

Building Tax Capacity at Scale: Evidence from Technology Investments in Ghana

James Dzansi (IGC) Anders Jensen (HKS) David Lagakos (UC San Diego) Henry Telli (IGC)

June 2018

1 Introduction

Increased capacity to collection leads to funding of public goods which have high returns, and contributes to the wider state-development process (Besley and Persson, 2010; Schumpeter, 1918). At the heart of any strong tax system lies the ability to accurately observe the activities of economic agents, and the ability to use this information as inputs into an efficient billing, collection, and enforcement process. Constraints on either of these two capacities can have large adverse impacts. First, it can drive a wedge between (potentially well-intended) official, statutory policies and the policies that are effectively implemented in the field (with potential regressive impacts); second, it increases the discretionary power of intermediaries (e.g. local tax collectors in the field); third, it increases the cost of the collection process itself, and can decrease citizen trust and engagement with wider tax system.

These adverse impacts are strongly present in the context of local property taxation. In Ghana, the context for our study, at baseline we find that: 76% of tax payments are made based on presumptive, rather than official, statutory tax schedules;¹ the cost of collection is 55% of collected revenue;² and, citizens report low trust in local tax officials and lack awareness of local tax setting. Appropriately designed information technologies (IT) have the potential to alleviate these capacity constraints. Indeed, they can massively increase the coverage of taxpayer information, while reducing both the fixed cost and marginal cost of acquiring this information; and, they can map this digitized information into an automated billing, collection, and enforcement process. Strong local tax capacity is integral to the wider agenda of decentralization of revenue and spending powers, which the World Bank Development Report 1999/2000 notes as having beneficial impacts on “political stability, public sector performance, equity, and macroeconomic stability.” (World Bank, 1999) In the African context, where the pace of decentralization has stagnated in recent years (Gadenne and Singhal, 2014), the introduction of appropriately designed IT systems has been hailed as a potentially transformative tool to alleviate constraints on local tax capacity (APTI, 2017), with large recent investments in several countries.³ But despite technology investments holding the promise to significantly improve tax and broader state capacities, there is no well-identified evidence on the returns to such investments.

¹Calculated as: $[\text{share properties unassessed}] * [\text{share assessed that are billed}] / (\text{numerator} + \text{share properties assessed})$

²For comparison, the IRS cost of collection is 0.5 percent of revenue; the average developing country’s central tax cost of collection is 2-3% (Jensen and Lagakos).

³Including Burundi, Tanzania, Rwanda, Nigeria (Lagos State), Sierra Leone, Cote d’Ivoire, Senegal, Malawi

The optimistic view surrounding IT solutions has gained momentum from recent technological advances. The first step towards building an effective property tax system is the creation and maintenance of a tax registry based local property maps. Rapid growth of cities and urban sprawl imply that manually entered local property maps are both costly and quickly outdated; in contrast, recently developed geographic information systems can drastically reduce the time and cost required to identify new properties and integrate them into pre-existing digital property maps. The second step is to provide valuations for properties in the tax registry. Often, this is handled by a limited number of central government officials, who have to perform the inspections manually in person. Instead, new computer assisted mass appraisal technologies allow local governments to directly collect the required data and produce the property values. The final step in a property tax system is the billing, collection, and enforcement of tax payments. Owing in part to the absence of digital registries and valuations, this process is also manual, providing intermediaries (local budget officers and local tax collectors) with significant discretion over who to tax and significant bargaining power in the negotiation of formal and informal payments with individual households. IT-based processes, in contrast, allow for tax bills to be generated for all properties, easy(/ier) recording and monitoring of payments, and automated identification and monitoring of tax defaulters. These technologies could significantly improve local tax collection, both by increasing the gross amount collected through greater coverage (share of properties in the registry) and recovery (share of bills that are paid), and by reducing the cost. In addition, the transition away from a manual, potentially discretionary process towards a technology-based, more automated process may also increase public confidence in property tax systems and lead to greater engagement with local governments.

At the same time, there are a number of reasons to be skeptical about the potential for IT-solutions. First, their implementation entails solving a complex mix of technical and logistical challenges, with the concern that the undertaking might fail unless all components are well-implemented (Kremer, 1993). Anecdotal evidence suggests that the impact of IT-systems for the collection and enforcement process is severely limited if not combined with efforts to digitize the property registry. Related, attempts to implement IT systems in the past for tax systems have been plagued by overly complex functionality, and lack of compatibility with local needs and infrastructure (Fish and Prichard, 2017). Second, resistance by those whose vested economic or political interests are threatened (such as local tax collectors and officials) may block the introduction of the new technologies or limit their effectiveness (Krusell and Rios-Rull, 1996; Acemoglu and Robinson, 2000). Third, even assuming positive revenue impacts, cost-effectiveness is unclear. Taken together, there is therefore limited robust evidence to support either the enthusiasts or the skeptics of technology-based tax collection processes.

In this project, we contribute toward filling this gap, by presenting evidence from a large-scale experimental evaluation of the revenue collection and broader state-building impacts of improved technology in the local tax setting of Ghana. In collaboration with the Ministry of Finance and the Ministry of Local Government, we randomized the order in 64 which local governments were offered a domestically-developed IT system to collect property and businesses taxes. The randomization was done at the level of entire local governments, covering 7.25 million people. Local governments in Ghana enjoy full discretion over the spending of local taxes, and property and business taxes together constitute over 75% of locally collected revenue. At the same time, there is wide-spread consensus among government ac-

tors that local taxation is severely constrained and carries untapped potential (Government of Ghana, 2014). Randomizing at the full scale of local governments lets us overcome one concern often voiced about randomized trials in developing countries: that studying small-scale pilots may not provide accurate forecasts of performance when governments must implement the same intervention at a larger organizational scale (Muralidharan and Niehaus, 2017). At the same time, our experimental sample covers local governments which differ significantly both in baseline 'administrative capacity' and in economic structure (including urbanization rates), allowing for a rich heterogeneity analysis. Both these features increase confidence that our results can speak to potential impacts of similar technologies in other settings.

The IT technology offered under the intervention was locally developed by a private Ghanaian firm, in consultation with central and local governments. We argue that this technology can be effectively implemented and have sustainable impacts for several reasons. First, the technology features IT systems to support all three critical components of the tax collection process identified above: creation of a registry; digitized valuations; billing, payment, and enforcement. Indeed, this technology was conceived in the context of a prior IT system, developed jointly by the central government and donors, which exclusively enhanced the third component and failed to be adopted. Less than 15% of districts currently have systems where the three components are supported by any technology. In addition, inadequate databases for property registration, valuations, and collection is overwhelmingly cited as the most important constraint on revenue collection by the local chief executives, administrators, and legislators. At the same time, low incentives to collect revenue and lack of qualified staff are cited as the two least important constraints. Second, the technology is tailored to the local setting, aiming to be compatible both with the local IT staff and infrastructure capacity, and with the legal tax structure which varies significantly across districts. Third, although the technology is offered for free in the context of our evaluation, the private firm will eventually seek to recover its investment through commission rates. This structure implies there are clear incentives for the developer to sustain implementation over the long run, through continued technical assistance. These two features contrast with previous technology adoption settings, where anecdotal evidence suggests that international IT solutions are often incompatible with local needs, while donor IT solutions often fail to sustain implementation of IT systems beyond the short-run due to shifts in policy priorities and a collapse in technical support to local governments' tax administrations (Prichard, 2014). Fourth, in the African context where central governments often exert significant *de facto* pressure on local governments' policies, the intervention benefits from the buy-in of relevant central actors: the Ministry of Finance (which oversees the public financial management of districts); the Ministry of Local Government (which oversees the chief executives in the districts); and, the Local Government Services (which oversees all administrative positions in the districts).

At baseline, less than 15% of local governments have technology-assistance for the three core components of the tax collection process. We expect significant, but not complete take-up after one and a half years of rollout. Conversion rates in other settings are slow. In the US, it took a decade to achieve wide-spread adoption of computer mass appraisal systems (Youngman, 2016); in Brazil, less than 9 percent of municipalities adopted subsidized technology enhancements over a 12 year period (Gadenne, 2016). On the other hand, the likely inability to reach 100% transition from manual to technology-based collection reflects the non-trivial logistical, administrative and political challenges of adopting

overhauling IT solutions in Africa (Fish and Prichard, 2014). We therefore focus throughout the paper on intent-to-treat analysis. These estimates yield the relevant policy-parameter, as they represent the average return to as-is implementation when local newly developed IT solutions become available, and are net of the political and logistical challenges that constrain adoption and effective implementation of such new technologies in practice.

2 Revenue Performance

We organize our analysis around three main dimensions: revenue performance, cost effectiveness, and citizen engagement and expenditure. Beginning with revenue performance in the first section, we begin by simply studying the impacts on gross aggregate revenue collection. We have access to monthly administrative data that is centrally collected by the Ministry of Finance. The administrative nature of the data should ensure high quality of the measured outcomes. The 'high frequency' of the data will also allow us to study the time-series impacts on revenue collection: first, through possible contrasts between short-run gains and 'medium-run' sustained impacts; second, through changes in the month-to-month variance. With regards to the latter, at baseline, the funds received by the average local government are extremely volatile, with a coefficient of variation of 0.65. Uncertainty regarding the amount of funding is thought to be associated with lower quality delivery of public goods (including non-completion which leads to fully wasted public expenditure: see Williams (2017)).

Next, we unpack the aggregate gross revenue impacts by studying two 'mechanisms'.⁴ First, we measure the impact of technology on the coverage ratio, defined as the percentage of properties in a district which are subject to taxes and are sent a bill. A property is subject to tax either if it is formally assessed and therefore falls under the official, statutory schedule; or, if it is unassessed, in which case it falls on the presumptive schedule. At baseline, the coverage ratio was 48 percent: 18 percent of properties were taxed on the official schedule, and 30 percent on the presumptive schedule. We expect the technology-based GIS systems to significantly expand the property registry, and therefore the coverage ratio to increase. The second 'mechanism' is the recovery rate, defined as the share of bills sent that are paid before the end of the fiscal year. At baseline, the average recovery rate is 70 percent.⁵ The automation of the billing, collection, and enforcement processes is conjectured to increase the recovery rate, both on pre-existing properties and on newly registered properties.

We also examine the distributional impacts of enhanced technology. We first study the distributional impacts within districts, between schedules. Indeed, while the presumptive tax schedule is based on limited property characteristics, the formal, official rates are based on a large number of observable inputs which allow for more progressive tax schedules. Through the digitized registry and the computer assisted appraisals, we expect the share of properties subject to formal, statutory tax schedules to increase, in relation to the overall number of taxed properties. This formalisation of taxes through the transition from presumptive to official schedules will be associated with increased 'de jure' progressivity.

⁴We can summarize the different effects in the following 'net revenue' formula:

$$\text{Net Revenue} = [\text{Value}] \times [\text{Coverage}] \times [\text{Tax Rate}] \times [\text{Recovery}] - \text{Cost}$$

⁵The seemingly high recovery rate has also been found in other property tax contexts (Khan, Khwaja, and Olken, 2016).

Second, we study the distributional impacts within districts and schedules, between households. In particular, we investigate whether the impact of technology on de facto, effective taxation is larger for poorer than for richer households, when proxied for by education level or occupation status. A priori, the individual level distributional impacts of enhanced technology are not clear. If the main impact of technology is through increased coverage, and the newly registered properties are less valuable than the pre-existing ones, then technology may have regressive impacts on effective taxes paid along the income distribution. On the other hand, if the main impact is through increased recovery conditional on coverage, it is plausible that richer households are the primary target of the increase in effective taxation. Finally, we study the distributional impacts between districts. In particular, we investigate whether treatment effects vary significantly as a function of district-level baseline characteristics, such as administrative 'capacity' and per capita income.

3 Costs and Leakage

Cost-effectiveness is determined by gross revenue impacts, on the one hand, and by potential impacts on 'official' costs and leakage on the other hand. We investigate cost-effectiveness in the following section. We measure official costs relative to gross revenue collected. By far, the most important cost-component is the salary paid out to field collectors. Prior to treatment, this cost represents 55% of gross revenue, on average. This high cost-ratio is consistent with a baseline setting in which field collectors derive rents from having more local information about current and potential taxpayers than the district's central office. By allowing the central officers to visualize electronic, up-to-date maps of properties in their district and by digitizing the inputs required for valuation, the intervention will significantly reduce the asymmetry of information, with a conjectured decrease in the cost-ratio. Total cost, however, may go up or down. It may go down if, through efficiency enhancement triggered by automation of the collection process, the intervention may reduce reducing the effective number of hours worked required (either decrease the hours worked per collector, or decrease the number of hired collectors). On the other hand, if technology and labor are complements in the 'revenue production function', then the intervention may increase the total number of tax collectors. Taken together, the increases in coverage and recovery combined with the decreases in costs can lead to significant impacts on net revenue (gross revenue minus costs): supposing each of the components changed by 20%, that would imply a 74% increase in net revenue.

Turning to leakage, there are two main margins of change: 'in-field' leakage due to corrupt behavior of local tax collectors (e.g. collecting 100\$ but only remitting 70% to the central office); and, 'central office' leakage due to corrupt behavior of local administrators (e.g. receiving 70\$ but only recording 40\$ on the official revenue accounts). We rely on two imperfect strategies to investigate whether technology reduces leakage along these two margins. First, in the case where aggregate payment rates based on resident surveys did not change (increased) in treated districts but aggregate gross collected revenue increased (did not change), we can infer a reduction (increase) in leakage of funds between the taxpayer and local government coffers. Using this method, however, does not allow us to distinguish between 'in-field' and 'central' leakage. Using this method, however, does not allow us to distinguish between 'in-field' and 'central' leakage. The second strategy is therefore to construct proxies for these two margins. We collect survey data asking officials (executive, administrator, legislator) whether leakage

of funds specifically in the field is an important constraint on revenue collection. In addition, we make use of the irregularities uncovered by the Auditor General in their periodic audits of local governments. The audits are comprehensive and made publicly available, and several audit outcomes relate directly to 'in-field' versus 'central' leakage. Per example, the audits report any value books which are not accounted for in the final bank deposits made by the budget or finance officer. We have digitized the irregularities up until 2016 and will continue digitizing the Auditor General reports as they become available.

4 Wider 'State-Building' Impacts: Citizens, Leaders, Expenditure

In the final section, we study the impacts of state capacity investments on citizen engagement, leadership preferences, and expenditure. Citizen behavior could be impacted through two main channels: first, the increase in direct taxation could lead citizens to engage, according to the "tax-participation" hypothesis (Weigel, 2018); second, in a low-capacity setting, the investment in technology can be interpreted by citizens as a signal about the state's increased overall capacities (Coate and Morris, 1995). We provide a reduced-form investigation, by studying the impact of the treatment on: citizens': level of interaction with tax officials (in absolute, and relative to interaction with other government officials); competency assessment of the local tax administration (in absolute, and relative to other government departments); level of awareness about the local tax system and local public goods projects; and, level of trust in local officials (in absolute, and relative to neighbours or people met for the first time). At baseline, the level of engagement in the typical district is low: 59% of citizens have never engaged with the local tax authority; 92 percent have never heard of the 'fee fixing resolution', the basic procedure through which property rates are established, and 0.25% have ever attended the public meetings where citizens can debate the resolutions with local government; citizens trust local officials as much as complete strangers and 30% less than their neighbours; 13% of citizens are aware of any public goods projects in a particular category (transport infrastructure; education; public hygiene; waste management; health care). We also test the impacts of state capacity investments on leadership's aversion to tax and discount factor, two preference-parameters which have been identified in the theoretical literature as key constraints to investments in a state's tax capacity (Besley and Persson, 2010). In particular, we study the treatment impacts on the willingness to pay for central transfers (over direct taxes), and on the discount factor (revenue today versus more revenue in the future) of each district's top executive, administrator, and legislator.⁶ Interestingly, at baseline, the discount factor seems to be more present than aversion to tax: 12 percent of leaders would give up revenue in the present period to receive it through transfers rather than through taxes; 26 percent of leaders would choose revenue in the current period over a larger amount of revenue in the future period.

We finish this section by testing the impacts on expenditure. The main hypothesis is that as a result of the signalling of government capability, the increase in direct taxation and citizen engagement, expenditure will be shifted towards types of public goods that are more (directly) beneficial to citizens. We test this hypothesis in two complementary ways. First, we study whether the share of total

⁶In a non-experimental design, we can also investigate the impacts of these leadership preferences on the adoption of technologies and on the effectiveness of the implementation: the first, by correlating decision to adopt with the preference measures; the second, by interacting treatment status with the preference measures when studying revenue outcomes.

expenditure devoted to recurrent expenditures (maintenance of administrative facilities, purchase of officials' transport, and salary compensation) decrease and the share devoted to delivery of goods and services increases. Second, and perhaps as the most direct test, we leverage information collected on both citizens' and leaders' ranking of priority spending across categories of goods and services delivered. Interestingly, at baseline, there is a systematic divergence between the goods categories that citizens and leaders consider to have highest priority: in particular, the top 2 categories chosen by a majority of residents (road building and water provision) are absent from the top three choices of the top chief executive, chief administrator, and chief legislator. We study both whether the gap narrows between leaders' and citizens' stated public goods preferences, and whether the actual share of expenditure on goods on services preferred by citizens increases in treatment districts. Finally, to understand whether quality of public goods expenditure increases, we investigate whether the treatment increases the completion rate of infrastructure projects. A priori, the direction of the impact is not clear. On the one hand, increased availability of revenues should relax funding constraints and lead to more projects being funded until completion. On the other hand, increased citizen engagement may actually exacerbate the importance and uncertainty of local political bargains which have been associated with non-completion (Williams, 2017). While the investigation in this section remains reduced-form, the analysis importantly causally investigates the additional mechanisms, above and beyond direct impacts on revenue performance, through which investments in capacity to tax contribute to the building of a capable state and the wider development process (Schumpeter, 1918).