreplaced. The lines show the local average teardown rate, and the shaded areas show local 95% confidence intervals for this estimate. Local statistics are calculated using an Epanechnikov kernel with bandwidth of 300m.

Figure 5.14: Heights in converted vs non-converted slums

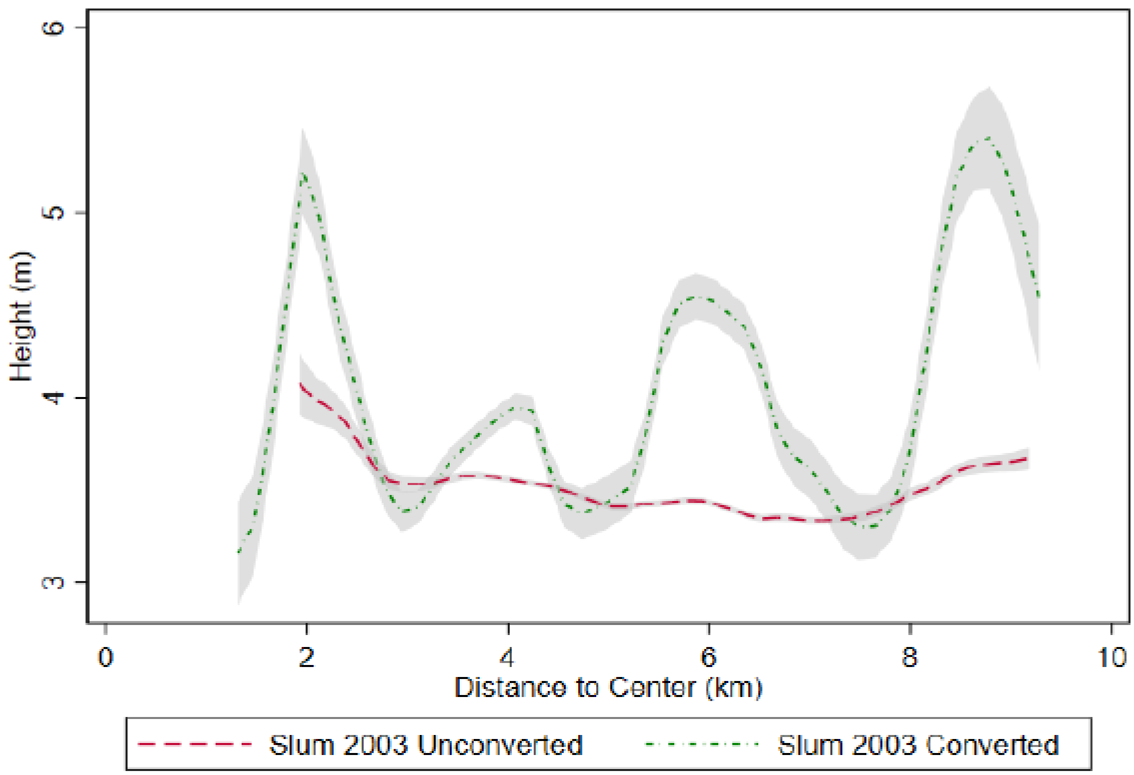

Notes: This figure shows heights of 2015 buildings along distance to the city centre for the 2003 slum sector. The 2003 slum sector is further broken down into unconverted slums (areas that remained classified as slum in 2011 by the IPE) and converted slums (areas that were not classified as slum in 2011 by the IPE). The lines show the local average building height, and the shaded areas show local 95% confidence intervals for this estimate. Local statistics are calculated using an Epanechnikov kernel with bandwidth of 300m.

Figure 5.15: Ownership of slums

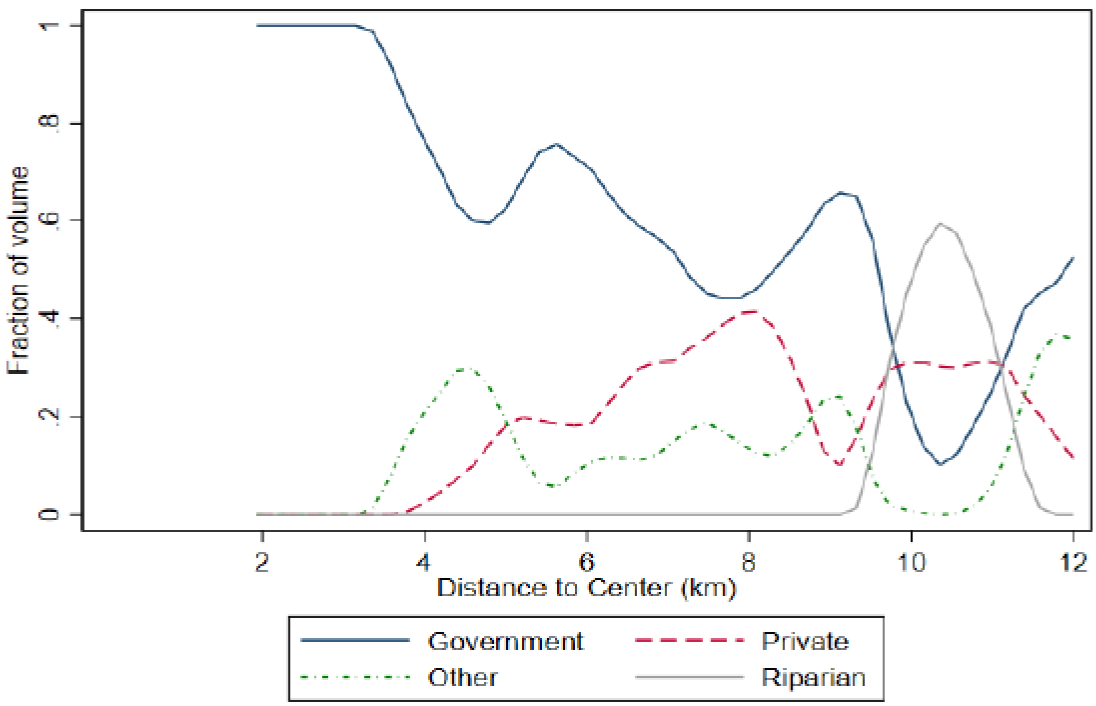

Notes: This figure shows the fraction of 2015 building volume along distance to the city centre for slum land by land tenure categories. Land tenure of slums is based on IPE (2013) and author calculations on volume. Government refers to slum land where tenure rights are held by a government entity, excluding the National Housing Corporation public housing. Private refers to slum land held exclusively by private individuals or entities. Other includes land that has a mixture of private and government ownership. Riparian refers to slums along rivers. From Kenya (2018) it seems that ownership of riparian slums is a mixture of mostly private property and possessory rights, with some government owned tracts. Local shares are calculated using an Epanechnikov kernel with bandwidth of 300m.

5.B Tables

Table 5.1: Formal sector

	Ln formal BVAR	Ln formal CAR	Ln formal height	Ln formal redeveloped height quantile: 80th pctile	Ln formal unchanged height quantile: 20th pctile	Ln land price m^2, \$2015	Ln formal space-rent: m^3, \$2015
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Distance to centre	-0.0493	0.0239	-0.0763	-0.101	-0.0763	-0.172	-0.0855
	(0.0124)	(0.00888)	(0.00791)	(0.00521)	(0.00223)	(0.0476)	(0.031)
Intercept (x = 0) for typical item	0.747	-1.672	2.487	3.315	2.131	7.254	3.148
	(0.0829)	(0.0571)	(0.0537)	(0.0356)	(0.015)	(0.277)	(0.071)
Ruggedness & elevation	✓	✓	✓	✓	✓	✓	✓
Other Controls						✓	✓
Observations	5435	5435	5435	4621	5079	136	361
R-squared	0.035	0.043	0.203			0.292	0.244

Notes: All columns except 7 are based on 2015 data for observations inside the 2003 extent of the city. Column 7 is based on NORC data and restricted to observations inside the 2003 extent of the city. Standard errors in parentheses. Standard errors are robust and in columns 1–3 are clustered based on a 750m×750m grid. Errors in cols. 4 and 5 are not clustered; in column 6 they are clustered at the neighbourhood area from the online listing service; and in column 7 at the census enumeration area. Details and full results for columns 4, 6, and 7 are in Table 5.10. Reported intercepts are predicted values at the city centre for median ruggedness and elevation for cols. 1–5. In those columns, for height regressions we only include grids for which there is cover. For BVAR and CAR regressions about 5% of observations have no cover and hence volume (e.g. playing fields, overpasses, small parks and the like). A Tobit including these as "censored" at 0 yields almost identical slope coefficients. In column 6, the reported intercept is the predicted sales price at the centre per m^2 for a lot based on median area, ruggedness, and elevation for the sample. In column 7, the intercept is the rent per cubic meter of space for a house with typical characteristics in the formal sector: mean values for categorical variables and median for continuous for the NORC formal sector, for the regressions where we force the gradient to have the model slope of 0.071 solved in the next section.

Table 5.2: Slum sector

	Ln slum BVAR	Ln slum CAR	Ln slum height	Ln slum space-rent: m^3, \$2015
	(1)	(2)	(3)	(4)
Distance to centre	-0.0948	-0.103	0.0107	-0.0094
	(0.0172)	(0.0171)	(0.00871)	(0.0243)
Intercept (x = 0) for typical item	1.275	-0.116	1.375	1.886
	(0.1176)	(0.1049)	(0.0686)	(0.0581)
Ruggedness & elevation	✓	✓	✓	✓
Other Controls				✓
Observations	958	958	958	439
R-squared	0.104	0.142	0.017	0.391

Notes: All columns are based on 2015 data for observations inside the 2003 extent of the city. Robust standard errors are in parentheses. Errors in columns 1–3 are clustered based on a 750m×750m grid and in column 4 are clustered at the enumeration area. Details and full results are in Table 5.10. Reported intercepts are predicted values at the city centre for median ruggedness and elevation for cols. 1–3. In column 4, the intercept is the rent per m^3 of space for a house with typical characteristics in the slum sector: mean values for categorical variables and median for continuous in the NORC slum sector for a regression where we force the gradient to have the model slope of 0.

Table 5.3: Model parameters

	Formal	Slum
a. Model parameters from gradients		
Volume m^3, per m^2 land	$\frac{dv_F(x,\tau_i)}{dx}\frac{1}{v_F} = \frac{-\theta}{\gamma-1} = -0.101$ Eqn (5.25): Table 5.1 col. 4	$\frac{dv_I(x,t)}{dx}\frac{1}{v_I} = \frac{-\theta\alpha}{\alpha-1} = -0.0948$ Eqn (5.4); Table 5.2 col. 1
Land-price, per m^2 land	$\frac{dR_F(x,\tau_i)}{dx}\frac{1}{R_F} = \frac{-\theta\gamma}{\gamma-1} = -0.172$ Eqn (5.25); Table 5.1 col. 6	
Solutions: $\gamma = 1.703$, $\alpha = 3.983$, $\theta = 0.071$		
b. Levels of variables and further parameters		
Volume m^3, per m^2 land	$v_F(x,\tau_i) = \left[\frac{p(x,\tau_i)}{\gamma}\frac{\Phi}{\kappa_F}\right]^{1/(\gamma-1)} = 6.80$ Eqn (5.10); Table 1 col. 4 for height, times median formal CAR	$v_I(x,t) = \left[\frac{a_I p(x,t)}{\kappa_I \alpha}\right]^{\alpha/(\alpha-1)} = 3.58$ Eqn (5.4): Table 5.2 col. 1
Space-rent, per m^3 volume	$p(x,t) = 23.29$ Table 5.1 col. 7	$p(x,t)a(v_I(x,t)) = \kappa_I\alpha = 6.59$ Eqn (5.6): Table 5.2 col. 4
Solutions: $p(x,t) = 23.29$, $a_I = 0.734$, $\kappa_F/\Phi = 3.554$		
Levels evaluated at $x = 0$, $t = \tau_i = 2015$, in US\$		

Table 5.4: Dates of development

	Distance from centre, x						
	0.5km	1.5km	2.5km	3.5km	10.5km	12.5km	14.5km
τ_1: Formalisation, $D(x) = 0$	1917	1924	1932	1940	1994	2009	2025
τ_2: First redevelopment, $D(x) = 0$	2006	2014	2022	2029	2083	2099	2114
Model prediction with space-rent path:							
$p(x,t) = p(0,2015)exp(\hat{p}(t-2015) - \theta x) = 23.29exp(0.0092(t-2015) - 0.071x)$							

Table 5.5: The value of land formalized at different dates

	Distance from centre, x						
	0-1km	1-2km	2-3km	3-4km	4-5km	5-6km	6-7km
$PV(x, z = 2015)$	1297	1092	919	774	652	549	462
$PV(x, z = \infty)$	377	343	312	284	258	235	213
Slum land, km^2, 2011	0	0.0024	0.24	1.07	2.22	1.9	1.32
No. slum households, 2009	0	0	2920	29,070	45,810	33,100	28,390
Lower bound on D (5.18)	493	395	314	248	194	150	113
Present values in 2015 in \$2015 per m^2							
Date of formalisation, z							

Table 5.6: Mapping Correspondence 2003

a. Weighted by building						
	Algo=-1	Algo=0	Algo=1	Algo=2	Algo=3	Algo=4
Manual=0	280	413	41	24	30	10
Manual=1	0	12	729	20	3	3
Manual=2	0	10	13	315	0	6
Manual=3	0	13	3	0	145	1
Manual=4	0	82	69	57	47	84
b. Weighted by area (m^2)						
Algo=-1	Algo=0	Algo=1	Algo=2	Algo=3	Algo=4	
Manual=0	12708	27194	3303	3295	1982	686
Manual=1	0	691	112262	6055	347	242
Manual=2	0	2343	1545	94617	0	515
Manual=3	0	347	202	0	15562	23
Manual=4	0	5308	5543	10978	4429	8950
Mapping definitions: -1 demolition or infill; 0 redevelopment; 1 one to one match; 2 one to many match; 3 many to one match; 4 many to many match.						

Table 5.7: Matching all areas 2003

a. Weighted by building			
Algo=0	Algo=1	Recall	
Manual=0	413	105	0.8
Manual=1	117	1495	0.93
Precision	0.78	0.93	F=0.93
b. Weighted by area (m^2)			
Algo=0	Algo=1	Recall	
Manual=0	27194	9266	0.75
Manual=1	8689	261270	0.97
Precision	0.76	0.97	F=0.97

Table 5.8: List of public uses

Recreational	**Public utility**
a) Impala club, Kenya Harlequins, and	a) Dandora dump ($0.5km^2$)
Rugby Union of East Africa ($0.14km^2$)	b) Sewage works ($0.25km^2$)
b) Golf Course ($0.9km^2$)	c) Kahawa Garrison
c) Arboretum ($0.25km^2$)	**Public use**
d) Central park, Uhuru park, railway club,	a) Communications Commission of Kenya ($0.1km^2$)
railway golf course ($0.5km^2$)	b) Langata Womens prison ($0.2km^2$)
e) Nyayo stadium ($0.1km^2$)	c) Nairobi and Kenyatta hospitals,
f) City park, Simba Union, Premier Club ($1.1km^2$)	Milimani Police Station, Civil Service club
g) Barclays, Stima, KCB, Ruaraka, Utali clubs,	d) Mbagathi hospital, Kenya Medical Research
and FOX drive in cinema ($0.3km^2$)	Institute, Monalisa funeral home
Undeveloped	e) National museums of Kenya
a) Makdara Railway Yard ($1km^2$)	f) Kenya convention centre and railway museum
b) John Michuki Memorial Park ($0.1km^2$)	g) Industrial area prison
Special use (includes poorly traced areas)	h) Mathari mental hospital, Mathare police station,
a) State House	traffic police, Kenya police, Ruaraka complex, and
b) Ministry of State for Defence	National youth service
c) Forces Memorial Hospital and	i) Jamahuri show ground
Administration Police Camp	**Educational (not primary and secondary schools)**
d) Langata Army Barracks	a) University of Nairobi and other colleges
e) Armed Forces	b) Kenya Institute of Highways & Built Technology
f) Moi Airbase	c) Railway Training Institute
	d) Kenya Veterinary Vaccines Production Institute
	e) Moi Forces Academy
	f) NYS engineering, Kenya Institute of Monetary Studies,
	KCA university, KPLC training, Utali college

Table 5.9: Monocentricy of population in developing and developed country cities

	Developing country cities (1)	Developed country cities (2)
Distance to centre (km)	-0.330 (0.018)	-0.138 (0.013)
Observations	2721	2473
R-squared	0.17	0.13

Notes: All observations are 1km×1km grid cells. The dependent variable is always log population as measured by Landscan. The developing country cities (col 1) are Luanda, Dar es Salaam, Abidjan, Nairobi, Kano, Dakar, Addis Abba, Ibadan, Yaound, Douala. The developed (col 2) are Atlanta, DC, Phoenix/Mesa, Montreal, Rome, Detroit, San Francisco, San Diego, Athens, Lisbon. City fixed effects are included in all columns. Significance levels for the null hypothesis that the true parameter is zero are denoted with asterisks: $^{*}p < 0.10$;$^{**}p < 0.05$;$^{***}p < 0.01$.

Table 5.10: Full key regressions in Tables 1 and 2

	Ln land sales price (USD per m^2)	Ln Formal redeveloped height; quantile: 80th pctile	Ln Slum BVAR		Ln formal space-rent per m^3 vol. in \$2015	Ln slum space-rent per m^3 vol. in \$2015
	(1)	(2)	(3)		(4)	(5)
Distance to centre	-0.172	-0.101	-0.0948	Distance to centre	-0.0855	-0.00943
	(0.0476)	(0.00521)	(0.0172)		(0.0310)	(0.0243)
SD elevation (km)	-0.0114	0.0222	-0.0138	Ln SD elevation	0.00226	-0.138
	(0.0531)	(0.0116)	(0.0136)		(0.144)	(0.120)
Elevation (km)	0.00535	321.7	-1142.5	Ln elevation	2.145	6.294
	(0.00178)	(197.9)	(404.3)		(2.122)	(1.204)
Lot size	-0.0302			No written tenancy	-0.428	-0.320
	(0.103)				(0.133)	(0.209)
Lot size^2	-0.00108			No piped water	-0.335	-0.278
	(0.00159)				(0.152)	(0.109)
Coordinates estimated	-0.468			Ln# Floors	0.287	0.126
	(0.372)				(0.108)	(0.126)
				# Bathrooms		
				One	-0.162	-0.0714
					(0.144)	(0.125)
				Two+	0.0098	0.296
					(0.220)	(0.132)
				Structure type		
				Single-story Shared facil.	0.0893	0.260
					(0.175)	(0.116)
				Multi-storey	0.165	0.347
					(0.200)	(0.133)
				Room in house		-0.603
						(0.169)
				Type of walls		
				Some mud	0.404	-0.839
					(0.215)	(0.211)
				Wood only	-0.370	-0.228
					(0.225)	(0.234)
				Iron or tin	0.400	-0.375
					(0.196)	(0.192)
Constant	-1.328	2.744	3.163	Constant	-12.46	-43.72
	(3.234)	(0.329)	(0.700)		(15.68)	(8.862)
Month F.E.s	✓					
Observations	136	4621	958	Observations	361	439
R^2	0.292		0.104	R^2	0.244	0.391

Notes: Columns 1-3 are based on 2015 data for observations inside the 2003 extent of the city. Columns 4 and 5 are based on NORC data and restricted to observations inside the 2003 extent of the city. Standard errors in parentheses. Errors in col. 1 are clustered at the neighbourhood area from the on-line listing service; in col. 2 are not clustered; in col. 3 are clustered based on a 750m×750m grid; and in cols. 4 and 5 at the census enumeration area. In column 2, for height regressions, we only include grids for which there is cover. For the BVAR regression, column 3, about 5% of observations have no cover and hence volume (e.g. playing fields, overpasses, small parks and the like). A Tobit including these as 'censored' at 0 yields almost identical slope coefficients. Monthly fixed effects in col. 1 refer to the month of first listing. In predicting the intercept in Table 1 we assign the last 3 months a weight of one-third each and the rest 0. Significance levels for the null hypothesis that the true parameter is zero are denoted with asterisks: $^{*}p < 0.10;^{**} p < 0.05;^{***} p < 0.01$.

Table 5.11: Robustness

	Standard	Adding distance to industrial	New CBD Definition	Census slum definition
	(1)	(2)	(3)	(4)
(1) Ln land price (\$2015 per m^2)				
Distance to centre	-0.172	-0.187	-0.17	n/a
	(0.0476)	(0.0534)	(0.0474)	
Distance to Industrial Centre		0.046		n/a
		(0.0974)		
'Intercept' for typical item	7.254	7.308	7.242	n/a
	(0.277)	(0.279)	(0.275)	
Observations	136	136	136	n/a
R^2	0.183	0.177	0.183	
(2) Ln formal redeveloped height; quantile: 80th pctile				
Distance to centre	-0.101	-0.107	-0.101	-0.09
	(0.0052)	(0.0095)	(0.0047)	(0.0049)
Distance to Industrial Centre		0.007		
		(0.0107)		
'Intercept' for typical item	3.315	3.368	3.298	3.225
	(0.036)	(0.072)	(0.032)	(0.033)
Observations	4621	4621	4621	4133
(3) Ln slum BVAR				
Distance to centre	-0.009	0.001	-0.014	
	(0.0172)	(0.0242)	(0.0181)	(0.0205)
Distance to Industrial Centre		-0.016		
		(0.0203)		
'Intercept' for typical item	1.275	1.162	1.286	1.12
	(0.118)	(0.199)	(0.120)	(0.152)
Observations	958	958	958	1590
R^2	0.101	0.101	0.105	0.087
(4) Ln formal space-rent m^3 in \$2015				
Distance to centre	-0.086	-0.048	-0.086	n/a
	(0.0310)	(0.0571)	(0.0307)	
Distance to Industrial Centre		-0.051		n/a
		(0.0750)		
'Intercept' for typical item	3.148	3.142	3.128	n/a
	(0.071)	(0.072)	(0.071)	
Observations	361	361	361	n/a
R^2	0.216	0.216	0.217	
(5) Ln slum space-rent m^3 in \$2015				
Distance to centre	-0.009	0.001	-0.014	n/a
	(0.0243)	(0.0416)	(0.0245)	
Distance to Industrial Centre		-0.013		n/a
		(0.0450)		
'Intercept' for typical item	1.886	1.859	1.886	n/a
	(0.058)	(0.087)	(0.058)	
Observations	439	439	439	n/a
R^2	0.371	0.37	0.371	

Notes: This table runs four robustness checks where each panel corresponds to a column in Table A2.5. Column 1 repeats the regressions from Table 5.10, column 2 adds a control for the distance to the industrial centre in kilometers, column 3 uses an alternative definition of the city centre based on recent nightlight values rather than 1992 values, and column 4 uses an alternative definition of slums based on the 2009 Kenyan Census unplanned enumeration areas. Panels 1-3 are based on 2015 data for observations inside the 2003 extent of the city. Panels 4 and 5 are based on NORC data and restricted to observations inside the 2003 extent of the city. Standard errors in parentheses. Errors in panel 1 are clustered at the neighbourhood area from the on-line listing service; in panel 2 are not clustered; in panel 3 are clustered based on a 750m×750m grid; and in panels 4 and 5 at the census enumeration area. In panel 2, for height regressions, we only include grids for which there is cover.

Table 5.12: Robustness

	Standard	Adding distance to industrial	New CBD Definition	Census slum definition
	(1)	(2)	(3)	(4)
(1) Ln land price (\$2015 per m^2)				
Distance to centre	-0.172	-0.187	-0.17	n/a
	(0.0476)	(0.0534)	(0.0474)	
Distance to Industrial Centre		0.046		n/a
		(0.0974)		
‘Intercept’ for typical item	7.254	7.308	7.242	n/a
	(0.277)	(0.279)	(0.275)	
Observations	136	136	136	n/a
R^2	0.183	0.177	0.183	
(2) Ln formal redeveloped height; quantile: 80th pctile				
Distance to centre	-0.101	-0.107	-0.101	-0.09
	(0.0052)	(0.0095)	(0.0047)	(0.0049)
Distance to Industrial Centre		0.007		
		(0.0107)		
‘Intercept’ for typical item	3.315	3.368	3.298	3.225
	(0.036)	(0.072)	(0.032)	(0.033)
Observations	4621	4621	4621	4133
(3) Ln slum BVAR				
Distance to centre	-0.009	0.001	-0.014	
	(0.0172)	(0.0242)	(0.0181)	(0.0205)
Distance to Industrial Centre		-0.016		
		(0.0203)		
‘Intercept’ for typical item	1.275	1.162	1.286	1.12
	(0.118)	(0.199)	(0.120)	(0.152)
Observations	958	958	958	1590
R^2	0.101	0.101	0.105	0.087
(4) Ln formal space-rent m^3 in \$2015				
Distance to centre	-0.086	-0.048	-0.086	n/a
	(0.0310)	(0.0571)	(0.0307)	
Distance to Industrial Centre		-0.051		n/a
		(0.0750)		
‘Intercept’ for typical item	3.148	3.142	3.128	n/a
	(0.071)	(0.072)	(0.071)	
Observations	361	361	361	n/a
R^2	0.216	0.216	0.217	
(5) Ln slum space-rent m^3 in \$2015				
Distance to centre	-0.009	0.001	-0.014	n/a
	(0.0243)	(0.0416)	(0.0245)	
Distance to Industrial Centre		-0.013		n/a
		(0.0450)		
‘Intercept’ for typical item	1.886	1.859	1.886	n/a
	(0.058)	(0.087)	(0.058)	
Observations	439	439	439	n/a
R^2	0.371	0.37	0.371	

Notes: This table runs four robustness checks where each panel corresponds to a column in Table A2.5. Column 1 repeats the regressions from Table 5.10, column 2 adds a control for the distance to the industrial centre in kilometers, column 3 uses an alternative definition of the city centre based on recent nightlight values rather than 1992 values, and column 4 uses an alternative definition of slums based on the 2009 Kenyan Census unplanned enumeration areas. Panels 1-3 are based on 2015 data for observations inside the 2003 extent of the city. Panels 4 and 5 are based on NORC data and restricted to observations inside the 2003 extent of the city. Standard errors in parentheses. Errors in panel 1 are clustered at the neighbourhood area from the on-line listing service; in panel 2 are not clustered; in panel 3 are clustered based on a 750m×750m grid; and in panels 4 and 5 at the census enumeration area. In panel 2, for height regressions, we only include grids for which there is cover.

Table 5.13: Quality of infilled land

	Elevation (m)	Water indicator 90x90m	SD elevation in 90x90m	SD elevation in 450x450m	Elevation difference to mean of 450x450m
	(1)	(2)	(3)	(4)	(5)
Distance to centre	-0.764	0.000449	-0.0131	-0.0239	0.0141
	(1.169)	(0.00162)	(0.0137)	(0.0452)	(0.00470)
Infilled=1	35.24	0.0971	0.421	1.392	-0.827
	(5.134)	(0.0217)	(0.106)	(0.280)	(0.103)
Distance to Center (km) × Infilled=1	-7.552	-0.00678	-0.0386	-0.207	0.0224
	(0.220)	(0.00112)	(0.00576)	(0.0107)	(0.0141)
Constant	1687.7	0.0566	1.717	4.721	0.0303
	(4.734)	(0.00962)	(0.0728)	(0.233)	(0.0261)
Observations	120299	120299	120299	120299	120299
R^2	0.015	0.005	0.002	0.004	0.003

Notes: All columns are based on 30m×30m grid observations on formal land inside the 2003 boundary and within 10km of the city centre. Standard errors in parentheses are robust and clustered based on a 450m×450m grid. The dependent variables are; the absolute elevation in meters above sea level (col 1); an indicator of the presence of rivers, streams, or lakes inside the surrounding 90m×90m block (col 2); the standard deviation in elevation across cells in the surrounding 90m×90m block (col 3); the standard deviation in elevation across cells in the surrounding 450m×450m block (col 4); and the relative elevation compared to the mean of cells in the surrounding 450m×450m block in meters (col 5). 'Infilled' indicates cells that are infill in 2015 and greenfield in 2003, i.e. the building did not previously exist and has not replaced another building. Significance levels for the null hypothesis that the true parameter is zero are denoted with asterisks: $^{*}p < 0.10;^{**}p < 0.05;^{***}p < 0.01$.

Table 5.14: Redeveloped height and land quality

	Ln formal redeveloped height
Distance to centre	-0.0660
	(0.00796)
Ln elevation (m)	-0.000137
	(0.000265)
Water indicator 90x90m	-0.00499
	(0.0547)
S.D. elevation in 90x90m	0.0266
	(0.00682)
S.D. elevation in 450x450m	0.00862
	(0.00697)
Elevation difference to mean of 450x450m	0.00410
	(0.00288)
Constant	2.331
	(0.439)
Observations	33701
R^2	0.053

Notes: This table based on 30m×30m grid observations on formal redeveloped land inside the 2003 boundary and within 10km of the city centre. Standard errors in parentheses are robust and clustered based on a 450m×450m grid. The dependent variable is the log of redeveloped building height. The rows are; distance to the city centre in kilometres; the absolute elevation in meters above sea level; an indicator of the presence of rivers, streams, or lakes inside the surrounding 90m×90m block; the standard deviation in elevation across cells in the surrounding 90m×90m block; the standard deviation in elevation across cells in the surrounding 450m×450m block; and the relative elevation compared to the mean of cells in the surrounding 450m×450m block in meters. Significance levels for the null hypothesis that the true parameter is zero are denoted with asterisks: $^{*}p < 0.10;^{**}p < 0.05;^{***}p < 0.01$.

Table 5.15: Comparison of slum land on one side of border to formal sector infill land on the other

	Elevation (m)	Water indicator 90x90m	SD elevation in 90x90m	SD elevation in 450x450m	Elevation difference to mean of 450x450m
	(1)	(2)	(3)	(4)	(5)
Distance to centre=3-4	0.806	-0.0866	0.468	0.397	0.145
	(0.474)	(0.0933)	(0.169)	(0.0899)	(0.141)
Distance to centre=4-5	0.487	-0.420	0.439	-0.117	0.667
	(1.919)	(0.207)	(0.626)	(0.701)	(0.428)
Distance to centre=5-6	1.978	-0.249	0.458	0.504	-0.102
	(2.902)	(0.231)	(0.725)	(0.799)	(0.550)
Distance to centre=6-7	0.236	0.201	0.935	0.999	-0.0776
	(3.393)	(0.300)	(0.683)	(0.723)	(0.567)
Distance to centre=7-8	-2.433	0.0997	1.645	2.403	0.533
	(3.297)	(0.298)	(0.729)	(0.837)	(0.588)
Distance to centre=8-9	-1.291	0.0558	2.308	3.038	0.306
	(3.748)	(0.314)	(0.937)	(1.036)	(0.683)
Distance to centre=9-10	-2.200	0.0562	1.825	2.298	0.0277
	(3.747)	(0.310)	(0.844)	(1.025)	(0.885)
Distance to centre=3-4 × Slum in 2015=1	-0.799	0.0844	0.0265	-0.125	0.0183
	(1.237)	(0.0828)	(0.301)	(0.296)	(0.280)
Distance to centre=4-5 × Slum in 2015=1	-0.468	0.291	-0.144	0.321	-0.157
	(0.603)	(0.0819)	(0.449)	(0.554)	(0.121)
Distance to centre=5-6 × Slum in 2015=1	-0.561	0.0597	-0.0597	-0.183	0.315
	(1.434)	(0.0738)	(0.320)	(0.426)	(0.335)
Distance to centre=6-7 × Slum in 2015=1	0.211	-0.194	0.326	0.408	0.510
	(1.120)	(0.0666)	(0.132)	(0.153)	(0.201)
Distance to centre=7-8 × Slum in 2015=1	0.525	-0.0246	-0.103	-0.337	-0.0156
	(0.561)	(0.0720)	(0.207)	(0.405)	(0.255)
Distance to centre=8-9 × Slum in 2015=1	-1.215	0.107	-0.696	-0.903	0.330
	(1.688)	(0.0729)	(0.589)	(0.593)	(0.274)
Distance to centre=9-10 × Slum in 2015=1	-0.317	-0.0542	-0.157	-0.136	0.638
	(1.604)	(0.0452)	(0.286)	(0.369)	(0.555)
Constant	1652.3	0.288	1.017	1.879	-0.509
	(2.362)	(0.211)	(0.537)	(0.569)	(0.430)
Observations	10167	10167	10167	10167	10167
R^2	0.995	0.416	0.369	0.584	0.081

Notes: All columns are based on 30x30m grid observations on land inside the 2003 boundary and within 10km of the city centre. Further, observations are restricted to be within 300m of a government owned slum boundary, and for formal land only include cells that are infill in 2015 and greenfield in 2003. Standard errors in parentheses are robust and clustered based on a 450m×450m grid. The dependent variables are; the absolute elevation in meters above sea level (col 1); an indicator of the presence of rivers, streams, or lakes inside the surrounding 90m×90m block (col 2); the standard deviation in elevation across cells in the surrounding 90m×90m block (col 3); the standard deviation in elevation across cells in the surrounding 150m×150m block (col 4); and the relative elevation compared to the mean of cells in the surrounding 150m×150m block in meters (col 5). Rows show estimates of the mean for each kilometre bin from the city centre on both the formal and slum sides of the boundary. Significance levels for the null hypothesis that the true parameter is zero are denoted with asterisks: $^{*}p < 0.10;^{**}p < 0.05;^{***}p < 0.01$.

5.C Theory

5.C.1 Further details of the model

Details of the open city equilibrium model underpinning the model of the text are as follows:

Households:

At date t a representative urban household living at distance x from the centre receives net income $w(t,x)$. The indirect utility function is $U^*(p(x,t),w(x,t))$ and the space-rent $p(x,t)$ must be such that at all occupied places $U^*(p(x,t),w(x,t)) = u_0$. In Sections 5.2.1-5.2.3 we assume that housing is normal, $w(x,t)$ is increasing through time, and u_0 is constant. Household demand for quantity adjusted housing volume is, by Roy's identity, $s(x,t)a = -\frac{\partial U^*}{\partial p(x,t)} / \frac{\partial U^*}{\partial w(x,t)} = \frac{\beta w(x,t)}{p(x,t)}$. The last of these equations assumed preferences are Cobb-Douglas, i.e. $U^*(p(x,t),w(x,t)) = p(x,t)^{-\beta}w(x,t)$, with corresponding direct utility function $U(s(x,t)a, c(x,t)) = (s(x,t)a)^{\beta}c(x,t)^{1-\beta}\beta^{-\beta}(1-\beta)^{\beta-1}$, where $c(x,t) = w(x,t) - p(x,t)s(x,t)a$.

From Section 5.2.4 onwards the rates of growth of $w(x,t)$ and form of U^* are such that space-rent grows at constant exponential rate through time and declines exponentially with x, interpreted as distance from the centre. The simplest primitives supporting this are:

i gross wage $W(t)$ equal to exogenous labour productivity and growing exponentially at constant rate $\hat{w}$, so $W(t) = W(0)\exp(\hat{w}t)$

ii Commuting costs taking fraction $1 - \exp(-\delta x)$ of gross wage, so net wage $w(x,t) = W(0)\exp(\hat{w}t)\exp(-\delta x)$

iii Cobb-Douglas preferences.

With the open city assumption $U^* = p(x,t)^{-\beta}w(x,t) = u_0$. This gives equilibrium space-rents $p(x,t) = \overline{p}\exp(\hat{p}t - \theta x)$ where $\hat{p} = \frac{\hat{w}}{\beta}$, $\theta = \frac{\delta}{\beta}$, and $\overline{p} = \frac{W(0)}{u_0}^{\frac{1}{\beta}}$. Notice that $\hat{w}$ can be interpreted as growth of the ratio of gross wages to outside utility.

Labour and population:

In the open-city equilibrium population is endogenous. Population at a point is $\frac{v}{s}$, total volume supplied divided by consumption of floor space per household. Total city population at date t is therefore

$$L(t) = \sum_{i=1}^{imax(t)} \int_{x_{i+1}(t)}^{x_i(t)} \frac{n(x)v_F(x,\tau_i)}{s_F(x,t)}dx + \int_{x_1(t)}^{x_0(t)} \frac{n(x)v_I(x,\tau_i)}{s_I(x,t)}dx$$

In this expression x is distance from the centre, and $n(x)$ is land area (in a circular city, proportional to x). The first term integrates over land in its i-th generation of development at date t (i.e. land in interval $(x_{i+1}(t), x_i(t))$, and sums over generations up to that which has been redeveloped the most times, denoted $imax(t)$. The second term is slum population.

5.D Data

This section has four components. The first discusses and describes the sources for all data used in this paper. The second deals with measures on cover/footprint and volume we use to analysis. The third gives the algorithm used to extract unchanged buildings, redeveloped buildings and infill from the overlay of 2003 and 2015 depiction of building polygons. The last reports some regression and welfare results.

5.D.1 Data sources

Building data:

We use two cross sections of data that delineate every building footprint in the city of Nairobi. The first is based on tracings of buildings from aerial photo images for 2003 that we received from the Nairobi City Council. Although no explicit metadata was provided, as far as we can tell this data was created by the Japan International Cooperation Agency (JICA) and the Government of the Republic of Kenya under the Japanese Government Technical Cooperation Program, and based on aerial images taken in February 2003 at a scale of 1:15,000. We base this off documentation from the Center for Sustainable Urban Development (CSUD) at Columbia University, who use a highly detailed building density and land-use map from the JICA (Williams & Klopp 2014). Further, we do our own data quality check by comparing the digital tracings to very high-resolution imagery from Google Earth. By examining areas that changed from 2002-2003 and from 2003-2004 we confirm that our data of building outlines matches those that exist in 2003, but did not exist in 2002, and does not include those that were yet to be built in 2003 and appeared in 2004. The second cross section comes from January 2015, when imagery at (10-20cm resolution) was recorded and digitized into building footprints by a Nairobi based company Ramani Geosystems.

The footprint data describe only the area on the ground that each building occupies while we are interested in the complete volume of each building. To address this need we supplement the 2-dimensional building data with 2015 building height data derived from LiDAR (0.3-1m resolution), also produced by Ramani Geosystems. Without direct measurements of heights in 2003, we interpolate by assigning to each building in a grid square in a sector (slum or formal) the average height of unchanged buildings in the same sector over queen neighbouring grid squares.

Slum and land use maps:

We focus on a definition of slums provided IPE Global under the Kenya Informal Settlements program (KISIP). IPE mapping of informal settlements was done using satellite imagery and topographic maps. Their approach was to identify slums as "unplanned settlements" which have some aspects of low house quality, poor infrastructure, or insecure tenure. To incorporate this definition of slums into our database we created shape files by manually digitizing KISIP documentation that contain detailed maps of all identified informal settlements in Nairobi (Limited & Consultants 2013). There remains an issue of tight delineation of slum areas, where boundaries are drawn to outline the slum areas leaving a lot of empty land residual in the formal sector that we define as the complement to slums. To offset this, we adjust the IPE slum boundaries by first classifying buildings as slum if their centre lies within the original slum boundary, and then assigning each 3m$\times$3m pixel of non-built land to slum if the nearest building is classified as slum, and formal otherwise.

A secondary set of maps that we use comes from the Center for Sustainable Urban Development (CSUD) at Columbia University. The CSUD maps land-use in 2003, including slums, based on a more detailed, copyrighted, land-use map created by the JICA and the Government of Kenya under the Japanese Government Technical Cooperation Program which was published and printed by the survey of Kenya 1000 in March 2005 (Williams & Klopp 2014). In principle, polygons are categorized as slums if they seemed to contain small mostly temporary buildings that are randomly distributed in high-density clusters. We use this set of slums to offer a descriptive comparison of how slums have changed on the extensive margin, but for our analysis we defer to a single definition based on IPE due to discrepancies in the definition of slum across the data sources. We also make use of the CSUD land-use map to identify areas that we remove from our formal classification. The areas that we chose to remove are listed in Table 5.9 and are areas in permanent public use.

Household Survey:

In order to get estimates on slum and formal household rents we use a cross section of georeferenced household level data from the 2012 'Kenya: State of the Cities' survey by the National Opinion Research Center (NORC) (Gulyani et al. 2012). This is the first survey to record household rent (with detailed house and some neighbourhood characteristics) for a sample that is stratified between slum and formal areas (based on the 2009 Census) covering Nairobi. We use the survey responses for monthly household rent and total household floor space to get a measure of the monthly rent per square meter of floor space. We then convert to annual rents in 2015 USD by appreciating 8% per year for 3 years, by 12 months a year, and at an exchange rate of 100 Kenyan shillings to the US dollar. Also included in this survey were geo-coordinates taken at the time of survey, however we found these to be imprecise when compared to the location of the enumeration area (EA) that the household was recorded to reside in. We correct household coordinates if they fall outside of their EA by replacing them with the EA's centroid coordinates.

Vacant land price listings:

We also require data on land values in order to calibrate the model, for this we rely on property values that have been scraped from (Property24 2015) over the period September 2014 to November 2015. This data source provides us with vacant land listings recording information on asking price and plot area and location, all of which are provided for in over 80% of the listings. Prices are listed as 1,000 Kenyan shillings per acre and we convert these to 2015 USD by 4047m^2 to an acre, and 100 Kenyan shillings to a US dollar. The locations are descriptive and so we entered geo-coordinates by manually searching the addresses and location descriptions. These listings are only found in the formal sector.

SRTM elevation:

Elevation and ruggedness measures used in regression tables are calculated from the Shuttle Radar Topography Mission (SRTM), a grid of 1 arc-second wide cells (or roughly 30 metres in Nairobi) published by the (USGS 2005). Elevation is simply the mean of these cells in each of our 150m$\times$150m gridcells, while we measured ruggedness as the standard deviation in elevation within each 150m$\times$150m gridcell.

SPOT Imagery:

We also use high resolution SPOT5 and SPOT6 images of Nairobi for 2004 and 2015 respectively. The raw imagery was created by Airbus Defence and Space and we used it as reference to manually trace roads and define their widths in order to come up with estimates of the extent of road coverage in both the early and late time periods. Alternative sources, like Open Streetmap, were unsuitable as they did not allow us to make the comparison across time.

5.D.2 Measures of cover and volume

Our unit of analysis is 150m×150m grid squares. For calculating cover within the grid square in a usage, each of these is broken into 2,500 3m×3m cells and use type classified by what is at the centroid of the 3m square in each period. There are three uses: vacant land, slum area and formal. For each 150m×150m square we sum across the 2,500 cells to get total use of each type. Most 150m×150m squares are either all slum or all formal sector. However, there are about 12% which are mixed grid squares, for which we record the cover or volume of slum and formal separately.

Having summed the total area of use of each type in 3m×3m squares in each 150m×150m square, these are averaged for 150m×150m squares whose centroid falls in a narrow distance ring. That sum is then divided by the total number of 150m×150m grid squares in that distance band. For volume for 2015, for each 3m×3m square which is formal sector, we have the height of the building at the centroid of that square. Volume for that 3m×3m square is nine times the height in meters of the building from LiDAR data. We then sum across the grid squares occupied with formal usage for 150m×150m grid squares in each distance ring and then average by the total number of 150m×150 m grid squares in the ring. For 2003 we have no height data. To infer 2003 heights, we use what we think is an upper bound on height: the height of unchanged buildings, where we presume demolished buildings between 2003 and 2015 are likely to be of lower height than those which survive. To assign a height to a 3m×3m square in 2003 in formal sector usage, we take the average height in 2015 of all buildings that were there in 2003 for all 3m×3m formal sector unchanged buildings in the own 150m×150m grids square and its 8 queen neighbours. Height is the height assigned to each 3m×3m square in usage in a distance ring from the centre averaged over all such cells, to effectively get a coverage weighted average of individual building heights.

How do we measure change between 2003 and 2015? For demolition, at the 3m×3m level the square is defined as demolition if its centroid is covered by a 2003 building that has become open space. Demolished coverage is lost 2003 cover; demolished volume is assessed as before using the average height of unchanged buildings in the neighbourhood. Infill is new buildings that do not overlap with any 2003 buildings; a 3m×3m square is infill if its centroid is covered by such a building on 2015 where there was no building in 2003. Infill cover and volume are assessed from 2015 data. Net redevelopment in coverage takes coverage in the new 2015 buildings and subtracts the coverage of old 2003 buildings. So for each 150m×150m meter square we have for redeveloped buildings, we have total coverage in 2003 measured at the 3m×3m level (centroid covered by the old 2003 building(s)) and we have total coverage in 2015 measured at the 3m×3m squares (centroid covered by the new replacement 2015 building(s)). Net redevelopment at the 150m×150m square is the difference. In general, the same buildings are drawn in 2015 to have modestly more coverage than in 2003 so coverage change is likely to be an upper bound. Net volume change again assigns heights in 2003 to the 3m×3m coverage based on neighbourhood averages for unchanged buildings and uses 2015 height information on the new buildings.

5.D.3 Overlaying buildings

We match buildings across time by overlaying 2015 and 2003 building polygon data in order to track the persistency, demolition, construction and reconstruction of buildings over time. Since buildings are not identified across time our links rely on a shape matching algorithm. For each building, the algorithm determines whether it was there in the other period, or not, by comparing it with the buildings that overlap in the other time period.

This task is not straightforward, since the same building can be recorded in different ways depending on the aerial imagery used, whether building height was available, and the idiosyncrasies of the human digitizer.

Data and definitions:

For 2003 we use the building dataset received from the Nairobi City Council with digitized polygons for every building, roughly 340,000 in the administrative boundary of Nairobi. For 2015 we use the dataset that was created by Ramani Geosystems using imagery (10-20cm resolution). The nomenclature we use is as follows. First, a *trace* is the collection of polygon vertices that make up its outline. A *shape* is the area enclosed by the trace, and can be thought of as a representation of the rooftop of a building. A *cavity* is an empty hole completely enclosed in a shape. A *candidate pair* is the set of any two shapes in different time periods which spatially intersect. A *link* is the relationship between a set of candidates in one period to a set of candidates in the other time period.

Pre-processing:

Before running our shape matching algorithm we clean up the data sets. First we take care of no data areas. There are some areas that were not delineated in 2003, including the Moi Air Base, and Nairobi State House. We drop all buildings in these areas for both 2003 and 2015, amounting to roughly 1,500 buildings from the 2015 data, and 100 buildings from 2003. Next we deal with overlapping shapes, an issue arising in the 2015 data, although not that for 2003. This is most often the same building traced multiple times. We identify all such overlapping polygons and discard the smaller version until no overlaps remain; about 1,400 buildings from the 2015 data this way. We also drop small shapes, in part because the 2015 data has many very small shapes, while the 2003 data does not. In order to avoid complications of censoring in the 2003 data, we simply drop all shapes that have an area of less than 1m^2. We drop two small buildings in 2003, and 462 small buildings in 2015.

Another issue is that buildings are often defined as contiguous shapes in 2003, but broken up in 2015. For the majority of buildings we cannot aggregate the broken up pieces in 2015 since it is hard to identify such cases in general. To match these cases across time we rely on our one to many, and many to many matching algorithms defined below. However, in the specific case where a building is completely enclosed in another the task is much easier. First, we find all cavities present in each period, then we take all building shapes that overlap with the cavities in the same time period. After identifying all shapes that intersect a cavity, we redefine both shapes, the original shape containing the cavity and the shape intersecting it, as a single new shape.

Shape Matching Algorithm:

After the pre-processing of each cross-section is complete, we run our shape matching algorithm to establish links between buildings across time periods. For any given building we consider five possible scenarios; that it has a link to no building, that it has a link to one building (one to one match), that it has a link to multiple buildings (one to many), that it is part of a group of buildings that match to one building (many to one), or that it is a part of a group of buildings that matches to a group of buildings (many to many). We follow and approach similar to Yeom et al. (2015) however, due to the inherent difficulty of inconsistent tracings we contribute to their method by introducing the one to many and many to many approaches. We assign each link a measure of fit that we call the overlay ratio. We then choose optimal links based on the overlay ratio. Finally, we categorize links as matched or not using a strict cut-off on the overlay ratio of 0.5. Other cu-offs such as 0.4, 0.6 and 0.7 produced more errors in categorization.

Candidates:

For all buildings A in the first time period, and B in the second time period we identify the set of candidates:

$$CP = \{(A, B);\ Area(A \cap B) \neq 0\}$$

For each candidate pair we find the ratio of the intersection area over the area of each shape, so if shapes A and B intersect, we find $r_{AB} = \frac{Area(A \cap B)}{Area(A)}$ and$r_{BA} = \frac{Area(A \cap B)}{Area(B)}$. We link all shapes which do not belong to a candidate pair to the empty set.

One-to-One Matching:

First we consider candidate pairs to be links on their own. For each pair, we calculate the overlay ratio as the intersection area over union area, so if A and B are candidate pair, we find:

$$R_{AB} = \frac{Area(A \cap B)}{Area(A \cup B)} = \frac{Area(A \cap B)}{Area(A) + Area(B) - Area(A \cap B)}$$

One-to-Many Matching:

For each time period separately, we identify all candidate pair links for which their intersection to area ratio is above threshold θ. For shape A we define a group $= \{B; r_{BA} \geq \theta\}$. Now we calculate the overlay ratio of one to many links as the intersection area over union area ratio:

$$R_{AG} = \frac{Area(A \cap \cup_{B \in G} B)}{Area(A \cup \cup_{B \in G} B)} = \frac{\sum_{B \in G} Area(A \cap B)}{\sum_{B \in G} Area(A \cup B)}$$

Many-to-Many Matching:

Here we have two cases, one when the shapes are fairly similar, which we capture in previous sections (one to one, or many to one). The other is inconsistent shapes that form the same structure. To capture these we consider both time periods at the once, we clean the candidate pair list, keeping links for which either ratio is above a threshold θ_1:

$$LC = \{(A, B);\ r_{AB} \geq \theta_1 \mid r_{BA} \geq \theta_1\}$$

Then we condition to only keep shape for which the total ratio intersection is above threshold θ_2, so shape A will be included if $\sum_{B \in \{x|(A,x) \in LC\}} r_{AB} \geq \theta_2$. Now we are left with a new candidate list, which we convert to sets $LC = \{(\{A\}, \{B\})\}$ and start merging them:

$$LC = \{(G_i \cup G_j, H_i \cup H_j)\} \cup LC/\{(G_i, H_i), (G_j, H_j)\}, i \neq j \text{ if } G_i \cap G_j \neq \emptyset \mid H_i \int H_j \neq \emptyset$$

We keep doing this until we can no longer merge any two rows. At this point we calculate the overlay ratio of many to many links as the intersection area over union section ratio:

$$R_{GH} = \frac{Area(\cup_{A \in G} A \cap \cup_{B \in H} B)}{Area(\cup_{A \in G} A \cup \cup_{B \in H} B)}$$

ICP Translation:

We encounter a problem when the two shapes or groups of shapes are similar but do not overlap well, this usually stems from the angle at which the images were taken, and is especially prevalent with tall buildings. To address this issue, we translate one trace towards the other, and then recalculate the overlay ratio. As in Besl & Mckay (1992), we use the iterative closest point (ICP) method to estimate this translation. To perform the ICP we ignore any cavity points as we found they often cause less suitable translation. We found that for similar shapes this will optimize the intersection area.

Optimal Linking:

In the end, we rank all links by their overlay ratio. We iteratively keep the link with the highest overlay ratio, or discard it if at least one of the buildings in the link has already been confirmed in a separate link. From the list of optimal links, we define a link to be a match if its overlay ratio, or the overlay ratio after ICP translation is above 0.5. We then define all matched candidates as unchanged, and the remaining candidates as redeveloped. All buildings that were not considered as candidates are defined as infill, if from 2015, and demolished, if from 2003.

Accuracy Assessment:

In order to assess the performance of the polygon matching algorithm we manually classified links between 2003 and 2015 for a random sample of buildings. We sampled 48 150m$\times$150m grid cells, stratifying over slum, non-slum within 3km, non-slum within 6km, and non-slum further than 6km to the centre. The sample consists of over 2,250 buildings in 2003 and 3,500 buildings in 2015.

Results:

We first break down matches by their mapping type. There are five types of manual link: redeveloped/infill/demolished (0), one to one match (1), one to many match (2), many to one match (3), and many to many match (4). For the algorithm we further split (0) into infill/demolished (-1) and redeveloped (0). Table 5.6 shows the correspondence between the two mappings by building (a) and roof area (b). We can see that most errors come from the one to one matches, however, the many to many matches have the worst performance. Overall the diagonal values are quite high, which means not only are we matching buildings well, but also the algorithm is recognising the clumping of buildings as a human does (bear in mind that, for example, the one to one matches which we 'misclassify' as many to many will still be classified as match in the final data). Finally, we have perfect correspondence for demolition and in 2015 nearly perfect for infill.

Next we compare buildings that were matched by the algorithm and those matched manually. For now we use a cut-off of the overlay ratio of 0.5, later we explore the effect of different cut-offs on performance. As seen in Table 5.6 infill and demolition are classified with almost perfect correspondence. For this reason we ignore buildings with these mappings and focus on accuracy of redevelopment and unchanged. In Table 5.7 we condense mappings 1, 2, 3, and 4 into category 1, while redevelopment, or category 0, remains the same.

We define precision P (negative predictive value NPV) as the fraction of buildings classified as unchanged (redeveloped) by the algorithm that are correct, recall R (true negative rate TNR) as the fraction of buildings classified as unchanged (redeveloped) by hand that the algorithm gets correct, and the F1 score (F) as the weighted average of the two.

$$
\begin{aligned}
P &= \frac{\text{True Positives}}{\text{Positive Predictions}}, \\
NPV &= \frac{\text{True Negative}}{\text{Negative Predictions}}, \\
R &= \frac{\text{True Positives}}{\text{Positive Condition}}, \\
TNR &= \frac{\text{True Negative}}{\text{Negative Condition}}, \\
F &= \frac{\text{2*P*R}}{\text{P+R}}
\end{aligned}
$$

The confusion matrix in Table 5.7 is done across all sampled buildings in 2003 and weights observations by buildings (1) and roof area (2). The F1 score is high in both cases, but in part this is due to relative success classifying unchanged buildings: precision for buildings that were classified as redeveloped by the algorithm is 76% of buildings and 72% of roof area, while recall of true redeveloped buildings is 83% of buildings and 74% of roof area.

In our first attempt we arbitrarily picked 50% as a cut off of the overlay ratio. Here we take a closer look at this choice. Using our manually classified links we can maximize the F1 score with respect to the cut off. In Figure 5.6 we plot the F1 score weighted by roof area against cut-offs of the overlay ratio for the 2003 data. We find that the highest F1 score comes just below 50% suggesting our first estimate was not far off.

In Figure 5.6 we plot lines for each method of calculating the overlay ratio: without ICP, with ICP, and the maximum of the two. Around 50% we can see that the maximum performs best, but with only a very slight improvement over the ICP alone, which is in turn marginally better than without the ICP.

5.D.4 Context issues: Monocentricity, and real estate data

Monocentricity

We look at a set of cities in the developing versus developed world in terms of their degree of monocentricity using a traditional indicator: steepness of the population density gradient, representing how sharply density declines from the city centre. A flatter slope indicates a lack of centrality, or force of pull of the centre. The sample cities is based upon the dataset of World Urbanization Prospects: The 2014 Revision, produced by the Population Division of the UN. The selection procedure follows three rules: (1) Select 10 cities from SSA and another 10 from Europe or USA; (2) Given the population of Nairobi is 3.7 million in the dataset, we select cities in the range from 1.5 million to 5.5 million; (3) Prioritize cities with population closer to the population in Nairobi to the sample.

Population is measured by Landscan which captures the ambient population at a 1km grid square level. We exclude grids with zero population and go out to 10km from city centre, cutting out water areas. To define the city centre, we:

1. Find all light grids with light radiance readings in the top 1% ranking within the city.
2. Define light clusters based on spatial rook contiguity for just these 1%.
3. Find the largest cluster with the highest sum of light radiance.
4. Define the grid with the highest radiance in the largest cluster as the center.

We then run a regression where the dependent variable is ln of grid pop. With a 10km radius, we can have up to about 315 sq per city. Regressions have city fixed effects and the basic explanatory variable is distance from the city centre. In Table 5.9 we show the results for developed vs developing country cities. Results in columns 1 and 2 show a distinctly steeper slope for developing country cities. In developing country cities, Nairobi with gradient slope of -0.273 has 5th steepest gradient out of 10. In developed countries only San Francisco at -0.346 is steeper.

Real estate data and robustness

We use quarterly reports on the Nairobi housing market from HassConsult, a market leader in the Kenyan Real Estate sector.[37] These reports have been an important source of information for Kenyan real estate investors and homeowners for over ten years. The primary source of data is based on listings available in public sources (newspapers, magazines, social

[37] All original reports were last downloaded from the HassConsult report archive on May 5th, 2020: hassconsult.co.ke/real-estate/hass-index/16-report-archives.

media, and online property portals). These listings are cleaned by HassConsult to remove properties and vacant land that are deemed atypical or duplicates, and then averaged in a way that ensures that price changes are not driven by composition effects.

The Hass Composite Sales and Lettings indices go back to 2000, and over its first 17 years used 163,000 cleaned listings after having dropped 15% due to duplications and 4% as outliers. Assuming that the collection rate has increased since the initial years, we can give a lower bound to the number of property listings used each year of 10,000. The Hass Land Composite Index goes back to 2007, and over its first ten years used 58,000 cleaned land listings. Assuming that the collection rate has increased since the initial years, we can give a lower bound to the number of land listings used each year of 5,800. The indices are broken down geographically into 18 Nairobi Suburbs[38] which is what we use and 14 Nairobi Satellite Towns.[39] We geolocate these neighbourhoods by a manual google maps search. Each neighbourhood had a distinct polygon which we were able to trace and georeference neighbourhood boundaries. We use the centroid of these boundaries for any spatial analysis of the HassConsult data.

We first looked for bubbles. Based on Glaeser et al. (1992), bubbles are hard to define; but we looked at two common indicators for which there are data: overbuilding and price paths. The evidence suggests there is no overbuilding, and price increases are not high and see quite steady around the 2015 time period. First on sustained overbuilding, as noted in the Introduction, built volume increased over the 12 years by 59%, while population growth increased by 67%. Second, we investigated the rate of price increases using data from HassConsult. For looking at bubbles, researchers look at property value data, to see if rates of increases are very high or if a bubble has burst. For the full set of communities, prices rose 2.38 fold from 2007 to 2019 for an average annual increase of 7.5%. For years around our Q4 2015 data from the Q4 2014, Q4 2015 and Q4 2016 archived reports, the prior year's annual rate of property value changes are respectively 8.3, 9.6 and 7.6, so 2015 looks like a typical year. The official inflation rate varies a lot annually but averaged 7.4% p.a. from 2010 to 2019, with about a 6% rate around 2015. The property price changes are not indicative of any bubble; they seem steady and not high.

We also used the HassConsult reports on rental prices to look at real rates of increase in space-rents. One index uses the full sample of communities, where over time more

[38] The HassConsult suburbs are: Donholm, Eastleigh, Gigiri, Karen, Kileleshwa, Kilimani, Kitisuru, Langata, Lavington, Loresho, Muthaiga, Nyari Estate, Parklands, Ridgeways, Spring Valley, Upperhill and Westlands.

[39] The HassConsult satellite towns are: Athi River, Juja Town, Kiambu Town, Kitengela, Limuru, Mlolongo, Ngong Town, Ongata Rongai, Ruaka, Ruiru, Syokimau, Thika Town and Tigoni.

properties in outlying lower quality housing areas are added. Rental prices for this sample of communities have risen 1.85 fold over 2007-2019, which gives a 5.3% average annual rental price increase. Another uses more traditional central communities without this bias.[40] House rental price increases for each of these communities range from 1.96 to 3.25 fold over the period 2007-2019, or 5.8% to 10.3% average annual rental price increases. Inflation was averaging 7-8% p.a., so the 0.94% annual rate of real space-rent increase that we use falls between these numbers.

Finally, using our georeferenced neighbourhoods and the HassConsult report on land prices for Quarter four of 2015, we take the average listing price per square meter of vacant lots in the 18 Nairobi Suburbs mentioned above to look at land price gradients. This serves as a robustness check of the analysis of property24 data for which we have microdata, but may not cover as broad of a sample as does the HassConsult data. We measure the distance from the neighbourhood centroid to the city centre. The estimated gradient from the simple regression of ln price per square meter on distance to centre in kilometres gives a similar slope -0.198 (0.0309); within one standard error of that estimated with property24 data.

5.E Land Quality

In order to address concerns about heterogeneity in land quality driving some of our empirical results this appendix examines the role of geographical characteristics. Our underlying data has a fine resolution, based on 30m×30m cells. We analyse how geography varies over small spatial scales for slum boundaries and infill, or greenfield developments. Cells are defined as exclusively slum or formal based on where their centroid lies. Elevation is measured using the SRTM data at the 30m×30m cell, and its mean and standard deviation are calculated over moving windows of 90m×90m, 150m×150m, and 450m×450m (USGS 2005) around the own cell to measure ruggedness at different scales. To determine the relative local elevation we calculate the difference between own elevation and mean elevation in the 150m×150m and 450m×450m windows. Similarly using the CSUD land-use map we digitize water bodies (rivers, lakes, ponds, etc.) and distinguish whether a cell contains water in its own cell or 90m×90m window (Williams & Klopp 2014). Further, cells are classified as infill if the cell contained only infilled buildings in 2015 and no buildings in 2003. With a range of window sizes, we focus on the 'small' 90m×90m and the 'large'

[40] The central neighbourhoods we choose are: Lavington, Kileleshwa, Kilimani, Westlands, and Muthaiga. From 2007 to 2019 rental prices in these areas rose by 2.69, 2.30, 1.96, 3.25, and 2.66 fold respectively, so the average annual rate of change in rent for each was 8.6%, 7.2%. 5.8%, 10.3%, and 8.5% respectively.

which is taken as 450m×450m except in order to avoid excessive overlap when considering boundary analysis where we use 150m×150m.

First we consider the role of land heterogeneity in explaining infill in areas near to the city centre, which according to the model should have been developed before those further out. In Table 5.13, for all formal sector land, we show that infill occurs on lower quality land, and this differential is largest near the centre. Near the centre, infill tends to occur on land with higher elevation, nearer to water, and that is more rugged. Infill also tends to be on land that is lower than other cells in its 450m×450m window (effect dissipating over distance), however the differential gradient is insignificant and so only suggestive. Together this suggests that infill occurs in places that are generally higher up, but locally downhill. Proximity to water and ruggedness are more easily interpreted as raising costs to development and here we find strong and significant results.

Next in Table 5.14, we look at whether heights of redeveloped buildings vary with land quality. If during the initial formal development, a sunk cost is paid to prepare the land; draining swamps, levelling the land, etc., then during successive periods of redevelopment the height of buildings should not vary at a given distance. In the table we show that only one of our five land quality measures significantly affect redeveloped building heights. There the standard deviation of elevation in the small scale (90m×90m) is associated with higher redeveloped buildings, rather than reduced height.

Finally, we are concerned that results in our welfare analysis may be driven by land quality. That is there may be a correlation between government owned slum land (where formalisation costs are high) and low quality land (where construction costs are high for natural reasons). In particular, if central slum land was on worse land than neighbouring greenfield formal this could be partially responsible for the gap in land rents. We focus on infill land because, over the past period, it is the relevant comparison for what could have been slum redevelopment. In Table 5.15 we look at the sample of cells that are within 300m of a government slum boundary, restricted to either government slum cells or formal infilled land. We run a fixed effects regression so the analysis focuses on variation within arbitrary 300m×300m blocks. Within these neighbourhoods we compare cells in and outside of slums at distance bins from the city centre for different land quality outcomes. Results show that, especially for slums inside of 6km, where we focus our welfare analysis, there is very little difference in land quality between slums and neighbouring formal land. Inside 6km slums are only found to be nearer to water in the bin from 4-5km, and in the bin 2-3km slums are actually less rugged and on higher local land in the 450m×450m window.

So for the 20 possible coefficients there is only one suggesting significantly lower quality in slums compared to neighbouring formal areas.

Bibliography

Ahabwe, J., Kikonyogo, H., Manwaring, P., Mugangaizi, R. R., Mutebi, J., Nuweabine, D. M. & Regan, T. (2020), Enhancing property tax compliance in Kampala: Initial evidence from taxpayer consultations and administrative data. IGC policy note 19147.

Ahlfeldt, G. & McMillen, D. (2014), New estimates of the elasticity of substitution of land for capital. Working paper No. 14GA1, Lincoln Land Institute.

Ahlfeldt, G. & McMillen, D. (2018), 'Tall buildings and land values: Height and construction cost elasticities in chicago, 1870–2010', *Review of Economics and Statistics* **100**(5), 861–875.

Ahlfeldt, G., Redding, S., Sturm, D. & Wolf, N. (2015), 'The economics of density; evidence from the berlin wall', *Econometrica* **83**, 2127–2189.

Ali, D. A., Deininger, K. & Wild, M. (2018), Using Satellite Imagery to Revolutionize Creation of Tax Maps and Local Revenue Collection. World Bank Policy Research Working Paper 8437.

Ali, D., Collin, M., Deninger, K., Dercon, S., Sandefur, J. & Zeitlin, A. (2016), 'Small price incentives increase women's access to land titles in Tanzania', *Journal of Development Economics* **123**, 107–122.

Alvaredo & Gasparini (2015), Recent trends in inequality and poverty in developing countries, *in* A. Atkinson & F. Bourguignon, eds, 'Handbook of Income Distribution', Vol. 2, Elsevier, chapter 9.

Amis, P. (1984), 'Squatters or tenants: The commercialization of unauthorized housing in nairobi', *World Development* **12**(1), 87–96.

Andema, F. & Haas, A. (2018), Efficient and effective municipal tax administration: a case study of the Kampala Capital City Authority. IGC Cities that Work Case Study.

Avenancio-Leon, C. & Howard, T. (2020), The Assessment Gap: Racial Inequalities in Property Taxation. working paper.

Bachas, P., Gadenne, L. & Jensen, A. (2019), Informality, Consumption Taxes and Redistribution. Working paper.

Balan, P., Bergeron, A., Tourek, G. & Weigel, J. (2020), Local Elites as State Capacity: How City Chiefs use Local Information to Increase Tax Compliance in the D.R. Congo. CEPR Discussion Paper DP15138.

Basurto, P., Dupas, P. & Robinson, J. (2018), Decentralization and efficiency of subsidy targeting: Evidence from cheifs in rural Malawi. Working paper.

Becker, G., DeGroot, M. & Marschak, J. (1964), 'Measuring Utility by a Single-Response Sequential Method', *Business Science* **9**(3), 226–234.

Behrens, J. (1977), 'Property Tax Administration: Use Of Assessment Ratios', *Journal of Education Finance* **3**, 158–174.

Bergeron, A., Tourek, G. & Weigel, J. (2019), The elasticity of tax compliance: evidence from randomized property rates. Pre-analysis plan (RCT ID AEARCTR-0003818).

Berry, J., Fischer, G. & Guiteras, R. (2020), 'Eliciting and Utilizing Willingness-to-Pay: Evidence from Field Trials in Northern Ghana', *Journal of Political Economy* **128**(4), 1436–1473.

Besl, P. & Mckay, N. (1992), 'A method for registration of 3-d shapes', *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(2), 239—256.

Besley, T. (1995), 'Property Rights and Investment Incentives: Theory and Evidence from Ghana', *Journal of Political Economy* **103**(5), 903–937.

Besley, T. & Ghatak, M. (2010), Property rights and economic development, *in* D. Rodrik & M. Rosenzweig, eds, 'Handbook of Development Economics', Vol. 5, Elsevier, chapter 68, pp. 4525–4595.

Besley, T. & Persson, T. (2014), 'Why Do Developing Countries Tax So Little?', *Journal of Economic Perspectives* **28**(4), 99–120.

Bezu, S. & Holden, S. (2014), 'Demand for second-stage land certification in Ethiopia: Evidence from household panel data', *Land Use Policy* **41**, 193–205.

Bird, J. & Venables, A. (2020), 'The efficiency of land-use in a developing city: traditional vs modern tenure systems in Kampala, Uganda ', *Journal of Urban Economics* . forthcoming.

Braid, R. (2001), 'Spatial growth and redevelopment with perfect foresight and durable housing', *Journal of Urban Economics* **49**, 425–452.

Briggs, J. (2011), 'The Land Formalisation Process and the Peri-Urban Zone of Dar es Salaam, Tanzania', *Planning Theory and Practice* **12**, 115–153.

Brockmeyer, A., Ramirez, K., Estefan, A. & Serrato, J. C. S. (2020), Taxing Property in Developing Countries: Theory and Evidence from Mexico. working paper.

Brockmeyer, A., Smith, S., Hernandez, M. & Kettle, S. (2019), 'Casting a Wider Tax Net: Experimental Evidence from Costa Rica', *American Economic Journal: Economic Policy* **11**(3), 55–87.

Brueckner, J. & Rosenthal, S. (2009), 'Gentrification and neighbourhood housing cycle', *Review of Economics and Statistics* **91**, 725–743.

Case, K. (2019), The value of land in the united states: 1975-2005, *in* G. Ingram & H. Y., eds, 'Land Policies and Their Outcomes', Lincoln Institute of Land Policy.

Castro, L. & Scartascini, C. (2015), 'Tax compliance and enforcement in the pampas evidence from a field experiment', *Journal of Eocnomic Behaviour & Organization* **116**, 65–82.

Cattaneo, M., Galiani, S., Gertler, P., Martinez, S. & Titiunik, R. (2009), 'Housing, health, and happiness', *American Economic Journal: Economic Policy* **1**(1), 75–105.

Cavalcanti, T., Da Mata, D. & Santos, M. (2019), 'On the determinants of slum formation', *The Economic Journal* **129**, 1971–1991.

Cheeseman, N. (2008), 'The kenyan elections of 2007: An introduction', *American Economic Journal: Economic Policy* **2**(2), 166–84.

Collier, Glaeser, Venables, Blake & Manwaring (2017), Secure, legally enforceable and marketable land rights for urban development. Working paper.

Combes, P., Duranton, G. & Gobillon, L. (2017), The production function for housing: Evidence from france. CEPR Discussion Paper No. DP11669.

Combes, P., Duranton, G. & Gobillon, L. (2019), 'The costs of agglomeration: House and land prices in french cities', *Review of Economic Studies* **86**, 1556–1589.

Davis, M. (2006), *Planet of Slums*, Verso, London.

Del Carpio, L. (2014), Are The Neighbors Cheating? Evidence from a Social Norm Experiment on Property Taxes in Peru. working paper.

Duflo, E., Galiani, S. & Mobarak, M. (2012), Improving access to urban services for the poor: Open issues and a framework for a future research agenda, Technical report, Cambridge, MA.

Duranton, G. & Puga, D. (2015), Urban land use, *in* G. Duranton, J. Henderson & W. Strange, eds, 'Handbook of Regional and Urban Economics', Elsevier.

Etherton, D. (1971), Mathare valley: A case study of uncontrolled settlement in nairobi, Technical report, Nairobi, Keny.

Feenstra, R., Inklaar, R. & Timmer, M. (2015), 'The next generation of the penn world table', *American Economic Review* **105**(10), 3150–82.

Ferreira, P., Monge-Naranjo, A. & Torres De Mello Pereira, L. (2016), Of cities and slums. FRB St. Louis Working Paper No. 2016-22.

Field, E. (2007), 'Entitled to Work: Urban Property Rights and Labor Supply in Peru', *The Quarterly Journal of Economics* **122**(4), 1561–1602.

Franzen, R. & McCluskey, W. (2017), *Property Tax in Africa: Status, Challenges, and Prospects*, Lincoln Land Institute of Land Policy, Cambridge, Massachusetts.

Fujita, M. (1982), 'Spatial development processes with renewal in a growing city', *Journal of Economics* **12**, 22–52.

Galiani, S. & Schargrodsky, E. (2010), 'Property rights for the poor: Effects of land titling', *Journal of Public Economics* **94**, 700–729.

Gechter & Tsivanidis (2017), Efficiency and Equity of Land Policy in Developing Country Cities: Evidence from the Mumbai Mills Redevelopment. World Bank WPS8163.

Geosystems, R. (2015), Nairobi building footprints and heights. Nairobi, Kenya: Multiple Electronic Files.

Glaeser, E., Huang, W., Ma, Y. & Shleifer, A. (1992), 'A real estate boom with chinese characteristics', *Journal of Economic Perspectives* **31**(1), 93–116.

Glennerster, R., Kremer, M., Mbiti, I. & Takavarasha, K. (2011), Access and quality in the kenyan education system: A review of the progress, challenges and potential, Technical report, Nairobi, Kenya.

Goldstein, M. & Udry, C. (2008), 'The Profits of Power: Land Rights and Agricultural Investment in Ghana', *Journal of Political Economy* **116**, 981–1022.

Gollin, D., Kirchberger, M., & Lagakos, D. (2017), In search of a spatial equilibrium in the developing world. NBER working Paper 23916.

Gordon, R. & Li, W. (2009), 'Tax structures in developing countries: Many puzzles and a possible explanation', *Journal of Public Economics* **93**, 855–866.

Gulyani, S., Ayres, W., Struyk, R. & Zinnes, C. (2012), Kenya state of the cities baseline survey 2012-2013, Technical report, Washington DC. Ref: KEN_2012_SOCBL_v01_M.

Gulyani, S. & Talukdar, D. (2008), 'Slum real estate: The low-quality high-price puzzle in nairobi's slum rental market and its implications for theory and practice', *World Development* **36**(10), 1916–37.

Harari, M. (2020), 'Cities in bad shape: Urban geometry in india', *American Economic Review* **110**, 2377–2421.

Harari, M. & Wong, M. (2018), Slum upgrading and long-run urban development: Evidence from indonesia. Working paper.

Hass, A. (2017), Property taxes: Exploring the untapped potential for the city of hargeisa. IGC Cities that work.

Hass, A. (n.d.), From moving vehicles to moving people: Designing a mass public transportation system for kampala.

Henderson, J. V., Regan, T. & Venables, A. (2020), 'Building the City: From Slums to a Modern Metropolis', *Review of Economic Studies* . forthcoming.

Henderson, J. V., Venables, A., Regan, T. & Samsonov, I. (2016), 'Building functional cities', *Science* **352**(6288), 946–947.

Henn, S. (2020), Complements or substitutes? how institutional arrangements bind chiefs and the state in africa. Working paper.

Huchzermeyer, M. (2007), 'Tenement city: The emergence of multi-storey districts through large-scale private landlordism in nairobi', *International Journal of Urban and Regional Research* **31**, 714–732.

International Association of Assessing Officers (2017), Standard on mass appraisal of real property, Technical report, Kansas City, MO. ISBN 978-0-88329-2075.

International Association of Assessing Officers (2018), Standard on automated valuation models, Technical report, Kansas City, MO. ISBN 978-0-88329-245-7.

IPUMS (2019). Minnesota Population Center. Integrated Public Use Microdata Series, International: Version 7.2 [dataset].

Jibao, S. & Prichard, W. (2016), 'Rebuilding Local Government Finances After Conflict: Lessons from a Property Tax Reform Programme in Post-Conflict Sierra Leone', *The Journal of Development Studies* **52**(12), 1759–1775.

Joireman, S. F. & Vanderpoel, R. S. (2011), In search of order: State systems of property rights and their failings, *in* S. F. Joireman, ed., 'Where There is No Government: Enforcing Property Rights in Common Law Africa', Oxford University Press.

Kangave, J., Nakato, S., Waiswa, R., Nalukwago, M. & Zzimbe, P. L. (2018), What Can We Learn from the Uganda Revenue Authority's Approach to Taxing High Net Worth Individuals? ICTD Working Paper 72.

Kelly, R. (2003), Property taxation on kenya, Technical report. worldbank.org/publicsector/decentralization/June2003Seminar/Kenya.pdf.

Kenya (2018), Report on riparian lands conservation and management policy framework, Technical report.

Kenya, C. (1999), Population and housing census, Technical report.

Kenya, K. (2009), Kenya population and housing census, Technical report.

Klopp, J. M. (2000), 'Pilfering the public: The problem of land grabbing in contemporary kenya', *Africa Today* **47**(1), 7–26.

Knebelmann, J. (2019), Taxing property owners in Dakar. IGC Policy brief 50415.

Kopanyi, M. (2015), Local revenue reform of Kampala Capital City Authority. IGC working paper C-43306-UGA-1.

Kremer, M. & Snyder, C. (2018), Worst-case bounds on r&d and pricing distortions: Theory with an application assuming consumer values follow the world income distribution. Working paper.

Lall, S., Henderson, J. V. & Venables, A. (2017), Africa's cities: Opening doors to the world, Technical report, Washington, D.C. openknowledge.worldbank.org/handle/10986/25896.

Limited, I. G. P. & Consultants, S. (2013), Consultancy services for city/municipal situational analysis of conditions of informal settlements in 15 municipalities, Technical report, Nairobi, Kenya.

Manara, M. (2020), Formal, informal, and legitimate. Implementing land institutions in Dar es Salaam. working paper.

Manara, M. & Pani, E. (2020), Complementarities in the co-evolution of property and credit institutions: how financial actors manage risk in the nascent credit market of Dar es Salaam. working paper.

Manara, M. & Regan, T. (2020), Determinants of demand for title deeds in urban Tanzania. working paper.

Manwaring, P. & Regan, T. (2020), Enhancing property tax in Kampala: Successes, challenges, and next steps for increasing municipal revenue. IGC policy brief 43448.

Marx, B., Stoker, T. & Suri, T. (2013), 'The economics of slums in the developing world', *Journal of Economic Perspectives* **27**(4), 187–210.

Marx, B., Stoker, T. & Suri, T. (2019), 'There is no free house: Ethnic patronage in a kenyan slum', *American Economic Journal: Applied Economics* **11**(4), 36–70.

McMillen, D. (2013), 'The Effect of Appeals on Assessment Ratio Distributions: Some Nonparametric Approaches', *Real Estate Economics* **41**(1), 165–191.

McMillen, D. & Singh, R. (2020), 'Assessment Regressivity and Property Taxation', *Journal of Real Estate and Financial Economics* **60**, 155–169.

McMillen, D. & Weber, R. (2010), 'Ask and ye shall receive? Predicting the successful appeal of property tax assessments', *Public Finance Review* **38**(1), 74–101.

Medard, C. (2006), City planning in nairobi: The stakes, the people, the sidetracking, *in* D. Rodrigues-Torres, ed., 'Nairobi Today: The Paradox of a Fragmented City', Mkuki na Nyota Publishers.

Michaels, G., Nigmatulina, D., Rauch, F., Regan, T., Baruah, N. & Dahlstrand-Rudin, A. (2020), Planning ahead for better neighbourhoods: Long run evidence from Tanzania. working paper.

Michalopoulos, S. & Papaioannou, E. (2013), 'Pre-colonial ethnic institutions and contemporary African development', *Econometrica* **81**(1), 113–152.

Moore, M. & Wilson, P. (2017), How Can Governments of Low-Income Countries Collect More Tax Revenue? ICTD Working Paper 70.

Moses, M. & Chiwambo, D. (2018), 'Assessing land titles application and uptake in regularised informal settlements in Tanzania', *International Journal of Urban Sustainable Development* **10**(3), 279–291.

Mwaura, A. (2006), Policy review for zones 3, 4 and 5, nairobi, kenya 42nd isocarp congress. www.isocarp.net/Data/case_studies/886.pdf.

Naritomi, J. (2019), 'Consumers as Tax Auditors', *American Economic Review* **109**(9), 3031–3072.

Narubi, A. (2015), An evaluation of property rate collection and enforcement of devolved systems of governance: A case study of nairobi city-county. MA thesis, University of Nairobi, Department of Real Estate and Construction Management.

Niehaus, P., Atanassova, A., Bertrand, M. & Mullainathan, S. (2013), 'Targeting with Agents', *American Economic Journal: Economic Policy* **5**(1), 206–238.

Oates, W. & Fischel, W. (2016), 'Are local property taxes regressive, progressive, or what?', *National Tax Journal* **69**(2).

Oates, W. & Schwab, R. (2015), 'The window tax: A case study in excess burden', *Journal of Economic Perspectives* **29**(1).

Olima, W. (2001), The dynamics and implications of sustaining urban spatial segregation in kenya: Experiences from nairobi metropolis. www.begakwabega.com/documenti/dynamics-implications-urban-segregation.pdf.

Olken, B. (2009), 'Corruption perceptions vs. corruption reality', *Journal of Public Economics* **93**, 950–964.

Omar, H. (2017), 'Regularization Of Informal Settlements Should It Be A 'Demand Driven' Or 'Supply Driven': The Case Of Kimara Ward In Ubungo Municipality, Tanzania', *Researchjournali's Journal of Geography* **4**(7).

Parliament of Uganda (2005), 'The Local Governments (Rating) Act'. private interviews.

Payton, S. (2012), 'The Impact of Property Assessment Standards on Property Tax Burden: An Examination of Systematic Bias in a Market Value Versus a Nonmarket Value Assessment Standard', *Public Finance Review* **40**(5), 584–613.

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

Pomeranz, D. & Vila-Belda, J. (2019), 'Taking State-Capacity Research to the Field: Insights from Collaborations with Tax Authorities', *Annual Review of Economics* **11**, 755–781.

Prichard, W., Kamara, A. B. & Meriggi, N. (2020), Freetown just implemented a new tax system that could quintuple revenue. africanarguments.org/2020/05/21/freetown-just-implemented-a-new-tax-system-that-could-quintuple-revenue/.

Programme, U. H. S. (2003), The challenge of slums: Global report on human settlements 2003, Technical report, London and Sterling, VA.

Property24 (2015), Nairobi vacant land listings. www.property24.com.

Rigol, N., Hussam, R. & Roth, B. (2020), Targeting high ability entrepreneurs using community information: Mechanism design in the field. HBS Working Paper 20-082.

Romer, P. (1994), 'New goods, old theory, and the welfare costs of trade restrictions ', *Journal of Development Economics* **43**, 5–38.

Rosen, S. (1974), 'Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition', *Journal of Political Economy* **82**(1), 34–55.

Salon, D. & Gulyani, S. (2010), 'Mobility, poverty, and gender: Travel "choices" of slum residents in nairobi, kenya', *Transport Reviews* **30**(5), 641–57.

Sclar, E. D., G. P. & Carolini, G. (2005), 'The 21st century health challenge of slums and cities', *The Lancet* **365**(9462), 901–3.

Selod, H. & Tobin, L. (2018), The informal city. Working Paper No. 8482, World Bank Policy Research.

Sheuya, S. & Burra, M. (2016), 'Tenure Security, Land Titles and Access to Formal Finance in Upgraded Informal Settlements: The Case of Dar es Salaam, Tanzania', *Current Urban Studies* **4**, 440–460.

Slemrod, J., Rehman, O. U. & Waseem, M. (2020), 'How do taxpayers respond to public disclosure and social recognition programs? evidence from pakistan', *Review of Economics and Statistics* . forthcoming.

Southall, R. (2005), 'The ndungu report: Land and graft in kenya', *Review of African Political Economy* **32**(103), 142–51.

Syagga, P., Mitullah, W. & Karirah-Gitau, S. (2002), Prepared for government of kenya and un-habitat collaborative nairobi slum upgrading initiative, Technical report, Nairobi, Kenya.

Taubenbock, H., Kraff, N. & Wurm, M. (2018), 'The morphology of the arrival city—a global categorization based on literature surveys and remotely sensed data', *Applied Geography* **92**, 150–167.

The Standard (2019), Sonko seeks to raise land rates by 50pc. [Online; accessed 15-July-2020].

UN (2018), World urbanization prospects, Technical report. population.un.org/wup/.

UN Habitat (2011), Ghana housing profile, Technical report, Nairobi, Kenya.

UN Habitat (2016*a*), Slum almanac 2015/2016: Tracking improvement in the lives of slum dwellers, Technical report, Nairobi, Kenya. http://unhabitat.org/slum-almanac-2015-2016/.

UN Habitat (2016*b*), World cities report 2016: Urbanization and development, emerging futures, Technical report, Nairobi, Kenya. unhabitat.org/world-cities-report.

USGS (2005), Shuttle radar topography mission, 1 arc second scene srtm_f03_s002e036. Global Land Cover Facility, University of Maryland, College Park, Maryland.

Weigel, J. (2020), 'The participation dividend of taxation: How citizens in congo engage more with the state when it tries to tax them', *Quarterly Journal of Economics* . forthcoming.

Williams, M. & Klopp, J. (2014), 'Toward open source kenya: Creating and sharing a gis database of nairobi', *Annals of the Association of American Geographers* **104**(1), 114–30.

Wolff, S., Kuch, A. & Chipman, J. (2018), Urban land governance in dar es salaam: Actors, processes and ownership documentation. IGC working paper C-40412-TZA-1.

World Bank (2006), Where is the wealth of nations: Measuring capital for the 21st century, Technical report, Washington DC.

World Bank (2015), 'Nairobi building footprints circa 2003 and plot boundaries circa 2010'. Washington, DC: Multiple Electronic Files.

Yeom, J., Kim, Y. & Lee, J. (2015), 'Hierarchical icp matching to enable interoperability of multisource building gis data', *Advanced Science and Technology Letters* **89**, 55–60.

Zebong, N., Fish, P. & Prichard, W. (2017), Valuation for Property Tax Purposes. ICTD Summary Brief Number 10.