Tax Productivity and Economic Development: A Quantitative Macroeconomic Analysis

Radhika Goyal
IMF

Anders Jensen
Harvard University and NBER

David Lagakos
Boston University and NBER

Abdoulaye Ndiaye
New York University

Yuntian Zhao
New York University

March 2026

Abstract

This paper examines the role of government capacity to collect taxes, or *tax productivity*, in driving economic development. We first draw on new evidence from 100 countries to study how tax collection outcomes relate to tax administration inputs across the world income distribution. We document that tax collections relative to tax administration expenditures rise sharply with GDP per capita: the richest countries collect approximately ten times as much tax revenue per dollar spent as the poorest. We interpret this finding in a quantitative Ramsey model with a tax administration sector and public capital in the production function. We calibrate the model to match our cross-country evidence and other salient empirical patterns, including the increasing ratio of taxes collected to GDP in development. We then simulate the long-run aggregate effects of raising tax productivity. Quantitatively, our model predicts that increases in tax productivity in low-income economies can raise GDP per capita substantially more than suggested by Hulten's theorem.

Keywords: tax productivity, economic development, Ramsey taxation, macroeconomics

Email: rgoyal3@imf.org, Anders_Jensen@hks.harvard.edu, lagakos@bu.edu, an1490@stern.nyu.edu, and yz11552@stern.nyu.edu. For helpful comments we thank Augustin Bergeron, Corina Boar, Gustavo Ventura, Xavier Gabaix, Emilien Gouin-Bonenfant, Roozbeh Hosseini, Jordan Norris, Steven Pennings, Joachim Voth, James Robinson, Florian Scheuer, Tao Zha, and seminar/conference participants at the Amherst College, the Atlanta Fed, Copenhagen, Nottingham, NYU, the 2025 NBER Summer Institute, USC, the World Bank, and ASU. All potential errors are our own.

1 Introduction

One of the most prominent theories of long-run economic development places state capacity — and, in particular, the ability of governments to raise tax revenue — at the heart of the development process (see, e.g., Besley and Persson, 2011). Historical work by Mayshar, Moav, and Pascali (2022) shows that the most successful early states, such as Egypt, were those that were better able to raise revenues to fund public works and military ventures. Tilly (1990), Dincecco and Katz (2016), and Gennaioli and Voth (2015) provide evidence that the constant warfare of European states in the 16th and 17th centuries led them to build up state capacity, which may ultimately have sown the seeds of their future development. Recent empirical work by D'Arcy, Nistotskaya, and Olsson (2024) ties the historical creation of public land records—the foundation of property taxation—to long-run growth.

The link between a government's ability to collect taxes and its long-run level of development has received less attention in the macroeconomics literature. The foundational papers of Barro (1990), Baxter and King (1993), and Jones, Manuelli, and Rossi (1993) incorporate public capital stocks into the production function but abstract from any costs of collecting the taxes needed to finance public investments. Tax administration costs are similarly abstracted away in nearly all models of optimal taxation, whether focusing on infinitely lived agents (Judd, 1985; Chamley, 1986), overlapping generations (Conesa, Kitao, and Krueger, 2009), or Aiyagari-type heterogeneous agents (Auclert, Cai, Rognlie, and Straub, 2024). A notable exception is Besley, Ilzetzki, and Persson (2013), which incorporates the political economy of tax capacity into a growth model. They show that more severe fiscal constraints can lead an economy to reach a less desirable steady state, although they do not pursue the quantitative implications of their model.

This paper aims to assess the quantitative importance of government capacity to collect taxes, which we refer to as *tax productivity*, for economic development. To this end, we first collect data on tax collections relative to tax administration expenditures in a large set of countries and measure how this ratio varies with development, both across countries and within countries over the long run. We then use these data to help discipline a multi-sector Ramsey model of taxation, capital accumulation, and economic growth. The key novel feature of the model is a tax administration sector that can collect tax revenues from households, but only by using up labor inputs. Public goods are complementary to private capital in the production function and can only be provided through tax revenues. The model allows for exogenous cross-country differences in both tax collection productivity and final goods production. We use the model to simulate the long-run aggregate effects of raising tax productivity for countries of low income levels, where we suspect growth is most constrained by an inability to mobilize tax revenues.

We draw on novel data hand-collected from tax authorities —such as the U.S. Internal Revenue Service and Ghana's Revenue Authority—across 100 countries and over time, yielding 625 country-year observations. Our two main variables of interest are the total tax revenues and the administrative cost of collecting those revenues. Our outcome of interest is simply the ratio of total tax revenues to collection costs, representing the average amount collected per dollar spent on collection. This unit-free measure is intuitive and comparable across countries and over time within countries while being closely related to, yet distinct from, tax productivity.

Our main empirical finding is that tax collections per unit of expenditure are strongly increasing in GDP per capita. In terms of magnitudes, the poorest countries in the world collect about $30 for every dollar they spend collecting taxes; this number rises to $250-270 in the world's richest countries. We show that this pattern is not an artifact of differences in the scope of the revenue authorities' activities. For example, collections per expenditure remain strongly correlated with development even when we exclude tax administrations with limited scope (those not responsible for VAT or sales tax) or expanded scope (those responsible for lotteries or social security transfers). Exploiting the panel dimension of our data, we also examine how collections per expenditure vary with GDP per capita within countries over time. We find that countries tend to increase their collections per expenditure as they grow, with an elasticity that is even stronger than in the cross-section. This within-country pattern again holds when restricting the analysis to tax administrations with comparable duties. The patterns also hold for individual sub-national tax authorities over time, which we show by leveraging newly constructed historical panel data on taxes collected and expenses for each (independent) state tax administration in the United States and India.

We interpret these facts using a version of the neoclassical growth model with a tax administration sector that can raise tax revenues but requires labor inputs in the process. Production combines labor, private capital, and public capital to produce a single output good for consumption or investment, following Barro (1990) and Baxter and King (1993). Public capital can only be provided through tax revenue, and taxes can only be raised through costly efforts by the tax administration. The government solves a Ramsey problem, choosing optimal tax collection efforts to maximize the utility of an infinitely lived representative household, taking both its tax collection productivity and the final goods productivity as exogenous.

The model incorporates three additional features that make it suitable for analyzing low- and middle-income countries. First, we include a dual economy structure with a modern sector that uses both types of capital but faces income taxation and a less productive traditional sector that can avoid taxation entirely. This provides an additional constraint on developing country governments: if taxation becomes too heavy, households can shift labor from the modern to the traditional sector. Second, we allow for decreasing returns to scale in the tax administration sector, reflecting the fact

that further expansions of tax collection efforts require moving from easy-to-collect sources, such as taxes on international trade, to more challenging ones, like property taxation and income taxes on the self-employed (e.g. Jensen, 2022). Finally, we allow some fraction of tax collections to be lost as waste – a reasonable assumption for any government but particularly relevant for developing nations where corruption is widely viewed as more prevalent.

We analytically characterize the model, focusing first on a simple static version. One key lesson is that the degree of returns to scale in the tax administration matters substantially. In the case of constant returns to scale, and considering our cross-country evidence, the governments of developing countries could always collect around 30 dollars in taxes for every 1 dollar spent on collections. In this case, tax collection per se cannot be a serious constraint on the provision and development of public goods. A second insight is that while government waste is clearly a negative force, allowing for it in the model does not change the model's conclusions about the importance of raising tax productivity. The reason is that government waste affects the economy both before and after improvements in tax productivity, implying similar relative effects of tax productivity regardless of the presence of waste.

We calibrate the model to match key features of developing countries, including their low ratio of tax collections to expenditure, large traditional sectors, and low average tax rates. One key parameter choice is the elasticity of output to public capital, and we select a value of just under 0.2 based on evidence from a meta-analysis. Another key choice is the extent of decreasing returns to scale in tax collection, and we pick this to match evidence on returns to expanding tax collection efforts by the U.S. and several developing countries.

We then simulate the long-run aggregate effects of raising tax collection productivity. The model predicts that a 50 percent increase in tax productivity raises GDP per capita by about 5 percent. As a frame of reference, Hulten's theorem suggests gains of less than 0.1 percent of GDP, given that the tax collection sector employs less than 0.2 percent of the labor force. The model's GDP gains stem from public and private capital increases of 57 percent and 9 percent, respectively, as well as modest shifts of labor from the traditional to the modern sector. Tax collections relative to expenditures rise from around 30 to around 150, comparable to middle-income country levels.

We conclude that improvements in tax collection productivity in the world's poorest economies likely lead to significant positive long-run gains in GDP. The magnitudes are not large enough to suggest that low-income countries can become middle- or high-income nations solely by improving their tax collection sectors. However, the GDP gains are still substantially larger than suggested by Hulten's theorem, which would imply negligible aggregate gains from fixing a sector that employs only one out of a thousand workers on average. In ongoing work, we are calibrating the model to match middle- and high-income countries, both to assess the model's predictions for higher-

income countries against our novel data and other cross-country evidence and to quantify the gains from tax collection improvements in these settings.

Related Literature. Our paper contributes to the macroeconomic literature on taxation by analyzing the effects of a resource cost of collecting revenues. Straub and Werning (2020) overturns the famous result, inspired by Chamley and Judd, that in Ramsey models the optimal capital tax rate should be zero. We take a simple route and assume only taxation of labor and positive tax rates are always optimal in our setting, since public capital is complementary to private capital. Our study also abstracts from insurance considerations as in Mirrlees problems (Kapicka, 2013; Farhi and Werning, 2013; Golosov, Troshkin, and Tsyvinski, 2016; Boerma and McGrattan, 2020). Our paper is related to the study by Bhandari, Evans, Yao, and McGrattan (2024) incorporates tax evasion into a macro model calibrated to match U.S. evidence on the effects of tax audits. More generally, our work relates to the public finance literature on tax evasion (see Slemrod, 2007, and the references therein). Households in our model can evade taxes by working in the traditional sector, though only by accepting a drop in productivity. Similarly, governments can raise effective tax rates in our setup by spending more resources on tax administration, which is supposed to represent greater efforts to enforce the tax code.

Finally, our paper adds to the large literature on public finance and development focused on tax capacity and its determinants (Gordon and Li, 2009). Paths to increasing tax capacity emphasized by this literature include incentive pay for tax collectors (Khan, Khwaja, and Olken, 2015), collaboration with traditional leaders, (Balan, Bergeron, Tourek, and Weigel, 2022), information technology (Hjort and Tian, 2024; Okunogbe and Tourek, 2024), including using geo-referenced databases of taxpayers (Dzansi, Jensen, Lagakos, and Telli, 2025), and self-enforcing tax systems (Naritomi, 2019). Much of the recent literature has focused on the taxation of properties, which holds vast promise but is thought to be vastly under-collected in the developing world (Brockmeyer, Estefan, Arras, and Suárez Serrato, 2023). The same is true of taxes on labor income, particularly given the challenges in collecting from the self-employed (Jensen, 2022). Our study contributes by modeling and quantifying the long-run general equilibrium consequences of improving tax capacity, which is new to this literature.

2 Data and Empirical Findings

We begin by presenting the data we use in our analysis and defining the ratio of taxes collected to expenditures. We then investigate how this measure of collection per dollar of expenses varies with income per capita across space and time - between countries today, within countries over the

long run, and at both the national and sub-national levels.

2.1 Data

In this subsection, we describe the data collection exercise undertaken to measure the ratio of taxes collected to administrative expenses on collection in a large sample of countries, years, and levels of government. As the theory will make clear (Section 3 onward), this measure is closely related to, but also distinct from, tax productivity. One way to think about this measure, which will be helpful in the model, is that it captures the average (*not* the marginal) tax collection per dollar of administrative expenditure for a given tax authority at a point in time. It presents the advantage of being unit-free and comparable both across countries at a point in time and within countries over time.

We aim to collect data that allows us to construct the ratio of taxes collected to expenditure on tax collection. The numerator in this ratio captures taxes collected at the national level. The taxes that are typically collected at this level are on personal and corporate income, property, wealth and inheritance, and consumption of goods and services. Importantly, we include data on taxes collected at the border. We focus on actual taxes collected rather than projected or estimated taxes. Finally, whenever the data permits, we ensure that non-tax revenues are excluded.

The denominator in this ratio seeks to measure the annual expenditure of the national tax authority. More specifically, we seek to measure the sum of operating expenditures and capital expenditures. In the first category, by far the most important item is salary; other items include training and miscellaneous administrative expenses. The second category includes expenses that go towards purchasing and maintaining buildings and land, as well as other types of infrastructure. Expenses in information and communication technology (ICT) can be included in the first or second category, depending on whether they are long-term investments (e.g., purchase of a new server) or day-to-day expenses (e.g., software maintenance). The breakdown of collection expenses by type, which exists on an incomplete basis in our sample, reveals that by far the most important cost categories are salary and ICT. Whenever the data permits, we include as many of these cost categories as possible. In all settings, we are confident that salaries are included; in many settings, though not all, the cost measure will be more comprehensive and include capital expenses.

There are many similarities in tax authorities across countries, and this definition of collection cost allows us to capture most of them in a manner that is meaningfully comparable across borders and over time. Notwithstanding, there are limits to this measure. First, there are differences across countries in the range and nature of revenues administered. Here, the most significant distinction is probably between taxes administered by the customs authority (tariffs and other border taxes)

versus taxes administered by the domestic tax authority. There is variation across countries, and within countries over time, in the roles assigned to each institution; importantly, our measure seeks to capture the combined cost of collecting taxes domestically and at the border – such that differences across space and time in roles assigned to different national institutions should not affect our ratio. Another example is that some countries administer direct taxes at the national level, but some indirect taxes at the sub-national level (e.g., the collection of sales taxes is devolved to states in both India and the US). Second, some tax authorities are responsible for carrying out activities that are not directly related to taxation. For example, some tax authorities are also responsible for collecting social security contributions and disbursing transfers. Some, though not many, tax authorities are responsible for hosting and maintaining the population registry. We investigate the importance of these variations in institutional design across tax authorities in our regressions. Third, there are naturally other factors that influence this ratio and that are not related to tax productivity: for example, macroeconomic changes could, at least temporarily, create changes in the taxable base without corresponding adjustments to collection expenditures.

Cross-country data: Sources We collect data from several sources in a cross-section of countries. The most important source is the International Survey on Revenue Administration (ISORA), which is administered collaboratively by multiple international agencies, including the OECD, the IMF, and the Inter-American Center of Tax Administrations. Whenever available, we use data on both the numerator and denominator from this source. We complement this with data from individual countries based on the yearly government budget and national reports from the tax authority and national statistics office. For each country, we seek to collect multiple years of data between 2014 and 2017; we report the yearly average in that range to smooth out fluctuations that occur in the ratio due to abnormal levels of expenses or collection in any given year.

Our resulting cross-country sample contains data from 100 countries. The countries span a broad range of development levels: from lower income countries including Malawi, Kenya, Tanzania, Nigeria, Ghana, Bangladesh, Pakistan and India, to middle income countries including Indonesia, Thailand, China, South Africa, Brazil, Colombia, Mexico and Turkey, to high-income countries including South Korea, Russia, UK, Germany, France, Scandinavian countries and the United States. Combined, our sample of 100 countries in the cross-section accounts for approximately 77 percent of the world population and 67 percent of world GDP.

Historical country panel data: Sources For a subset of countries, historical data exists that allows us to construct measures of tax productivity over the long run within countries. We follow the same guidelines for data collection as for the cross-section of countries.

Our longest time series comes from the United States, where we can measure collection per dollar of administrative expenses back to 1867. We collect data separately for domestic taxes and border taxes. In the former case, we rely on annual reports published on a yearly basis by the Commissioner of the Internal Revenue Service (IRS). In the latter case, we rely on the annual reports of the Treasury on the state of the finances and data from the US Customs Service and the US International Trade Commission. In the historical years, we digitize the relevant data and harmonize it with the pre-existing digital sources for years closer to the present time. In the US case, we are helped by the fact that the underlying sources are often published in a highly consistent manner across a large range of years. We measure tax collections as gross collections before refunds are issued, but after including penalties and interest paid towards outstanding taxpayer arrears. The expenditure measure is comprehensive and includes all expenditures associated with the process of collecting internal revenue. Throughout the entire period studied, personnel costs (including for central officers and field agents) have been an important cost-component. Over time, other cost components have become gradually more important, including physical infrastructure and ICT.

We collect historical data for a sample of other countries around the world. The resulting sample consists of an unbalanced panel of countries: Australia (1972-2017); Bangladesh (1964-2015); Botswana (1971-2005); Canada (1924-2008); Denmark (1959-2006); India (1946-2016); Kenya (1962-1995); Malawi (1964-1996); Pakistan (1949-2005); South Korea (1987-2016); Sierra Leone (1964-1996); South Africa (1960-2007); Thailand (1959-2017); United Kingdom (1993-2017); United States (1877-2016); Uruguay (1980-2017). The median number of years per country is 52, which constitutes a meaningfully long time series.

For the historical panel, we rely on sources that are specific to each country. In most cases, constructing the series for a country required us to scan, digitize, and clean historical publications that existed on paper.[1] For example, in countries including Denmark, Bangladesh, Kenya, and Thailand, we digitized the country's entire series of annual reports produced by the national statistics office or the ministries of planning and economic development. In countries including Canada, India, and Pakistan, we digitized the country's full set of annual government finance account publications.

Historical sub-national panel data: Sources Up to this point, we have focused on collecting measures of tax collection and expenses at the national level for the national tax authority. However, many countries around the world implement a decentralized taxation system, where the mandate to design, implement, collect, and enforce certain taxes is devolved to the sub-national level. In these cases, the sub-national governments are also in charge of establishing their own tax

[1] Most of these physical publications were available in the Lamont Library at Harvard University.

authority, and, in principle, the concept of tax collection per dollar of expense applies at this level as well.

We create long-run panel data on tax productivity at the sub-national level for Indian and US states. We focus on the states because they are in charge of collection for the sales tax - a substantial source of taxation that is collected by the national tax authority in most other countries.[2] In the Indian case, we rely primarily on the 'Combined Finance and Revenue Accounts of the Central and State Governments' series published by the Comptroller and Auditor-General, which we digitize and harmonize for our purposes. In the US case, an important source of data was the historical series 'Book of the States' that is published annually by the Council of State Governments.

The resulting unbalanced panel dataset at the sub-national level contains measures of tax collection and expenses for 28 states in India between 1950 and 2017 and for 50 states in the US between 1993 and 2016.

2.2 Results: Tax collections per expenditure and development

We begin by investigating the relationship between the average collection per expenditure and development in the cross-section. Figure 1 shows the association across countries between the (log) of GDP per capita in USD and the log of the ratio of taxes collected to collection expenses. There is a clear positive association whereby tax authorities in countries with higher GDP per capita have a higher average collection per dollar of tax expenditure.

The magnitude of the differences in collections per dollar of expenditure between the world's poorest and richest countries is substantial. Countries like Burundi, Liberia, and Malawi collect between \$20 and \$30 for every dollar they spend collecting taxes. The world's highest values are primarily in high-income countries such as Sweden and the United States, which collect between \$250 and \$270 in taxes for every dollar spent. China, which is situated around the middle of the GDP per capita distribution, performs remarkably well (\$283) for its level of development. Other middle-income countries typically collect in the range between \$80 for every dollar spent, as in Indonesia, and \$150, as in Brazil.[3]

Table 1 investigates extensions and robustness checks for the pattern uncovered in the cross-section. Column (1) reports the slope coefficient (0.248) that corresponds to the full sample in

[2]Lower levels of sub-national government, such as the municipality or county, are also often in charge of implementing specific taxes (most importantly property tax) and implement their own tax authority to do so. Collecting measures of tax productivity at this level would be interesting but is also practically challenging since most countries do not systematically gather and centralize the required data from the (sometimes large) set of municipalities.

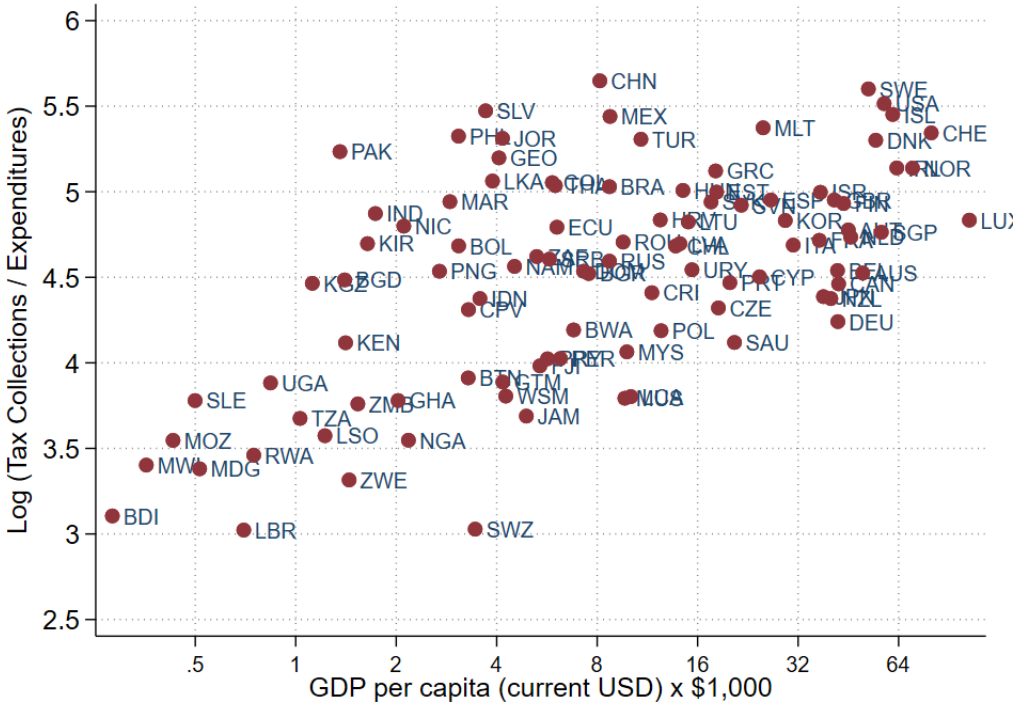
[3]The relationship generally appears to hold within regions of the world, with stronger slopes in Sub-Saharan Africa and Europe and Central Asia, and flatter slopes in South Asia, and the Middle East and North Africa.

Table 1: Cross-Section of Countries: Tax Collections per Expenditure and Income Per Capita

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Outcome	Log T/X	Log T/X	Log T/X	Log T/X	Log T/X	Log T/X	Log T/X
Log GDPpc	0.248***	0.654***	0.225***	0.221***	0.249***	0.260***	0.218***
	(0.0389)	(0.142)	(0.061)	(0.0586)	(0.0395)	(0.0398)	(0.0484)
N	100	35	65	99	97	90	73
R^2	0.293	0.390	0.174	0.321	0.296	0.326	0.222
Modif. to column (1)		Only low, lower-middle inc. countries	Only upper-middle, high inc. countries	Cont. for PIT, VAT, CIT, tariff	Excl. if no VAT or sales tax	Excl. if no customs	Excl. if include soc. sec.

Notes: This table provides regressions of the (log) ratio of taxes collected to expenses on collection on log GDP per capita. The sample is a cross-section of 100 countries between 2014 and 2017 (averaged at the country level when multiple data points exist within a country). Column 1 reports the basic OLS, which corresponds to Figure 1. Column 2 restricts the sample to low and lower-middle-income countries (with GDP per capita below $4,200). Column 3 restricts the sample to upper-middle and high-income countries (with GDP per capita above $4,200). Column 4 includes as controls the statutory tax rates for PIT, CIT, VAT, and tariff (Figure **??**). Column 5 excludes countries that do not have a VAT or general sales tax at the federal/national level. Column 6 excludes countries where the ratio of collection to expenses does not include customs. Column 7 excludes countries in which the tax authority is responsible for collecting and disbursing social security. Column 8 excludes countries in which the tax authority hosts (maintains and updates) the population registry. Robust standard errors are reported in parentheses. $*p < 0.10 * *p < 0.05 * * * p < 0.01$. The figure plots the (log) ratio of taxes collected relative to expenses on tax collection against GDP per capita in a sample of 100 countries.

Figure 1: Tax Collections Per Expenditure in the Cross-Section of Countries



Notes: This figure plots the (log) ratio of taxes collected relative to expenses on tax collection against GDP per capita in a sample of 100 countries. Data is collected for each country between 2014 and 2017, and the average is taken over the years if multiple data points exist per country.

Figure 1. Columns (2) and (3) estimate the relationship separately in low and lower-middle-income countries and upper-middle and high-income countries, respectively. The slope between tax collection to expenses and (log) GDP per capita is larger at lower levels of development (0.654 versus 0.225), but it remains significant in both samples.

In column (4), we control for the countries' statutory tax rates. We include as controls the top personal income tax rate, the main corporate income tax (CIT) and value-added tax (VAT) rates, and the average tariff rate.[4] This specification investigates whether the variation in collection per $ of expense is driven by variation in the amount that can be collected. It is based on the notion that countries at higher levels of development, due in part to stronger enforcement, can levy higher tax rates. Figure **??** shows that there is some positive relation between tax rates and GDP per capita in our sample, but it is limited - in particular, it applies only PIT, with flat rates for both CIT and

[4]The PIT, CIT, and VAT rates are collected from the global tax rate databases produced by international consultancies, including KPMG. The tariff rate is calculated as the unweighted average across all rates and comes from the World Bank's Development Indicators. In turn, the World Bank's data comes from the UN's Trade Analysis and Information Systems (TRAINS) database.

VAT and, as expected, decreasing rates for tariff. Moreover, the positive association between the PIT rate and development only exists for upper-middle and high-income countries, while the slope between tax collection per $ expense and GDP per capita is stronger in low and lower-middle-income countries. Consistent with these observations, column (3) shows that controlling for the statutory rates does not alter the slope coefficient substantially (0.221).

Columns (5) through (8) remove from the sample countries whose tax authorities differ in significant ways. In column (5), we remove countries where the national tax authority is not responsible for collecting the VAT or general sales tax. This leaves the main coefficient virtually unchanged (0.249). In column (6), we remove countries where the (domestic) tax authority is not also in charge of collecting border taxes – this decreases the sample to 90 countries and increases the slope coefficient slightly (0.260). In column (7), we remove countries where the tax authority is also in charge of collecting social security contributions - the reduction in sample size is significant (73), but the slope coefficient remains stable (0.218). [5]

The sample restrictions of Table 1 suggest that the positive association between the average collection per expenditure and development is a general feature in the data that holds across a variety of institutional settings and at different levels of GDP per capita.

We now turn to investigating patterns within countries over time. Figure 2 displays the evolution of taxes collected per dollar of collection expense at the federal level in the US between 1867 and 2017. To maintain consistency with the cross-country figure, the figure plots the ratio against the continuous measure of GDP per capita rather than calendar years. Similar to the cross-country figure, there is a clear, strong positive association: at higher levels of GDP per capita, the federal government in the US collects much more tax revenues per dollar of collection expenses. There is a clear (visual) break in trend, which corresponds to the immediate post-WWII period; the positive association is much steeper in the period prior to WWII. Interestingly, the range of observed values for our average collection per expenditure measure is comparable in the US long run and across countries today - varying in both cases between (log) 3 and (log) 6 values.

We estimate the relationship between our measure of average collection per expenditure and (log) GDP per capita for the full panel of countries in Table 2a. Around half of these are currently advanced economies, specifically: Australia (1972-2017), Canada (1924-2008), Denmark (1959-2006), South Korea (1987-2016), the United Kingdom (1993-2017), and the United States (1877-2016). The other half are best described as low- and middle-income economies: Bangladesh (1964-

[5] We also find similar patterns when removing tax authorities that are in charge of hosting and maintaining the national population registry. This only removes two countries, Sweden and Norway, and the slope coefficient is essentially unchanged (0.241). In our sample, no tax authority is directly in charge of conducting the national census, neither for citizens nor for firms.

Figure 2: Tax Collections per Expenditure in the United States: 1867 to 2017



Notes: This figure plots the association between the (log) ratio of taxes collected to collection expenses against GDP per capita (in constant 2011 USD) in the US between 1867 and 2017. Progressing along the x-axis is closely related to increases in the calendar year; the vertical dashed lines mark 10-year intervals.

2015), Botswana (1971-2005), India (1946-2016), Kenya (1962-1995), Malawi (1964-1996), Pakistan (1949-2005), Sierra Leone (1964-1996), South Africa (1960-2007), Thailand (1959-2017), Uruguay (1980-2017). All regressions include country-fixed effects, such that the slope coefficient captures the average association based on changes within a country over time. To account for serial correlation in both the outcome and regressor, standard errors are clustered at the country level. Column (1) indicates a positive association, which is highly significant. Interestingly, the magnitude of the within-country association (0.644) is much larger than the association across countries (0.248, column 1 in Table 1). The following columns show that the estimated slope coefficient is robust to removing the country-year observations where the domestic tax authority is not responsible for collecting border taxes (column 2, coefficient of 0.645) or is responsible for collecting social security contributions (column 3, coefficient of 0.691).

To complete the investigation, we turn our focus to studying the relationship between collection per expenditure and GDP per capita at the sub-national level of U.S. and Indian states. Conceptually, individual state tax authorities within a country may differ in some characteristics, which impact collection per expenditure similarly to how these characteristics vary across countries. As such,

Table 2: Panel Regressions: Tax Collections per Expenditure and Income per Capita

(a) Panel of Countries

	(1)	(2)	(3)
Outcome	Log T/X	Log T/X	Log T/X
Log GDP percap	0.644***	0.645***	0.691***
	(0.171)	(0.193)	(0.157)
Country FEs	X	X	X
Countries	16	14	13
N	625	546	519
R^2	0.566	0.545	0.578
Sample Restrictions		Exclude if no customs	Exclude if soc sec incl.

(b) Panel of Sub-National States

	(1)	(2)
Outcome	Log T/X	Log T/X
Log GDP percap	0.567***	0.278***
	(0.143)	(0.0606)
Sample	US states	Indian states
	(1993-2017)	(1950-2017)
Number of states	50	28
State FEs	X	X
N	850	964
R^2	0.719	0.777

Notes: This table provides panel regressions of the (log) ratio of taxes collected to expenses on collection on log GDP per capita. Panel A consists of an unbalanced panel of countries: Australia (1972-2017); Bangladesh (1964-2015); Botswana (1971-2005); Canada (1924-2008); Denmark (1959-2006); India (1946-2016); Kenya (1962-1995); Malawi (1964-1996); Pakistan (1949-2005); South Korea (1987-2016); Sierra Leone (1964-1996); South Africa (1960-2007); Thailand (1959-2017); United Kingdom (1993-2017); United States (1877-2016); Uruguay (1980-2017). All columns include country-fixed effects. Column 1 runs the basic OLS. Column 2 excludes country-years where customs are not included in the T/X ratio. Column 3 excludes countries where the tax authority is responsible for social security. Panel B consists of an unbalanced panel of sub-national states: US states between 1993 and 2016 (column 1) and Indian states between 1950 and 2017 (column 2). All regressions include state-fixed effects. Standard errors are clustered at the country or state level. * p¡0.10 ** p¡0.05 *** p¡0.01.

it is a meaningful extension of our main investigation to investigate whether we observe similar patterns at the sub-national level. Naturally, focusing on variation within states over time also allows us to control for all unobserved variation at the country-year level, which may confound our interpretation of the results in both the cross-section and panel of countries.

Table 2b reports the results of a basic regression, which includes state-fixed effects and clusters standard errors at the state level. Column 1 focuses on the sample of U.S. states (1993-2017), and column 2 focuses on Indian states (1950-2017). In both cases, we uncover a positive, precisely estimated coefficient. The coefficient within U.S. states over time (0.567) is comparable in magnitude to the country panel regression (Table 2a), while the coefficient within Indian states over time (0.278) is more in line with the country cross-sectional regression (Table 1).

In summary, based on novel data, we find that the collection per dollar of administrative expenditure, measured as the ratio of taxes collected to collection expenses, varies positively with development across space and time – across countries today, within countries over the long run, at both the national and sub-national levels of government.

3 Model of Tax Productivity and Development

We next turn attention to the interpretation of the patterns documented above, and their implications for economic development. Our interpretation is based on a version of the neoclassical growth model in which tax revenue finances public capital, which enters modern production à la Baxter and King (1993). The economy consists of a modern or a traditional sector, based on the long tradition in macroeconomic approaches to development starting with Lewis (1954). We adapt the Ramsey framework, which has seldom been applied in a cross-country quantitative setting.

3.1 Environment

3.1.1 Production and Taxation

Production occurs using one of two production technologies. Firms in the modern sector produce output according to a technology that features both private and public capital:

$$Y_M = F(A^M, k_G, k, L_M) \tag{1}$$

where A^M is a technology parameter, k_G denotes public capital, k is private capital, and L_M is the

effective labor employed in the modern sector. We assume that the production F is of constant returns to scale in private capital and labor (k, L_M). In each period t, modern firms hire workers at wage $w_{M,t}$ and rent capital at rate r_t. Their profit maximization problem is

$$\max_{k_t, L_{M,t}} F(A^M, k_{G,t}, k_t, L_{M,t}) - w_{M,t} L_{M,t} - r_t k_t. \tag{2}$$

Firms in the traditional sector use a linear production function:

$$Y_T = A^T L_T, \tag{3}$$

where A^T is the productivity parameter and L_T is the labor employed. Traditional firms hire workers at a fixed wage $w_T = A^T$.

Tax collections in the model work as follows. In each period t, Taxes T_t are collected on modern sector wages. However, the tax rate – $T_t / w_{M,t} L_{M,t}$ – depends positively on the tax collection personnel $L_{C,t}$ and a country-specific tax collection productivity A^C. Let $G(A^C, L_{C,t}) \leq 1$ denote the government tax productivity function that determines the rate of modern sector wages collected so that taxes $T_t = G(A^C, L_{C,t}) w_{M,t} L_{M,t}$. Without loss of generality, tax collection personnel are paid at the modern sector wage, and we assume that the tax productivity function has decreasing returns to scale in $L_{C,t}$.

Note that governments do not choose statutory tax rates in this model. The idea is that they can choose how much effort they want to apply via the number of tax collectors, $L_{C,t}$. This is meant to capture the idea that simply announcing a higher tax rate in low-income countries doesn't lead taxpayers to pay more.

3.1.2 Households

The modern and traditional sector workers live in a household that constitutes the representative consumer. The representative household owns all capital and supplies modern and traditional sector labor. The representative household maximizes lifetime utility subject to the intertemporal budget constraint,

$$\max_{\{c_t, k_{t+1}, b_{t+1}\}} \sum_{t=0}^{\infty} \beta^t u(c_t, L_{M,t}, L_{T,t}),$$
$$c_t + k_{t+1} + b_{t+1} \leq (1 - \omega_t)(1 + r_t - \delta) k_t + (1 - G(A^C, L_{C,t})) w_{M,t} L_{M,t} + w_T L_{T,t} + R_t b_t. \tag{4}$$

with initial private capital $k_0 = \bar{k}_0$. In this equation, r_t is the return on private capital, δ is the depreciation rate of private capital, and R_t is the return on bonds. The terms ω_t are a sequence of given wedges on the household's capital choice. These wedges capture for realism capital distortions in emerging markets, such as capital expropriation and the fact that the provision of public capital can distort private capital accumulation. Note that income taxation affects sectoral allocations.

3.1.3 Government and Public Capital Accumulation

The government hires tax collection personnel and collects taxes. The government can borrow to finance its overall budget. That is the budget constraint:

$$w_{M,t}L_{C,t} + R_t b_t + k^G_{t+1} \leq T_t + b_{t+1} + (1-\delta^G)k^G_t + \omega_t(1 + r_t - \delta)k_t \tag{5}$$

where δ^G is the depreciation rate of public capital. However, the government cannot borrow to finance public capital,[6] that is, public capital evolves according to

$$k_{G,t+1} = (1-\delta^G)\, k_{G,t} + T_t, \tag{6}$$

Higher tax revenues directly feed into more public capital, which in turn boosts modern production.

Market clearing in the goods market requires that

$$c_t + k_{t+1} + w_{M,t}L_{C,t} + k_{G,t+1} \leq (1 + r_t - \delta)k_t + w_{M,t}L_{M,t} + A^T L_{T,t} + (1-\delta^G)k_{G,t}. \tag{7}$$

3.1.4 Equilibrium

A competitive equilibrium in this economy consists of allocations $\{c_t, L_{T,t}, L_{M,t}, L_{C,t}, k_t, k_{G,t}, b_t\}$ and prices $\{w_{T,t}, w_{M,t}, r_t, R_t\}$ such that:

1. Firm optimality: Modern firms take prices $\{w_{M,t}, r_t\}$ and public capital $k_{G,t}$ as given, solve modern producer's problem (2) and make zero profit in equilibrium. Traditional firms hire labor at wage $w_t = A^T$,

[6] We can relax this assumption by assuming that the government borrows at a higher rate than the households to finance its budget. Note that in equilibrium, this market structure for government finances implies that the government can borrow short-term to finance inputs in tax collection –here public sector wages— while public capital is the sum of depreciated values of taxes collected. One interpretation for this model of tax collection and development is that there are payment in advance constraints for tax collectors that require debt financing.

2. Household optimality: representative households take prices $\{w_{T,t}, w_{M,t}, r_t, R_t\}$ and tax collection personnel $L_{C,t}$ for given and solve the consumer problem (4),

3. Budget constraints: government's budget constraint (5) and public capital accumulation equation (6) are satisfied,

4. Market clearing: markets clear (7).

3.2 The Government's Problem

To study optimal policy, we now formulate the Ramsey problem. The government (or tax authority) chooses sequences of tax collection personnel and private capital wedge $\pi = \{L_{c,t}\}$ to maximize the representative household's lifetime welfare $\sum_{t=0}^{\infty} \beta^t u(c_t(\pi), L_{M,t}(\pi) + L_{T,t}(\pi))$ where $c_t(\pi), L_{M,t}(\pi), L_{T,t}(\pi)$ are equilibrium allocations with tax collection.[7]

Following the primal approach, we formulate the government's problem as a choice on allocations directly. Conditions for these allocations to correspond to those chosen by the households in equilibrium result in an implementability condition that constrains the government's choices.

Lemma 1. *The implementability condition on the government's allocation choice is:*

$$\sum_{t=0}^{\infty} \beta^t \left[u_c(t) c_t - u_{L_M}(t) L_{M,t} - u_{L_{T,t}}(t) L_{T,t} \right] = u_c(0) R_0 k_0. \tag{8}$$

Proof. The household maximizes

$$\max_{\{c_t, k_{t+1}, b_{t+1}\}} \sum_{t=0}^{\infty} \beta^t u(c_t, L_{M,t}, L_{T,t})$$

subject to the sequence of budget constraints

$$c_t + k_{t+1} + b_{t+1} = (1 - \omega_t)(1 + r_t - \delta) k_t + (1 - G(A^C, L_{C,t})) w_{M,t} L_{M,t} + w_T L_{T,t} + R_t b_t,$$

with given k_0. Let λ_t be the Lagrange multiplier on the period–t constraint. The corresponding first–order conditions (FOCs) include: For consumption:

$$u_c(c_t, L_t) = \lambda_t.$$

[7] In this environment, many tax instruments are equivalent. An increasing tax on consumption is equivalent to taxes on private capital. In addition, we can show similarly to Chamley (1986) that optimal private capital taxes are zero starting in period one. Therefore, we focus on income taxation and the productivity of tax collection.

For capital:

$$\lambda_t = \beta\,\lambda_{t+1}\,(1-\omega_t)(1+r_{t+1}-\delta),$$

For bonds:

$$\lambda_t = \beta\,\lambda_{t+1}\,R_{t+1}.$$

For labor in each sector: The intratemporal conditions (from differentiating with respect to L_M and L_T) yield

$$u_{L_M}(L_M) + \lambda_t\,(1-G(A^C, L_{C,t}))\,w_{M,t} = 0,$$

$$u_{L_T}(L_t) + \lambda_t\,w_T = 0.$$

where $L_t = L_M + L_T$ is total labor. Multiply the period–t budget constraint by the discount factor and the multiplier:

$$\beta^t\lambda_t\Big[c_t + k_{t+1} + b_{t+1}\Big] = \beta^t\lambda_t\Big[(1-\omega_t)(1+r_t-\delta)k_t + (1-G(A^C, L_{C,t}))w_{M,t}L_{M,t} + w_T\,L_{T,t} + R_t\,b_t\Big].$$

Now we sum over $t = 0$ to ∞. Because the asset evolution terms (for k_{t+1} and b_{t+1}) and the Euler equations imply that the multipliers λ_t evolve consistently with the returns, the terms involving future assets cancel out telescopically. In other words, when you write the sequence $\beta^t\lambda_t\,k_{t+1}$ and $\beta^t\lambda_t, b_{t+1}$, the Euler equations ensure that almost all terms cancel with the corresponding terms coming from $\beta^{t+1}\lambda_{t+1}\,(1-\omega_t)(1+r_{t+1}-\delta)\,k_t$ and $\beta^{t+1}\lambda_{t+1}\,R_{t+1}\,b_t$. The remaining term is the one associated with the initial capital $\lambda_0\,R_0\,k_0$ and by assumption $b_0 = 0$. Rearrange the labor FOCs as follows:

$$\lambda_t\,(1-G(A^C, L_{C,t}))\,w_{M,t} = -u_{L_M}(c_t, L_t),$$

$$\lambda_t\,w_T = -u_{L_T}(c_t, L_t).$$

Replacing these terms in the sum over the budget constraints, the contributions from the wage income parts become:

$$\beta^t\lambda_t\,(1-G(A^C, L_{C,t}))\,w_{M,t}\,L_{M,t} = -\beta^t u_{L_M}(c_t, L_t)L_M,$$

$$\beta^t\lambda_t\,w_T\,L_{T,t} = -\beta^t u_{L_T}(c_t, L_t)L_T.$$

After the telescoping cancellation from asset terms and substituting the expressions for the wage incomes, the summed expression becomes

$$\sum_{t=0}^{\infty}\beta^t\Big[\lambda_t c_t - u_{L_M}(c_t, L_t)L_M - u_{L_T}(c_t, L_t)L_T\Big] = \lambda_0\,R_0\,k_0.$$

Recalling that $\lambda_t = u_c(c_t, L_t)$ (and writing $u_c(t)$ for brevity) and $\lambda_0 = u_c(0)$, we obtain

$$\sum_{t=0}^{\infty} \beta^t \left[u_c(t)c_t - u_{L_M}(t)L_M - u_{L_T}(t)L_T \right] = u_c(0) R_0 k_0.$$

Thus, the result.

□

This condition ensures that the chosen sequence of consumption, labor, and capital is consistent with the household's optimal choices, initial endowment, and the evolution of prices.

The government then chooses allocations in order to maximize the representative household's lifetime welfare subject to (i) the economy's resource constraints, (ii) the implementability condition (8), and (iii) the government's budget constraint.

Formally, the government's problem is:

$$\max_{\{c_t, k_{t+1}, L_{M,t}, L_{T,t}, L_{C,t}\}} \sum_{t=0}^{\infty} \beta^t u(c_t, L_{M,t}, L_{T,t})$$

subject to

$$c_t + k_{t+1} + \underbrace{F_{L_M}(k_t, k_t^G, L_{M,t})}_{w_{M,t}} L_{C,t} + k_{G,t+1} \leq F(k_t, k_t^G, L_{M,t}) + A^T L_{T,t} + (1-\delta)k_t + (1-\delta^G)k_t^G$$

$$k_{G,t+1} = (1-\delta^G)k_{G,t} + G(A^C, L_{C,t}) \underbrace{F_{L_M}(k_t, k_t^G, L_{M,t})}_{w_{M,t}} L_{M,t},$$

the implementability condition (8) holds.

For ease of notation, we next specify G with the decreasing returns to tax collection personnel $G(A^C, L_C) = A^C L_C^{\eta}$ where $0 < \eta < 1$. We derive the follow optimal tax collection formula at the steady state.

Proposition 1 (Steady-State)**.** *Suppose there exists an interior steady state, then the optimal tax collection choice is determined by:*

$$\frac{1-\beta\left(1-\delta_G\right)}{G_{L_C}(A^C,L_C)\,L_M} = \underbrace{\frac{F_G+1-\delta_G}{(F_k+1+-\delta)(1-\omega)}-1}_{\text{net return of public capital relative to private capital}} + \underbrace{\beta\left(\frac{1-\eta}{\eta}\right)F_{L_MG}L_{C,}}_{\text{net direct benefit from wage increases - personnel pay}} \tag{9}$$

We relegate the proof to Appendix A.1.

The optimal tax collection rule encapsulates several economic trade-offs. The left-hand side is an increasing function of tax collection personnel. The first term of the right-hand side is the net marginal return of private capital relative to public capital in terms of output. All else equal, a high marginal return on private capital (i.e., a high F_k) makes investment in private capital more attractive relative to public hiring for tax collection. In contrast, a high marginal return on public capital (i.e., a high F_G) means that public capital is particularly productive. This enhanced productivity makes public hiring more appealing. Finally, improved tax collection productivity (a higher A^C) directly raises the marginal effectiveness of tax collectors, enabling the government to gather higher tax revenues. Finally, a larger modern sector (a higher L_M) expands the taxable base, which naturally increases total tax collections. The second term is the direct net marginal benefit from wage increases minus personnel pay. A larger net marginal from wage increases, net of tax, personnel pay increases the government's incentive to hire more tax collectors. We next provide sufficient conditions under which such a steady-state exists.

Lemma 2. *Suppose that household preferences are separable in consumption and labor inputs, and of the type CRRA in consumption, and the capital wedge is efficient. Then, there exists an interior steady state.*

4 Quantitative Analysis

We now set the functional forms of the production technologies, utility function, and tax collection production function to perform our quantitative analysis. The tax productivity function has decreasing returns in tax administration labor L_C and is given by:

$$G(A^C,L_C) = A^C L_C^{\eta}, \quad \text{where } 0<\eta<1. \tag{10}$$

We follow the World Bank Long-Run Growth Model and consider modern-sector production function that combines public capital k^G, private capital k, and labor L_M, to produce output, with the TFP parameter A_M, the efficiency of public investment $\theta \leq 1$ which can capture corruption, the

elasticity of public capital ζ and the private capital share μ:

$$F_M(k^G,k,L_M) = A_M(\theta k^G)^{\zeta} k^{\mu} L_M^{1-\mu}, \quad \text{where } \zeta,\mu,\in(0,1). \tag{11}$$

The traditional-sector production function is linear, the utility of consumption c is CRRA with risk aversion coefficient γ, and the disutility of each labor input is isoelastic with Frisch elasticity of labor supply ε.

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad \gamma > 0, \qquad u(L) = \frac{L^{1+1/\varepsilon}}{1+1/\varepsilon}, \quad \varepsilon > 0. \tag{12}$$

4.1 Steady-State

A steady state is determined by the system of 9 equations and nine unknowns $(w_M, w_T, r, L_m, L_T, L_C, k, k_G, c)$

$$\text{optimal collection:} \quad \frac{1-\beta(1-\delta_G)}{G_{L_C}(A^C,L_C)L_M} = \beta\left[F_G + (1-\delta_G)\right] - 1 + \beta\left(\frac{1-\eta}{\eta}\right) F_{L_M G} L_{C,} \tag{13}$$

$$\text{private capital:} \quad 1 = \beta(F_k + 1 - \delta) + \beta\frac{1-\eta}{\eta} F_{L_M,k} L_C \tag{14}$$

$$\text{public capital:} \quad k_G = \frac{1}{\delta^G} G(A^C,L_C) w_M L_M \tag{15}$$

$$\text{return on private capital:} \quad r = F_{M,k} \tag{16}$$

$$\text{modern sector wage:} \quad w_M = F_{M,L} \tag{17}$$

$$\text{traditional sector wage:} \quad w_T = F_{T,L} \tag{18}$$

$$\text{modern labor:} \quad u_{L_M} = u_c w_M \times \left(1 - G(A^C,L_C)\right) \tag{19}$$

$$\text{traditional labor:} \quad u_{L_T} = u_c w_T \tag{20}$$

$$\text{agg. res. constraint:} \quad c + w_M L_C + \delta^G k_G = (r-\delta)k + w_M L_M + A^T L_T \tag{21}$$

4.2 Calibration of General Parameters

We begin by assigning some values to "general" parameters that apply to all countries. We deliberately taking standard choices wherever they are available. As each period represents one year, we set the discount factor, β, to be 0.96, and the depreciation rates of physical and public capital, δ and δ_G, to be 0.06. We assume that capital's share in the production function is 0.33, giving one-third of labor income to private capital. We pick a labor supply elasticity, ε, of 0.5, and a

coefficient of relative risk aversion, γ, of 2. We normalize the productivity of the traditional sector, A_T, to be 1.

Two important parameters for our analysis that are less standard in macroeconomics are the degree of returns to scale in the tax administration sector, η, and the elasticity of output to public goods, ζ. We discuss each of these in turn.

For the former, we set $\eta = 0.05$. In our model, η corresponds to the ratio of the marginal product of tax collections to the average product (e.g., total collections per tax administration expenditure, the subject of our empirical analysis). Several types of evidence point to marginal products being significantly smaller than average products. First, it is well established that tax authorities organize their resources and activities to prioritize the more easily taxable segments of the economy (Sugin, 2014; Christians, 2013; OECD, 2024; IMF, 2015; EU, 2015). For example, wage employees are easier to tax than self-employed (Kleven, Knudsen, Kreiner, Pedersen, and Saez, 2011), larger firms are easier to tax than smaller firms (Pomeranz, 2015; Kleven, Kreiner, and Saez, 2016), and consumers in large, modern retail stores are more easily taxed on their expenditure than consumers in small, traditional stores (Bachas, Gadenne, and Jensen, 2024). Most tax codes exempt firms below certain thresholds (of size, assets, or number of employees) and in certain sectors (often agriculture) from the main tax bases.[8] For individuals, almost all personal income tax (PIT) schedules around the world feature an exemption threshold, whose location is sometimes justified based on administrative resource considerations. [9]

Evidence for marginal collections exists for the property tax in Ghana, a low-income country (Dzansi, Jensen, Lagakos, and Telli, 2025, based on). In the cross-section of local governments, results suggest that the marginal collection per dollar of expenses for a local tax collector is 1.56. The findings are remarkably similar based on an RCT within a local government, where an experiment that increases both collections and expenses for a random subset of collectors yields a marginal collection per dollar of expenses of 1.62. In Ghana, average collections per dollar of expenses is 43 – leading to a ratio of marginal to average of 0.04. Moving to a middle-income

[8] Manuals for practitioners explicitly justify this policy design as driven by administrative cost-benefit considerations: in *The Modern VAT*, IMF (2001) states that "for the great majority of small traders, the cost of taxing them would exceed the VAT that could be collected, so an exemption threshold is efficient and sensible."

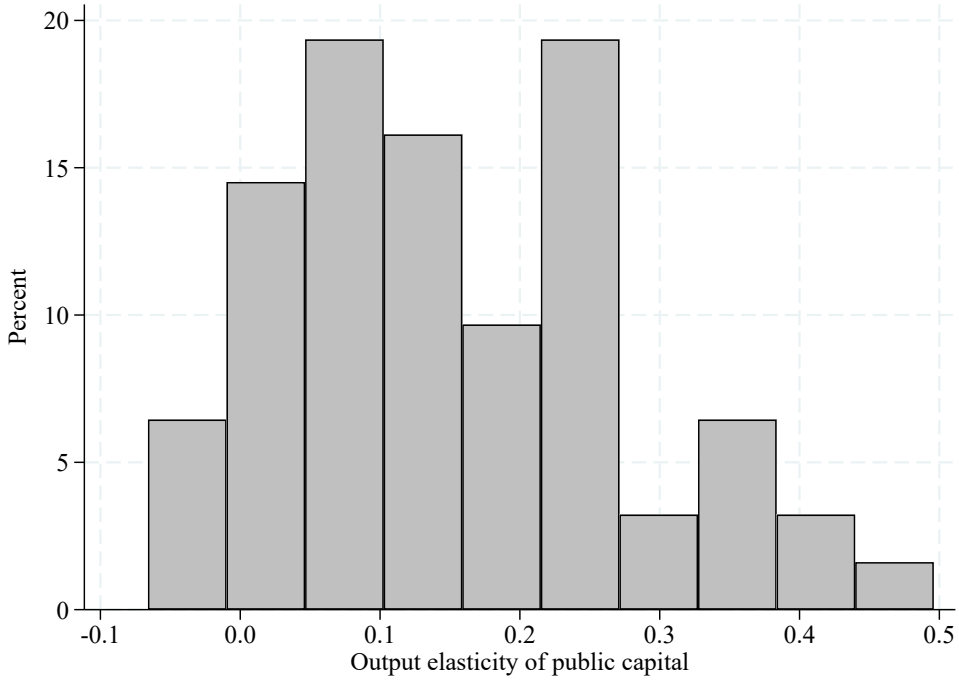
[9] In the early 20th century, when the US featured a large self-employed workforce that was hard to tax, Cordell Hull, the legislator who led the introduction of the US PIT, commented that "A lower exemption would be difficult of enforcement and also would entail so much administrative work as to make it unprofitable(...) a lower exemption than $4,000 would embrace more complicated taxpayers while this bill proposes to collect satisfactorily at the source [i.e., on employees]" (cited in Blakey, 1914). In the 2015 budget speech, while outlining a reform to the PIT exemption threshold, Ghana's deputy finance minister argued that, by practically exempting most self-employed and small businesses, "the increase in threshold would help ease the system and ensure efficient administration of the tax" (of Ghana, 2015). Finally, in the domain of compliance, a key policy prescription by international organizations is to segment taxpayers and focus enforcement resources on larger firms and more well-off individuals, through the creation of medium and large taxpayer offices (IMF, 2015; World Bank, 2024).

country, in Costa Rica, Brockmeyer, Smith, Hernandez, and Kettle (2019) experimentally estimates the impacts of a policy that increases enforcement among nil filers (registered taxpayers that declare a value of 0 for their taxable activities). The paper estimates a marginal collection per dollar of expenses of 4; with an average collection per expenses in Costa Rica of 82, the implied ratio of marginal to average returns is 0.05. Finally, evidence from a high-income country can be found in Boning, Hendren, Sprung-Keyser, and Stuart (2025), which estimates the marginal collections from random audits of individual income taxpayers in the US. The authors find that the marginal collection per dollar is 12.8. With an average collection of 260 in the US, the ratio of marginal to average returns in the US is 0.05. Our choice of $\eta = 0.05$ is consistent with these three pieces of evidence.

For the elasticity of output to public goods, ζ, we choose two values: a lower value of 0.17 and higher value of 0.42. There have been a number of attempts to estimate this elasticity in the literature, and with a very wide range of values, representing difficulties in estimation as well as the natural variation in the actual value of public spending. The World Bank uses these estimates to discipline its Long-Term Growth Model (see Devadas and Pennings, 2018), drawing on a meta-analysis by Bom and Ligthart (2014). For the low value we follow Devadas and Pennings (2018) in picking a value coming from studies of "core public capital (i.e., roads, railways, airports, and utilities)." Figure 3 plots the estimated elasticities from the meta-analysis; our lower value falls squarely in the middle of the range. One limitation is that 93 percent of these studies come from advanced economies, which is the main focus of our study. Another challenge is that there are natural challenges to identifying this parameter, and a significant number of the studies cited in the Bom and Ligthart (2014) pre-date the "credibility revolution" in economics.

Several more recent studies in less developed settings are informative about the output elasticity of public capital. However these do little to narrow the plausible values for ζ. Donaldson and Hornbeck (2016) study the expansion of the U.S. railroad network to estimate the contribution of transport infrastructure to agricultural land values. Their estimates imply a long-run output elasticity of public capital of around 0.6. By contrast, Faber (2014) examines the construction of China's national highway system and finds that while highway access generates gains for directly connected areas, the implied output elasticity for peripheral regions is closer to zero. Taken together these two studies confirm the need to consider a range of plausible ζ values rather than taking one value and treating it as definitive.

Figure 3: Bom and Ligthart (2014) Output Elasticities of Public Capital



Notes: This figure plots a histogram of output elasticities of public capital from the meta-analysis of Bom and Ligthart (2014).

4.3 Cross-Country Parameterization

We set other parameters to be country specific. One of these is θ^i which governs the fraction of government tax revenue that is lost to graft or waste. Our strategy is to take three separate targets for low-, middle-, and high-income economies, since it is widely agreed that government corruption is more prevalent in less developed economies. We target values of 0.58 for the low-income economies, 0.74 for the middle-income countries, and 0.84 for the high-income ones. These values are used by the World Bank's Infrastructure Efficiency Index (IEI); see (see Devadas and Pennings, 2018). To our knowledge there are no country-specific values of this parameter that have been constructed.

There are two remaining sets of parameters that we aim to target jointly: A_C^i, the tax productivity in country i, and A_M^i, the efficiency of production in country i's modern sector. We target these jointly to match two targeted moments in each country: the T/X ratio studied in the empirical portion of the paper above, and the second is the ratio of taxes to GDP in

work earlier in the paper, and

which is consistent with our data from the world's poorest countries, and a modern-sector employment share of 50 percent, which is a plausible value for countries in this income range. For example, this is approximately the agriculture sector's employment share in Sub-Saharan Africa, and most of agricultural production is at small-scale farms that may be hard to tax.

Figure 4

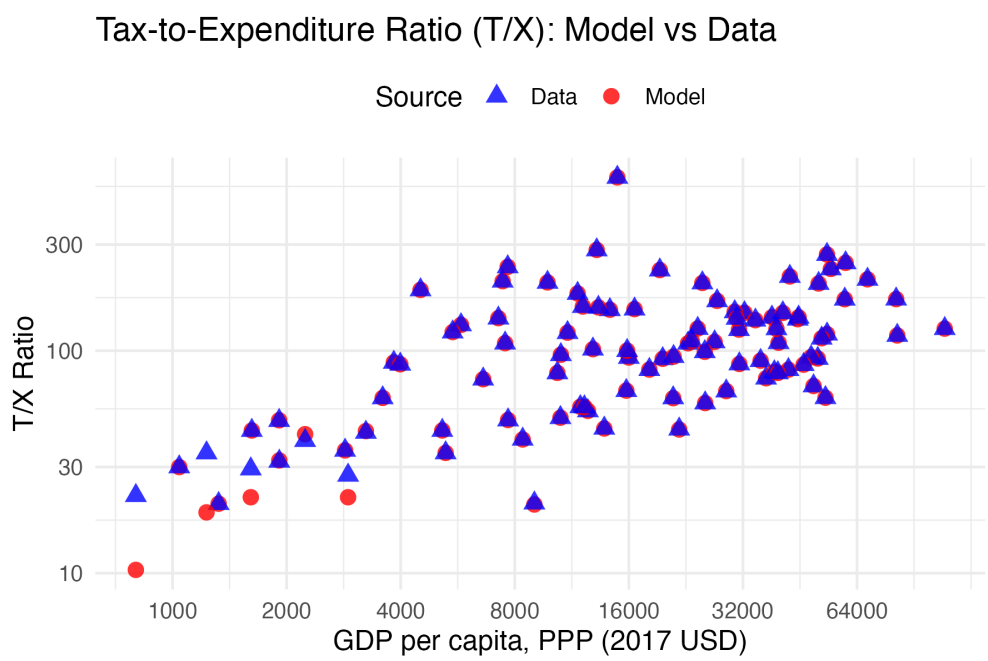
The calibrated model does not target the share of total employment in the administration sector, but it does make reasonable predictions for this statistic. In the model, around 0.1 percent of employment is in the tax administration sector, and this is broadly consistent with evidence from Sub-Saharan Africa. Tax collections relative to GDP are also not targeted directly. In our model, this is 2.9 percent. In the data, Sub-Saharan African countries have tax-to-GDP ratios averaging around 10 percent (Appendix Table 5). Though our model only features taxes on labor income, and expenditures that are entirely used for investments in public goods provision, with no role for transfers between households in the model. When matching our model's predictions to the share of tax revenues that go toward public investments, our model's predictions are closer to the mark. Evidence suggests that closer to one-third of tax revenues go into public investments (as opposed to military spending, transfers, and debt repayments); Comelli, Kovacs, Montoya Villavicencio, Sode, David, and Eyraud (see e.g. 2023, and the IMF Government Finance Statistics).

4.4 Quantitative Counterfactuals

Once the model is calibrated to match low-income countries, our main quantitative exercise simulates the long-run general equilibrium effects of raising the efficiency of the tax administration sector. We compute the effects in particular of an increase in A_C of 10 percent, 20 percent, and 50 percent. We consider that productivity increases in this range are eminently feasible given existing policy levers available to developing economies. For example, in an experiment by a subset of the authors (Dzansi, Jensen, Lagakos, and Telli, 2025), better databases of property owners linked to hand-held tablets with GIS capabilities significantly increased property tax collection outcomes in Ghana. The probability that tax collectors successfully delivered bills to the right property owner increased by 27 percent, and tax collections per worker increased by nearly 100 percent. Although these results do not literally map into changes A_C, they suggest that the range we choose is not implausible.

Table 3 reports the changes in GDP and other key variables in the new steady state resulting solely from a change in A_C. Even a 10 percent increase in A_C results in a long-run GDP gain of 1.2 percent, which rises to 4.7 percent from an increase in A_C of 50 percent. As a frame of reference, the GDP increases suggested by Hulten's theorem, in its simplest form, come from taking the

Figure 4: Model Moments and Data



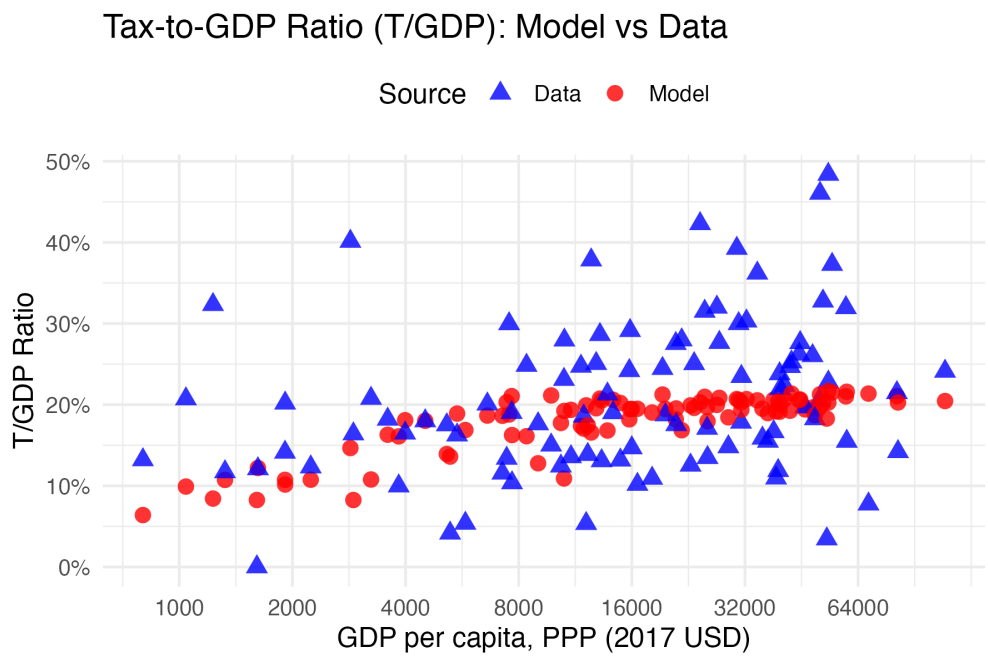
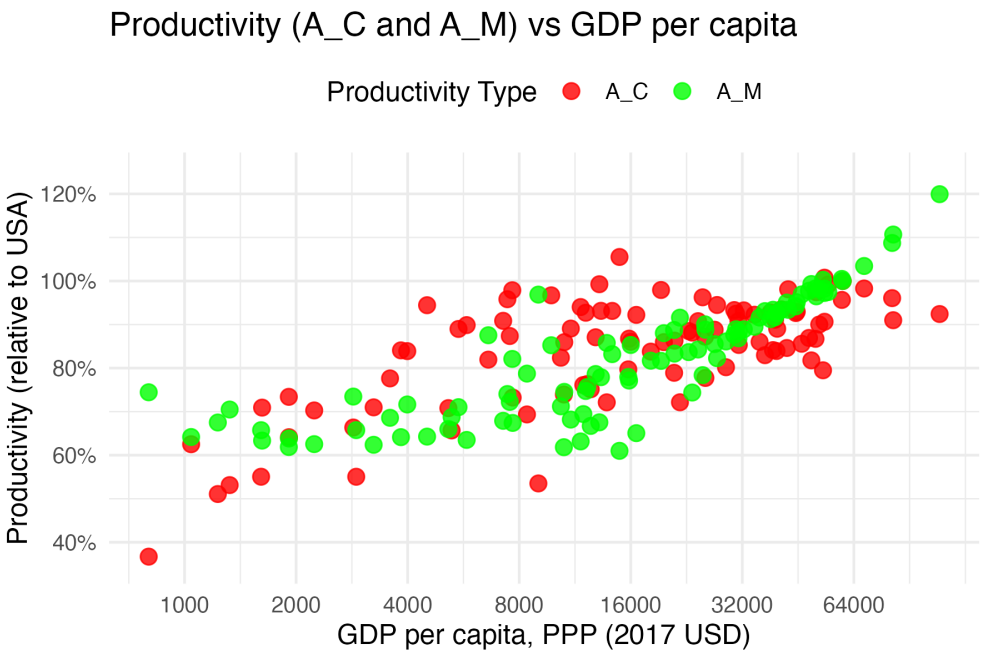


Figure 5: Implied Productivities, A_C and A_M



tax administration's Domar weight (which is 0.1 percent) and multiplying it by the increase in A_C. Using this simple yardstick, the GDP gains from raising A_C should be 0.11 percent and 0.15 percent, respectively. In other words, the GDP gains from improving a tiny sector should be tiny. Our model's predicted GDP increases, by comparison, are an order of magnitude larger.

How does the model predict such outsized long-term gains from upgrading productivity in the tiny tax administration sector? The answer comes not from within the sector itself, but from the resulting increase in tax revenues, which increase by 57 percent in the case of a 50 percent increase in A_C. The result is a similar percentage increase in public capital stocks, which raises the marginal product of private capital and stimulates an 8.6 percent increase in private investment. These are the two main forces, in a simple growth accounting sense, for why GDP increases so much in the long run. Shifts in labor across sectors do not play a central role, with only a modest increase in the size of the modern sector. Taxes per unit of expenditure increase by up to 155 percent, which looks more like the range of T/X in the middle-income countries.

Of course, middle-income countries differ in other ways than low-income countries. In the work in progress, we are working on computing and validating our model's predictions for the middle-income and high-income countries too.

Table 3: Effects of Raising Tax Productivity, A_c

(a) Lower ζ

ΔA_C	ΔGDP	ΔHulten$_{\text{im}}$	ΔHulten$_{\text{sm}}$	Δk	Δk_G	L_M/L	L_C/L	T/X
0	-	-	-	-	-	0.45	0.0014	31
10	1.1	0.01	0.31	2.4	12.3	0.45	0.0012	41
20	2.1	0.02	0.61	4.3	24.4	0.45	0.0009	54
50	4.4	0.05	1.53	8.6	57.2	0.46	0.0004	155

(b) Higher ζ

ΔA_C	ΔGDP	ΔHulten$_{\text{im}}$	ΔHulten$_{\text{sm}}$	Δk	Δk_G	L_M/L	L_C/L	T/X
0	-	-	-	-	-	0.43	0.0036	32
10	10.9	0.02	0.78	25.1	38.6	0.45	0.0030	43
20	20.4	0.05	1.55	45.6	76.6	0.46	0.0025	59
50	41.1	0.12	3.89	85.3	175.7	0.49	0.0008	221

Notes: This table reports the results of calibrating the model to a low-income country and raising A_C, using lower (Panel A) and higher (Panel B) values of ζ.

Table 4: Effects of Increasing Government Efficiency in Spending, θ

	θ	ΔGDP	Δk	Δk_G	L_M/L	L_C/L	T/X
θ_{low}	0.58	-	-	-	0.43	0.0036	32
θ_{middle}	0.74	48.5	118.9	120.0	0.51	0.0042	33
θ_{high}	0.84	76.8	184.6	186.0	0.55	0.0045	33

5 Conclusion

This paper develops the concept of tax productivity—the efficiency with which governments convert administrative expenditures into tax revenues—and studies its relationship with economic development. Using comprehensive data covering 100 countries across a wide range of income levels and spanning multiple decades, we document a robust positive association between tax collections per expenditure on collections and GDP per capita, observed consistently in both cross-sectional and longitudinal analyses. We interpret these empirical patterns through a quantitative growth model that incorporates tax administration into a Ramsey framework. The model highlights that increased tax productivity enhances a government's ability to finance productive public investments, thus contributing to higher levels of economic output. Preliminary quantitative findings indicate that improvements in tax productivity could substantially impact aggregate outcomes, and significantly more than suggested by the tax administration sector's tiny share of total employment.

References

AUCLERT, A., M. CAI, M. ROGNLIE, AND L. STRAUB (2024): "Optimal Long-Run Fiscal Policy with Heterogeneous Agents," Unpublished Manuscript, Harvard University.

BACHAS, P., L. GADENNE, AND A. JENSEN (2024): "Informality, Consumption Taxes, and Redistribution," *Review of Economic Studies*, 91(5).

BALAN, P., A. BERGERON, G. TOUREK, AND J. L. WEIGEL (2022): "Local Elites as State Capacity," *American Economic Review*, 112(3).

BARRO, R. J. (1990): "Government Spending in a Simple Model of Endogeneous Growth," *Journal of Political Economy*, 98(5).

BAXTER, M., AND R. G. KING (1993): "Fiscal Policy in General Equilibrium," *American Economic Review*, pp. 315–334.

BESLEY, T., E. ILZETZKI, AND T. PERSSON (2013): "Weak States and Steady States," *American Economic Journal: Macroeconomics*, 5(4), 203–235.

BESLEY, T., AND T. PERSSON (2011): *Pillars of Prosperity*. Princeton University Press.

BHANDARI, A., D. EVANS, Y. YAO, AND E. R. MCGRATTAN (2024): "Business Income Underreporting and Public Finance," Federal Reserve Bank of Minneapolis Working Paper.

BLAKEY, R. (1914): "The Income Tax Exemption," *Outlook*.

BOERMA, J., AND E. MCGRATTAN (2020): "Quantifying Efficient Tax Reform," Federal Reserve Bank of Minneapolis Working Paper.

BOM, P. R., AND J. E. LIGTHART (2014): "What Have We Learned From Three Decades of Research on the Productivity of Public Capital?," *Journal of Economic Surveys*, 28(5), 889–916.

BONING, W. C., N. HENDREN, B. SPRUNG-KEYSER, AND E. STUART (2025): "A Welfare Analysis of Tax Audits Across the Income Distribution," *Quarterly Journal of Economics*, 140(1).

BROCKMEYER, A., A. ESTEFAN, K. ARRAS, AND J. C. SUÁREZ SERRATO (2023): "Taxing Property: Theory and Evidence from Mexico," *Working Paper*.

BROCKMEYER, A., S. SMITH, M. HERNANDEZ, AND S. KETTLE (2019): "Casting a Wider Tax Net: Experimental Evidence from Costa Rica ," *American Economic Journal: Economic Policy*, 11(3).

CHAMLEY, C. (1986): "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, pp. 607–622.

CHRISTIANS, A. (2013): "Drawing the boundaries of tax justice," *Working Paper*.

COMELLI, F., P. KOVACS, J. J. MONTOYA VILLAVICENCIO, A. SODE, A. DAVID, AND L. EYRAUD (2023): "Navigating Fiscal Challenges in Sub-Saharan Africa: Resilient Strategies and Credible Anchors in Turbulent Waters," Departmental Paper 2023/007, International Monetary Fund.

CONESA, J. C., S. KITAO, AND D. KRUEGER (2009): "Taxing Capital? Not a Bad Idea after All!," *American Economic Review*, 99(1), 25–48.

D'ARCY, M., M. NISTOTSKAYA, AND O. OLSSON (2024): "Cadasters and Economic Growth: A Long-Run Cross-Country Panel," *Journal of Political Economy*, 132(11).

DEVADAS, S., AND S. PENNINGS (2018): "Assessing the Effect of Public Capital on Growth: An Extension of the World Bank Long-Term Growth Model," Policy Research Working Paper 8604, World Bank, Development Economics Development Research Group.

DINCECCO, M., AND G. KATZ (2016): "State Capacity and Long-Run Economic Performance," *Economic Journal*, 126(590).

DONALDSON, D., AND R. HORNBECK (2016): "Railroads and American Economic Growth: A "Market Access" Approach," *Quarterly Journal of Economics*, 131(2), 799–858.

DZANSI, J., A. JENSEN, D. LAGAKOS, AND H. TELLI (2025): "Technology and Tax Capacity," *NBER Working Paper No. 29923*.

EU (2015): "Tax Revenue Mobilization in Developing Countries," *Technicsl Report*.

FABER, B. (2014): "Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System," *Review of Economic Studies*, 81(3), 1046–1070.

FARHI, E., AND I. WERNING (2013): "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, 80(2), 596–635.

GENNAIOLI, N., AND H.-J. VOTH (2015): "State Capacity and Military Conflict," *Review of Economic Studies*, 82(4).

GOLOSOV, M., M. TROSHKIN, AND A. TSYVINSKI (2016): "Redistribution and Social Insurance," *American Economic Review*, 106(2), 359–386.

GORDON, R., AND W. LI (2009): "Tax Structures in Developing Countries: Many Puzzles and a Possible Explanation," *Journal of Public Economics*, 93(7-8).

HJORT, J., AND L. TIAN (2024): "The Economic Impact of Internet Connectivity in Developing Countries," *Annual Review of Economics*, Forthcoming.

IMF (2001): "The Modern VAT," *IMF Publications*.

——— (2015): "Current Challenges in Revenue Mobilization: Improving Tax Compliance," *Staff Report*.

JENSEN, A. (2022): "Employment Structure and the Rise of the Modern Tax System," *American Economic Review*, 112(1), 213–234.

JONES, L. E., R. E. MANUELLI, AND P. E. ROSSI (1993): "Optimal Taxation in Models of Endogenous Growth," *Journal of Political Economy*, 101(3), 485–517.

JUDD, K. L. (1985): "Redistributive Taxation in a Simple Perfect Foresight Model," *Journal of Public Economics*, 28(1), 59–83.

KAPICKA, M. (2013): "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach," *Review of Economic Studies*, 80(3), 1027–1054.

KHAN, A., A. KHWAJA, AND B. OLKEN (2015): "Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors," *Quarterly Journal of Economics*, 131(1).

KLEVEN, H. J., M. B. KNUDSEN, C. T. KREINER, S. PEDERSEN, AND E. SAEZ (2011): "Unwilling or Unable to Cheat?," *Econometrica*, 79(3).

KLEVEN, H. J., C. T. KREINER, AND E. SAEZ (2016): "Why can Modern Governments Tax so Much?," *Economica*, 83(330).

LEWIS, A. W. (1954): "Economic Development with Unlimited Supplies of Labor," *The Manchester School*, 22(2), 139–91.

MAYSHAR, J., O. MOAV, AND L. PASCALI (2022): "The Origin of the State: Land Productivity or Appropriability?," *Journal of Political Economy*, 130(4).

NARITOMI, J. (2019): "Consumers as Tax Auditors," *American Economic Review*, 109(9), 3031–72.

OECD (2024): "Tax Administration 2022: Comparative Information on OECD and Other Advanced and Emerging Economies," *Annual Report*.

OF GHANA, G. (2015): "Parliamentary Debates," *Official Reports*.

OKUNOGBE, O., AND G. TOUREK (2024): "How Can Lower-Income Countries Collect More Taxes?," *Journal of Economic Perspectives*, 38(1).

POMERANZ, D. (2015): "No Taxation without Information," *American Economic Review*, 105(8).

SLEMROD, J. (2007): "Cheating Ourselves: The Economics of Tax Evasion," *Journal of Economic Perspectives*, 21(1), 25–48.

STRAUB, L., AND I. WERNING (2020): "Positive Long-Run Capital Taxation: Chamley-Judd Revisited," *American Economic Review*, 110(1), 86–119.

SUGIN, L. (2014): "Payroll Taxes, Mythology, and Fairness.," *Harvard Journal on Legislation*, 51.

TILLY, C. (1990): *Coercion, Capital, and European States, AD 990–1990*. Blackwell.

WORLD BANK (2024): *Revenue Administration Handbook*. World Bank Publications.

A Online Appendix

A.1 Steady State

Proof. We begin by writing the first-order conditions of the government's problem for any period $t > 0$.

For the consumption choice, we have:

$$\beta^t u_c(t) + \beta^t \mu \Big[u_{cc}(t) c_t + u_c(t) - u_{M,cL_M}(t) L_M - u_{T,cL_T}(t) L_T \Big] = \lambda_t,$$

where λ_t denotes the Lagrange multiplier on the period-t resource constraint.

The FOC with respect to private capital k_{t+1} is

$$-\lambda_t + \lambda_{t+1} \Big[F_k(t+1) + 1 - \delta - F_{L_M k}(t+1) L_{C,t+1} \Big] + \gamma_{t+1} G(A^C, L_{C,t+1}) F_{L_M k}(t+1) L_{M,t+1} = 0.$$

Rearranging, we obtain

$$\lambda_t = \lambda_{t+1} \Big[F_k(t+1) + 1 - \delta - F_{L_M k}(t+1) L_{C,t+1} \Big] + \gamma_{t+1} G(A^C, L_{C,t+1}) F_{L_M k}(t+1) L_{M,t+1}.$$

For the choice of tax collection personnel L_C, the FOC is given by:

$$-\lambda_t w_{M,t} + \gamma_t G_{L_C}(A^C, L_{C,t}) w_{M,t} L_{M,t} = 0.$$

Since $w_{M,t} > 0$, it follows that

$$\lambda_t = \gamma_t G_{L_C}(A^C, L_{C,t}) L_{M,t}.$$

The FOC with respect to public capital k^G_{t+1} is

$$\begin{aligned} -\lambda_t - \gamma_t - \lambda_{t+1} F_{L_M G}(t+1) L_{C,t+1} + \Big[1 - \delta_G + G(A^C, L_{C,t+1}) F_{L_M G}(t+1) L_{M,t+1} \Big] \gamma_{t+1} \\ + \lambda_{t+1} \Big[F_G(t+1) + 1 - \delta_G \Big] = 0. \end{aligned}$$

Rearranging this condition yields

$$\lambda_t = \lambda_{t+1} \Big[F_G(t+1) + (1 - \delta_G) - F_{L_M G}(t+1) L_{C,t+1} \Big] - \gamma_t + \Big[1 - \delta_G + G(A^C, L_{C,t+1}) F_{L_M G}(t+1) L_{M,t+1} \Big] \gamma_{t+1}.$$

Dividing by λ_t and writing the resulting expression in terms of ratios, we have

$$1=\frac{\lambda_{t+1}}{\lambda_t}\left[F_G(t+1)+(1-\delta_G)-F_{L_MG}(t+1)L_{C,t+1}\right]-\frac{\gamma_t}{\lambda_t}+\left[1-\delta_G+G(A^C,L_{C,t+1})\,F_{L_MG}(t+1)L_{M,t+1}\right]\frac{\gamma_{t+1}}{\lambda_t}.$$

So

$$1=\frac{\lambda_{t+1}}{\lambda_t}\left[F_G(t+1)+(1-\delta_G)-F_{L_MG}(t+1)L_{C,t+1}\right]-\frac{1}{G_{L_C}(A^C,L_{C,t})\,L_{M,t}}+$$
$$\left[1-\delta_G+G(A^C,L_{C,t+1})\,F_{L_MG}(t+1)L_{M,t+1}\right]\frac{1}{G_{L_C}(A^C,L_{C,t+1})\,L_{M,t+1}}\frac{\lambda_{t+1}}{\lambda_t}.$$

At an interior steady state, the relevant ratios become time-invariant. In particular, the Euler equation implies that $\lambda_{t+1}/\lambda_t=\beta(1-\omega)(1+r-\delta)$. Moreover, by setting $L_{C,t}=L_C$ and $L_{M,t}=L_M$ in steady state and using the relation

$$\frac{\gamma}{\lambda}=\frac{1}{G_{L_C}(A^C,L_C)L_M}.$$

The public capital FOC simplifies to

$$\frac{1-\beta\left(1-\delta_G+G(A^C,L_C)\,F_{L_MG}L_M\right)}{G_{L_C}(A^C,L_C)\,L_M}=\beta\left[F_G+(1-\delta_G)-F_{L_MG}L_{C,}\right]-1 \tag{22}$$

Set $G(A^C,L_C)=A^CL_C^\eta$ with $0<\eta<1$

$$\frac{1-\beta\,(1-\delta_G)}{G_{L_C}(A^C,L_C)\,L_M}-\beta L_C\,F_{L_MG}/\eta=\beta\left[F_G+(1-\delta_G)-F_{L_MG}L_{C,}\right]-1 \tag{23}$$

i.e.

$$\frac{1-\beta\,(1-\delta_G)}{G_{L_C}(A^C,L_C)\,L_M}=\frac{F_G+1-\delta_G}{(F_k+1+-\delta)(1-\omega)}-1+\beta\left(\frac{1-\eta}{\eta}\right)F_{L_MG}L_{C,} \tag{24}$$

This completes the proof.

□

A.2 Solution Method of Steady State

The original 9-equation and 9-unknown system can be effectively characterized by reducing it to a simpler 3-equation and 3-unknown system. For convenience, we choose the variables (L_M,G,k) as our main unknowns.

To begin, we note that using G allows us to directly determine the tax administration labor, L_C, via the relationship:

$$L_C = \left(\frac{G}{A_C}\right)^{1/\eta}. \tag{25}$$

Furthermore, due to the constant returns to scale (CRS) property of the production function, we have $w_M L_M = (1-\mu)F$. Thus, the equation for public capital can be rewritten as:

$$k_G = \frac{1}{\delta_G} G(1-\mu)F. \tag{26}$$

Substituting this expression into the production function $F = A_M(\theta k_G)^{\zeta} k^{\mu} L_M^{1-\mu}$ yields a simplified form:

$$F = \left[A_M \left(\frac{\theta(1-\mu)G}{\delta_G}\right)^{\zeta} k^{\mu} L_M^{1-\mu}\right]^{\frac{1}{1-\zeta}}. \tag{27}$$

With this formulation, the factor prices can now be explicitly determined. The modern-sector wage w_M and rental rate of capital r are given by:

$$w_M = \frac{(1-\mu)F}{L_M}, \quad r = \frac{\mu F}{k}, \tag{28}$$

assuming the values of L_M and k are already known.

Subsequently, we incorporate the equilibrium conditions related to labor and consumption in the modern and traditional sectors. Consumption c and traditional labor L_T are obtained through:

$$c = \left[L_M^{-1/\varepsilon} w_M (1-G)\right]^{1/\gamma}, \quad L_T = (c^{-\gamma} w_T)^{\varepsilon}, \tag{29}$$

with the normalization $w_T = A_T = 1$.

Summarizing the above transformations, we express all relevant variables $(c, L_M, L_T, L_C, k_G, k, w_M, r)$ in terms of three primary endogenous variables (L_M, G, k). After suitable substitution and elimination, we are left with three essential equilibrium conditions: optimal tax collection, private capital accumulation, and the aggregate resource constraint.