

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

October 2026

Abstract

This paper examines the role of government capacity to collect taxes in driving economic development. We first draw on a new macro-historical database to document that tax collections relative to tax administration expenditures rise sharply with GDP per capita. We then interpret this finding in a neoclassical growth model with a tax administration sector in which labor inputs are required to raise tax revenues. Countries differ in their levels of tax productivity, capturing how effectively labor is used to raise taxes, as well as their levels of productivity in creating consumption and investment goods. Governments use tax revenues to provide public capital that complements private factors in production and to make transfers. 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.

Email: rgoyal3@imf.org, anders.jensen@hks.harvard.edu, lagakos@bu.edu, an1490@stern.nyu.edu, and yz11552@stern.nyu.edu. For outstanding research assistance, we thank Negin Jafari. For helpful comments we thank Marco Bassetto, Augustin Bergeron, Michael Best, Corina Boar, V.V. Chari, German Cubas, Xavier Gabaix, Emilien Gouin-Bonenfant, Roozbeh Hosseini, Jordan Norris, Steven Pennings, Valerie Ramey, James Robinson, Florian Scheuer, Gustavo Ventura, Joachim Voth, Tao Zha, and seminar/conference participants at Amherst College, ASU, the Atlanta Fed, Copenhagen, Houston, Nanyang Technological University, Nottingham, NYU, Cornell, the 2025 NBER Summer Institute Macro Public Finance, Tufts, USC, UC Irvine, the World Bank, the Tax and Development Conference at Williams College, the 2026 NBER Summer Institute Economic Growth and Long-Run Macroeconomic Development and the African Development Bank. All 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 et al. \(2022\)](#) shows that the most successful early states, such as Egypt, were those best able to raise revenue 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 laid the groundwork for their later 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.

By contrast, 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 [Aschauer \(1989\)](#), [Barro \(1990\)](#), and [Baxter and King \(1993\)](#) incorporate public capital into the production function but abstract from the costs of collecting taxes. Tax administration costs are likewise absent from nearly all macroeconomic analyses of taxation ([Judd, 1985](#); [Chamley, 1986](#); [Chari and Kehoe, 1999](#); [Conesa et al., 2009](#); [Straub and Werning, 2020](#)) and models of fiscal multipliers ([Ramey, 2011](#); [Farhi and Werning, 2016](#)). A notable exception is the paper by [Besley, Ilizetzi 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 assesses the quantitative importance of governments’ capacity to collect taxes, which we refer to as *tax productivity*, for economic development. Methodologically, we follow the growing literature on macroeconomic development that studies the role of productivity in specific sectors, such as agriculture ([Restuccia, Yang and Zhu, 2008](#); [Caselli, 2005](#); [Boppart, Kiernan, Krusell and Malmberg, 2023](#)), investment goods ([Hsieh and Klenow, 2007](#); [Herrendorf and Valentinyi, 2012](#)), or electricity ([Colmer, Lagakos and Shu, 2024](#)), in shaping long-run development. The main novel feature of our model is a tax administration sector that requires labor inputs (e.g., “tax collectors”) to raise tax revenues. Countries differ along two main dimensions. The first is economy-wide productivity, which governs the ability to produce consumption and investment goods. The second is tax productivity—modeled as the efficiency of labor in raising tax revenue—which determines the government’s capacity to provide public capital, a complement to private factors of production. We measure tax productivity across countries and quantify the long-run aggregate effects of raising it.

We begin by creating a new dataset of tax collections and tax administration expenses in a broad set of countries. We hand-collect these statistics from national tax authorities—such as the U.S. Internal Revenue Service and the Ghana Revenue Authority—for 100 countries in the cross-section and 16 countries in a long-run, historical panel. Our outcome of interest is 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 both across countries and over time within countries. As we argue in the quantitative analysis below, it is also informative about the level of tax productivity in each country.

Our main empirical finding is that tax collections per unit of expenditure are strongly increasing in GDP per capita. The poorest countries in the world collect about \$30 in revenue for every dollar spent on collection; this number rises to around \$270 in the richest. We provide evidence that this pattern is not an artifact of differences in the scope of the revenue authorities’ activities, such as whether they are responsible for value-added taxes or not. Exploiting the panel dimension of our data, we also examine how collections per expenditure vary with GDP per capita within countries (as well as U.S. and Indian States) over the long run. We find suggestive evidence that economies tend to increase their collections per unit of expenditure as they grow, with an elasticity that is larger than in the cross-section but less robustly estimated.

We interpret these facts using a neoclassical growth model that includes productive public spending and a tax administration sector. As in [Baxter and King \(1993\)](#) and related work in macro public finance, production combines labor, private capital, and public capital into a single output good usable for consumption or investment. Net of collection costs, the government uses tax revenue to invest in new public capital or to fund transfers. Households make standard consumption, saving, and labor-supply decisions, facing a marginal tax on their labor income and taking transfers as given.

We make three modeling choices about taxation that help the model capture public finance systems across different levels of development. First, rather than having the government simply announce a tax rate, we posit that the effective tax rate depends on both a country’s tax productivity and the number of tax collectors employed. Second, rather than explicitly modeling each type of tax and the rising difficulty of collecting more from any one of them, we allow for decreasing returns to scale in tax collection. This choice is consistent with evidence that expanding tax collection requires moving from easy-to-collect sources, such as taxes on international trade, to more challenging ones, such as income and consumption taxes (e.g., [Jensen, 2022](#)). Third, we allow a fraction of tax collections to be wasted, so not all public investment contributes to production. This fraction varies exogenously with the level of development, capturing potentially greater inefficiency in public spending in poorer economies.

We use the model to measure tax productivity in each country and to simulate the long-run effects of raising it. The magnitude of these effects is not obvious: the tax collection sector employs well under one percent of the workforce in every country, and a naive application of Hulten’s theorem would predict negligible aggregate gains from raising its productivity. To move beyond this back-of-the-envelope logic, we calibrate the model to match key features of the data, including GDP per capita, the ratio of tax collections to expenditure, the tax-to-GDP ratio, the employment share in tax collection, and the share of taxes devoted to transfers. One important parameter is the elasticity of output with respect to public capital, which we set based on evidence from a meta-analysis and other estimates. A second is the degree of decreasing returns to scale in tax collection, which we choose to match evidence on the returns to expanding collection efforts in the United States and several developing countries. As we show later, both parameter values strongly shape how much tax productivity matters for development.

The model matches its targets reasonably well and also replicates several important moments not targeted in the calibration. Private capital-to-labor ratios rise almost linearly with GDP per capita in both the model and data, with nearly identical levels. Empirical measures of public capital are complicated by a basic problem: public capital stocks are typically built up from investment spending, implicitly assuming that each dollar spent yields a dollar of capital ([Pritchett, 2000](#); [Cubas, 2020](#)). Our model reproduces the overall slope with respect to income per capita, though it under-predicts the measured stocks for lower-income countries. The model’s investment share of GDP is increasing in development, as in the national accounts, though at a higher level than in the data. Labor supply is modestly decreasing in development, consistent with the cross-country and time-series evidence ([Ramey and Francis, 2009](#); [Bick et al., 2018](#); [Boppart and Krusell, 2020](#)).

The implied tax productivity estimates exhibit even more variation with income per capita than do aggregate productivity estimates. In our preferred estimates, a regression of tax productivity on log GDP per capita has a slope of around 0.80; by comparison, a similar regression yields a slope of 0.34 for aggregate productivity. We take this as evidence that the tax collection sector is particularly unproductive in poor countries, consistent with the broad claims of the state capacity literature.

We assess the empirical validity of our model-based measure of tax productivity using external evidence from the International Monetary Fund’s International Survey on Revenue Administration (ISORA). For example, the survey records whether tax authorities use third-party information from employers and banks to automatically pre-fill tax returns and conduct compliance checks. We find that such practices are strongly associated with higher tax productivity. Tax productivity is also positively correlated with the use of data science and artificial intelligence tools, and with enforcement practices that are inexpensive but organizationally complex, such as offsetting tax ar-

rears across domestic and international tax bases. Because these practices require capability rather than resources, they speak directly to productivity rather than spending. The positive association with these determinants of tax productivity persists after controlling for determinants of aggregate, country-wide productivity; moreover, country-wide productivity is itself uncorrelated with the determinants of tax productivity. These results support the interpretation of our model's tax productivity as a distinct measure of government's productive capacity rather than a simple proxy for the country's overall productivity.

The model predicts that a 50 percent increase in tax productivity raises long-run GDP per capita in low-income countries by about 18 percent on average. Consumption also rises by around 18 percent, and labor supply falls by under one percent. Much of the gain comes from the direct increase in public capital per worker, and private capital accumulation also helps expand output. By contrast, Hulten's approximation theorem suggests GDP gains of less than 0.1 percent given the tax collection sector's very modest share of economic activity. The model predicts that middle-income countries gain far less from the same percent increase in tax productivity, and high-income countries have a small positive average consumption gain. Two main forces explain this weaker response in more developed economies. First, public capital expands less in richer countries, which leads to smaller increases in wages. Second, the tax burden increase is far more onerous, lowering after-tax wages and depressing labor supply.

We conclude that improvements in tax collection productivity in the world's poorest economies likely lead to significant positive long-run gains in GDP and consumption. As current governments are moving to digitize their records and adopt AI tools in their administrative processes, our findings highlight that any resulting gains in tax collection productivity could yield substantial growth benefits. We also show that public capital stocks in lower-income countries are inefficiently low relative to what a benevolent planner would choose. This inefficiency largely explains why the model's gains so far exceed the Hulten benchmark, which applies to efficient economies: the government under-provides an important input that is not priced in markets.

One limitation of our main model is that the tax base is proportional to the wage. In reality, as a country develops, workers move from subsistence agriculture, which is hard to tax, into wage work. To address this possibility, we extend the model to include a traditional sector that cannot be taxed, but whose productivity does not increase with development. In this extension, a 50 percent increase in tax productivity raises GDP in low-income countries by 6.2 percent and consumption by 5.0 percent, down from about 18 percent in the baseline but still far above the Hulten benchmark.

Another potentially worrying simplification in the main model is that tax collection inputs are exogenous. We consider two extensions with endogenous government choices of tax collection inputs. Holding other parameters at their calibrated values, low- and middle-income countries em-

ploy too few collectors relative to the welfare-maximizing level, while high-income countries employ too many; yet closing these gaps raises GDP by less than one percent, an order of magnitude less than raising tax productivity. We reach broadly similar conclusions in the Ramsey problem. In both cases, the small gains reflect the sharply decreasing marginal product of tax collectors.

Related Literature. Our paper builds on the macroeconomic literature in which taxation funds productive public capital (see [Glomm and Ravikumar, 1992](#); [Jones et al., 1993](#), in addition to the references cited above). Our work is particularly inspired by [Ramey \(2021\)](#)’s insight that the long-run effects of public capital provision can be sizeable, particularly when initial levels of public capital start below efficient levels. Whether this is the case in low-income countries today has received much less attention. Our paper is also closely related to work of [Cubas \(2020\)](#) arguing that public capital should be separated from private capital in development accounting.

Most macro public finance work has been focused on advanced economies. Some of the more related papers include [Leeper et al. \(2010\)](#)’s neoclassical analysis of the American Recovery and Reinvestment Act, and [Bhandari et al. \(2024\)](#)’s model of tax evasion, calibrated to evidence on the impacts of tax audits. To our knowledge, no prior paper in this literature has modeled a tax administration sector; for studies of advanced economies, where tax collection comes at low cost, this may not be such a grave omission.

Finally, our paper adds to the largely microeconomic 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 et al., 2015](#)), collaboration with traditional leaders ([Balan et al., 2022](#)), information technology ([Hjort and Tian, 2025](#); [Okunogbe and Tourek, 2024](#)), including using geo-referenced databases of taxpayers ([Dzansi et al., 2025](#)), and self-enforcing tax systems ([Naritomi, 2019](#)). Much of the recent literature has focused on the taxation of property, 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 productivity, which is new to this literature.

2 Cost of Collection Across Space and Time: New Facts

This section aims to provide new facts on the cost of tax collection across space and time. Based on a newly constructed dataset, we document that, far from costless as commonly assumed in

macro-public finance models, tax authorities in most countries incur substantial costs to collect their taxes. Moreover, far from uniform, we find that the cost of collection varies substantially across countries, in a manner that is closely associated with countries' level of development.

Our measure of cost is the ratio of tax revenues collected to tax administration costs in a given country (and year). This ratio is best understood as the average (*not* marginal) tax collection per dollar of administrative expenditure for a given authority at a point in time. The measure is unit-free and comparable across countries and, within countries, over time. As the theory will make clear (Section 3 onward), this measure is closely related to, but also distinct from, tax productivity. The patterns uncovered here motivate the model of tax productivity we develop, but the objective in this section is simply to provide new facts on this cost measure without any deeper interpretation.

2.1 Data

In this subsection, we describe the data collected to measure the ratio of tax revenue to tax administration costs across a large sample of countries. The numerator T captures taxes collected nationally—typically taxes on personal and corporate income; property, wealth, and inheritance; and consumption of goods and services. We also include border taxes, a major revenue source in many developing countries. Throughout, we use actual collections rather than projected or estimated figures and exclude non-tax revenue where the data permit.

The denominator measures the national tax authority's annual expenditure X , comprising operating and capital expenditures. Operating expenditures primarily cover salaries, alongside training and miscellaneous administrative expenses. Capital expenditures cover purchasing and maintaining buildings, land, and other infrastructure. Information and communication technology (ICT) expenses fall into either category, depending on whether they represent long-term investments (e.g., purchasing a server) or day-to-day expenses (e.g., software maintenance). Where this breakdown is available in our sample, salaries and ICT are by far the largest cost categories. We include as many categories as the data permit: salaries are included in all settings, while in many, though not all, the measure is more comprehensive and also captures capital expenses.

Tax authorities share many similarities across countries, and our definition of collection cost captures most in a meaningfully comparable way across space and over time. Nevertheless, the measure has limitations. First, countries differ in the range and nature of revenues administered. The most significant distinction is probably between taxes administered by customs authorities (tariffs and other border taxes) and domestic tax authorities. Because our measure combines domestic and border collection costs, these differences in institutional assignment should not affect our ratio. In some countries, direct taxes are administered nationally while certain indirect taxes are adminis-

tered sub-nationally—for instance, sales tax collection is devolved to states in India and the United States. Second, some tax authorities perform activities not directly related to taxation, such as collecting social security contributions and maintaining the population registry. We investigate the importance of these institutional variations in our regressions. Macroeconomic changes could temporarily alter the taxable base without corresponding adjustments to collection expenditures. To mitigate this concern, we average both T and X over multiple adjacent years.

Cross-country data We collect data from several sources in a cross-section of countries. The most important source is the International Survey on Revenue Administration (ISORA) ([Asian Development Bank et al., 2026](#)). ISORA is a cross-country survey that collects standardized information on how national or federal tax authorities operate. It is coordinated through a collaboration between the IMF, OECD, CIAT, IOTA, and the Asian Development Bank. 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 collect multiple years of data between 2014 and 2017; we report the yearly average over that range to smooth fluctuations in both T and X .

Our resulting cross-country sample contains data from 98 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, the UK, Germany, France, Scandinavian countries and the United States. These countries account for 77% of the world’s population and 67% of global GDP.

We combine the tax-administration data with macroeconomic data from the Penn World Table, version 11.0 ([Feenstra et al., 2015](#); [Groningen Growth and Development Centre, 2025](#)). We measure GDP per capita as expenditure-side real GDP at chained purchasing power parities divided by population, expressed in 2021 U.S. dollars.

Historical country panel data For a subset of countries, historical data exist that allow us to construct measures of collection cost 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 1877. We collect data separately for domestic taxes and border taxes. In the former case, we rely on annual reports published 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. For historical years, we digitize the data and harmonize it with the pre-existing digital sources for years closer to the present. For the US, underlying sources are often published consistently over many years. We measure gross tax collections before refunds, including penalties and interest paid on outstanding taxpayer arrears. Our comprehensive expenditure measure includes all costs associated with collecting internal revenue. Personnel costs, including those of central officers and field agents, have been important throughout the period studied. Other cost components, including physical infrastructure and ICT, have gradually become more important over time.

We collect historical data for a sample of other countries around the world. The resulting sample consists of an unbalanced panel of 16 countries: Australia (1972–2017); Bangladesh (1964–2015); Botswana (1971–2015); Canada (1954–2008); Denmark (1959–2006); India (1946–2016); Kenya (1962–1995); Malawi (1964–1996); Pakistan (1949–2005); South Korea (1968–2016); Sierra Leone (1964–1996); South Africa (1970–2007); Thailand (1959–2017); United Kingdom (1963–2017); United States (1877–2016); Uruguay (1960–2017). The average number of years per country is 55, which constitutes a meaningfully long time series. Around 40 percent of these are currently advanced economies, while the remaining ones are best described as low- and middle-income economies.

For the historical panel, we use country-specific sources, generally scanning, digitizing, and cleaning historical publications available on paper.¹ In Denmark, Bangladesh, Kenya, and Thailand, among others, we digitized the entire series of annual reports from national statistics offices or ministries of planning and economic development. In Canada, India, and Pakistan, among others, we digitized the full set of annual government finance account publications.


Historical sub-national panel. So far, we have measured tax collections and expenses for national tax authorities. However, many countries decentralize taxation, devolving the mandate to design, implement, collect, and enforce certain taxes to sub-national governments. These governments also establish their own tax authorities, so the concept of tax collection per dollar of expense applies, in principle, at the sub-national level as well.

We create long-run panel data at the sub-national level for Indian and US states. These states collect sales tax—a substantial revenue source administered nationally in most other countries.²

¹Most were available in the Lamont Library at Harvard University.

²Lower-level governments, such as municipalities and counties, also administer specific taxes through their own tax authorities. Measuring their collection costs would be interesting but practically challenging: most countries do not systematically gather and centralize the required data across their sometimes numerous municipalities.

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 98 countries. Data are collected for each country between 2014 and 2017, and the average is taken over the years if multiple data points exist per country.

For India, we primarily digitize and harmonize the Comptroller and Auditor-General's 'Combined Finance and Revenue Accounts of the Central and State Governments'. For the US, an important source is the historical 'Book of the States' series, published annually by the Council of State Governments. The resulting panel dataset at the sub-national level contains measures of tax collection and expenses for 30 states in India (1950–2017) and 50 U.S. states (1993–2017).

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.³

³We lack GDP per capita data for Papua New Guinea and Samoa, resulting in 98 countries in Figure 1.

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

Outcome	(1) Log T/X	(2) Log T/X	(3) Log T/X	(4) Log T/X	(5) Log T/X	(6) Log T/X	(7) Log T/X
Log GDP p.c.	0.344*** (0.048)	0.773*** (0.166)	0.172* (0.088)	0.334*** (0.0702)	0.347*** (0.048)	0.358*** (0.048)	0.309*** (0.058)
<i>N</i>	98	20	78	97	95	88	71
<i>R</i> ²	0.346	0.544	0.047	0.367	0.353	0.383	0.291
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 reports regressions of the log ratio of taxes collected to collection expenses on log GDP per capita. The sample is a cross-section of countries between 2014 and 2017 (averaged at the country level when multiple observations are available). Column (1) reports the basic OLS specification, 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) controls for statutory PIT, CIT, VAT, and tariff rates (Figure A.2). 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 collections to expenses does not include customs. Column (7) excludes countries in which the tax authority is responsible for collecting and disbursing social security. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 performs remarkably well (\$283) for a middle-income country. Other middle-income countries typically collect in the range between \$80 for every dollar spent, as in Indonesia, and \$150, as in Brazil.⁴

Table 1 investigates extensions and robustness checks for the pattern uncovered in the cross-section. Column (1) reports the slope coefficient (0.344) that corresponds to the full sample in 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

⁴The 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.

tax collection to expenses and (log) GDP per capita is larger at lower levels of development (0.773 versus 0.172), but it remains significant at 10% in both samples.

In column (4), we control for statutory tax rates using the global tax rate database of [Bachas et al. \(2024b\)](#): the top personal income tax (PIT) rate, main corporate income tax (CIT) and value-added tax (VAT) rates, and average tariff rate. This tests whether variation in collection per dollar of expense reflects differences in collectible amounts, since more developed countries can levy higher rates, partly through stronger enforcement. Figure [A.2](#) shows a limited positive relationship between tax rates and GDP per capita: only PIT rates increase, while CIT and VAT rates are flat and tariff rates, as expected, decline. Moreover, the positive PIT relationship exists only among upper-middle- and high-income countries, whereas the relationship between collection per dollar of expense and GDP per capita is stronger among low- and lower-middle-income countries. Consistent with these patterns, controlling for statutory rates in column (4) leaves the slope coefficient largely unchanged (0.334).

Columns (5) through (7) exclude countries whose tax authorities perform substantially different activities, potentially limiting comparability of administrative expenses. Column (5) excludes countries where the national tax authority does not collect VAT or general sales tax, leaving the coefficient virtually unchanged (0.347). Column (6) excludes countries where the domestic tax authority does not collect border taxes, reducing the sample to 88 countries and slightly increasing the slope coefficient (0.358). Column (7) excludes countries where the tax authority collects social security contributions; the sample falls substantially to 71 countries, but the coefficient remains stable (0.309).⁵ Overall, Table 1 suggests that the positive association between average collection per expenditure and development holds across institutional settings and levels of GDP per capita.

Results Based on Panel Data We now turn to estimating the relationship between average collection per expenditure and (log) GDP per capita in the full panel of countries. Similar in spirit to exploiting the substantial variation in income levels observed in the cross-section of countries, we use long-run variation in GDP per capita within each country to estimate the association between development and collection costs within countries over time.

To accommodate different coverage periods, we divide each country's series into two equal-length subperiods based on its first and last available years.⁶ In each country, this yields two non-overlapping segments averaging approximately 27 years each. Within each country-segment, we

⁵Excluding authorities that host and maintain the national population registry removes only Sweden and Norway, leaving the coefficient essentially unchanged (0.336). No tax authority in our sample directly conducts the national census of citizens or firms.

⁶For the US, we start in 1960 to maintain a time range comparable to other countries.

calculate long-run changes in $\log T/X$ and $\log$ GDP per capita. We pool all country-time segments ($N = 32$) and regress the change in T/X on the change in GDP per capita, while including time-segment dummies.⁷ The slope coefficient captures the association between collection per dollar of expense and development within a country over an average time horizon of 27 years. Differencing removes all time-invariant country characteristics by construction. The time-segment dummies absorb systematic differences between countries' earlier and later time periods. Standard errors are clustered at the country level.

Column (1) in Table 2 indicates a positive and significant long-run coefficient that is larger in magnitude (0.572) than in the cross-section. This pattern holds in columns (2) and (3), which remove country-year observations where the domestic tax authority is not responsible for collecting border taxes (column (2), coefficient of 0.614) or is responsible for collecting social security contributions (column (3), coefficient of 0.566).

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, 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.

Columns (4) and (5) report long-run regressions for US and Indian states, respectively, constructed as in the country panel.⁸ The US coefficient is large (1.358) and significant; India's coefficient (0.607) is closer to the country-panel estimate and also significant.

As an alternative to the long-run specifications, we can estimate annual panel regressions with year and country fixed effects, or year and state fixed effects in the samples for Indian and US states. In this specification, we find that the coefficient on $\log$ GDP per capita remains positive but is smaller and statistically significant only for US states (coefficient=1.017; s.e.=0.098), not for countries (coefficient=0.150; s.e.=0.216) or Indian states (coefficient= 0.024; s.e.=0.096). The full results are available upon request. The weaker annual results may reflect gradual adjustment in

⁷For Canada, the segments are $s_1 = [1954, 1981]$ and $s_2 = [1982, 2008]$; for Bangladesh, they are $s_1 = [1964, 1989]$ and $s_2 = [1990, 2015]$. For each country c and segment s , the long-run change in $\log(T/X)$ is the difference between the last and first observations: $\Delta(T/X)_{cs} = \log(T/X)_{c,s, \text{last}} - \log(T/X)_{c,s, \text{first}}$. We similarly calculate the change in $\log$ GDP per capita, $\Delta\text{GDPpccs}$. Using all 32 country-segments (two per country across 16 countries), we estimate $\Delta(T/X)_{cs} = \beta_0 + \beta_1 \cdot \Delta\text{GDPpccs} + \theta_s + \varepsilon_{cs}$, where θ_s denotes time-segment fixed effects.

⁸The shorter US panel allows only one long-run regression, covering 1993–2017. For India, each state has two segments: $s_1 = [1950, 1983]$ and $s_2 = [1984, 2017]$. We regress long-run changes across all 60 state-segments, including segment fixed effects and clustering standard errors by state.

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

Outcome:	(1)	(2)	(3)	(4)	(5)
	Log T/X				
Log GDP p.c.	0.572** (0.218)	0.614** (0.240)	0.566** (0.225)	1.358*** (0.489)	0.607*** (0.210)
<i>N</i>	32	28	26	50	60
Sample	Historical panel of countries			US states (1993–2017)	Indian states (1950–2017)
Countries/States	16	14	13	50	30

Notes: This table provides long-run regressions of the (log) ratio of taxes collected to expenses on collection on log GDP per capita within national and sub-national governments over time. Columns (1) to (3) focus on an unbalanced panel of countries: Australia (1972–2017); Bangladesh (1964–2015); Botswana (1971–2015); Canada (1954–2008); Denmark (1959–2006); India (1946–2016); Kenya (1962–1995); Malawi (1964–1996); Pakistan (1949–2005); South Korea (1968–2016); Sierra Leone (1964–1996); South Africa (1970–2007); Thailand (1959–2017); United Kingdom (1963–2017); United States (1877–2016); Uruguay (1960–2017). For each country, we calculate two long-run changes in both log T/X and log GDP per capita, by dividing the time series into two approximately equal time segments. We cluster standard errors at the country level and include time segment 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. Column (4) estimates a single long-run regression based on changes within individual US states between 1993 and 2017, with robust standard errors. Column (5) estimates long-run changes within Indian states between 1950 and 2017, based on the same specification as for the panel of countries (columns (1)–(3)): there are two time-segments for each Indian state ($s_1 = [1950, 1983]$, $s_2 = [1984, 2017]$); time-segment fixed effects are included; standard errors clustered at the state level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

tax administration or greater sensitivity to short-run fluctuations and measurement noise. However, year fixed effects also absorb common temporal changes that are not fully captured by the long-run specification’s time-segment dummies. We therefore consider the panel results to be suggestive but less robust across specifications than the cross-sectional evidence.

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: we find a robust pattern across countries today and a suggestive pattern within countries over the long run at both the national and sub-national levels of government.

While the cost of collection is closely related to tax productivity, it does not identify productivity on its own: interpreting the observed ratio requires additional assumptions about how tax revenues and collection expenses are determined. The advantage of the measure used here is its transparency and direct comparability across countries and over time. The systematic patterns we document suggest that tax productivity varies with economic development, but establishing this connection requires a model. The next section develops such a model to formalize tax productivity, interpret

the empirical patterns, and study its consequences for long-run economic growth.

3 Model of Tax Productivity and Development

We now turn to interpreting these patterns and exploring 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 the production function as in [Baxter and King \(1993\)](#). Unlike the main macro analyses of productive public capital, which focus on advanced economies and assume costless taxation, we model revenue as costly to raise. Specifically, a tax administration sector operates a technology that converts household income into tax revenue for the government. Countries differ both in their productivity in creating consumption and investment goods and in their tax productivity, which governs how easily they raise revenue. Countries also differ in the fraction of revenue that is diverted and in the share they devote to public goods rather than transfers.

3.1 Environment

Each economy has a unit measure of households with preferences that are separable in consumption and labor supply:

$$\sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\gamma}}{1-\gamma} - \alpha \frac{L_t^{1+1/\varepsilon}}{1+1/\varepsilon} \right], \quad (1)$$

where C_t and L_t are consumption and labor at time t . All countries have the same discount factor, β , curvature of utility of consumption, γ , disutility of labor supply, α , and Frisch elasticity of labor supply, ε . Households supply labor and capital in competitive markets and solve a standard consumption-savings problem. Their budget constraint is:

$$C_t + K_{t+1} = (1 + r_t - \delta)K_t + (1 - \tau_t)w_t L_t + TR_t \quad (2)$$

where K_t is the household's capital stock, r_t and w_t are the competitive rental rate of capital and wage rates, τ_t is the effective tax rate on labor income, and TR_t represents government transfers. We will define the effective tax rate below. From the household's perspective, τ_t is taken as given and represents the fraction of labor income taken by the government. The household also takes transfers as given.

The production function for the final good—which can be used for either consumption or investment—

is:

$$Y_t = F(K_t, K_{G,t}, L_{M,t}) = A_M (\theta K_{G,t})^\zeta K_t^\mu L_{M,t}^{1-\mu}, \quad (3)$$

where $K_{G,t}$ and $L_{M,t}$ are inputs of public capital and labor, A_M is the country's exogenous productivity level in final goods, and θ is the country's exogenous efficiency in turning public capital investments into real public capital. Parameter μ represents private capital's share in production and is common to all countries. The subscript M denotes labor supplied to this modern sector as distinct from labor employed in the government's tax administration sector (or, in an extension to follow, a traditional sector).

As in [Baxter and King \(1993\)](#), there are constant returns to scale in privately provided inputs: K_t and $L_{M,t}$. The power ζ governs the output elasticity of public capital, holding other inputs fixed. So long as ζ is positive, there are increasing returns to scale in all inputs (private and public combined).

The tax collection technology is given by:

$$\tau_t = G(L_{C,t}^i) = A_C^i (L_{C,t}^i)^\eta, \quad (4)$$

where A_C^i is exogenous tax productivity in country i , $L_{C,t}^i$ is labor input into tax collection (e.g. the number of "tax collectors"), and $0 < \eta \leq 1$ governs the returns to scale in tax collection. If $\eta = 1$, there are constant returns to scale in tax collection, and each additional tax collector raises the effective tax rate by as much as the last. If $\eta < 1$, there will be decreasing returns to scale, and each additional tax collector will prove less effective than prior ones. Later we argue that the latter is the empirically relevant case, and that it materially affects the paper's main conclusions.

For simplicity, we assume that governments exogenously decide the number of tax collectors each period and these choices are fully known by the households. Governments thus move first in determining the full schedule of tax rates each period, and households respond by solving their dynamic problems taking government behavior as given, as in a Stackelberg game. In the extended model to follow, we let the government's tax-collection inputs be chosen endogenously, though at some cost in transparency relative to this simpler setup.

The government budget constraint each period is:

$$TR_t + X_t + K_{G,t+1} = T_t + (1 - \delta_G)K_{G,t}. \quad (5)$$

where T_t represents total tax collection:

$$T_t = \tau_t w_t (L_{M,t} + L_{C,t}^i), \quad (6)$$

and X_t represents expenditures on tax collections:

$$X_t = w_t L_{C,t}. \quad (7)$$

The government thus uses its current resources each period to finance transfers, payments to tax collectors, and public capital for the next period. Its current resources are tax revenues from market labor and government tax collectors, plus the current undepreciated stock of public capital. Governments exogenously split their gross tax revenues between transfers and public capital investments:

$$TR_t = \xi T_t, \quad (8)$$

implying that:

$$K_{G,t+1} = w_t \left[(1 - \xi) \tau_t (L_{M,t} + L_{C,t}) - L_{C,t} \right] + (1 - \delta_G) K_{G,t}. \quad (9)$$

Economies thus exogenously differ in four parameters: final-goods productivity, A_M , tax productivity, A_C , public spending effectiveness, θ , and the transfer share of gross tax revenues, ξ . They also differ in their number of tax collectors, L_C , each period.

3.2 Equilibrium and Tax Collections per Expenditure

Our equilibrium concept is a sequence-of-markets equilibrium in which households solve their dynamic problem taking prices, tax rates, and transfers as given. Firms solve their static problems taking prices as given and all markets clear. Tax collections per expenditure in equilibrium are:

$$T_t/X_t = A_C(L_{C,t})^\eta (L_{M,t} + L_{C,t})/L_{C,t}. \quad (10)$$

There is a direct relationship, in other words, between tax collections relative to expenditures, the statistic we reported in the previous section, and tax productivity, A_C , in the model. Both are increasing in the wage rate of the economy, but these effects cancel out in the ratio. Labor in the modern sector depends on the level of income and the effective tax rate in equilibrium, and these depend endogenously on other parameters of the model.

The baseline model assumes that all labor income is taxable. In practice, however, some workers are employed outside the tax base, which we will call the traditional sector. By raising productivity

in the modern sector, public investment can increase the relative attractiveness of taxable employment and draw workers out of these untaxed activities. This labor reallocation will further increase the tax base. We explore this additional channel in Section 6.2.

3.3 Steady State

We characterize the steady state for a single economy. The effective tax rate is:

$$\tau = A_C(L_C)^\eta, \quad (11)$$

which will be positive so long as A_C and L_C are positive but could be greater than one in principle. In our quantitative work to follow, tax rates stay strictly between zero and one. A steady state $(C, K, K_G, L_M, Y, w, r, \tau)$ satisfies the household optimality conditions, government budget constraint, and resource constraint in steady state:

$$1 = \beta(1 + r - \delta), \quad (12)$$

$$\alpha(L_M + L_C)^{1/\varepsilon} = C^{-\gamma}(1 - \tau)w, \quad (13)$$

$$\delta_G K_G = w \left[(1 - \xi)\tau(L_M + L_C) - L_C \right], \quad (14)$$

$$C + \delta K + \delta_G K_G = Y, \quad (15)$$

where wage and rental rates satisfy the firm optimality conditions in steady state:

$$Y = A_M(\theta K_G)^\zeta K^\mu L_M^{1-\mu}, \quad (16)$$

$$r = \mu A_M(\theta K_G)^\zeta K^{\mu-1} L_M^{1-\mu}, \quad (17)$$

$$w = (1 - \mu) A_M(\theta K_G)^\zeta K^\mu L_M^{-\mu}. \quad (18)$$

Under mild conditions, we can show that the steady-state equilibrium exists and is unique.

Proposition 1. *If $\gamma > 1$ and $\mu + \zeta < 1$, then there exists a unique steady-state equilibrium.*

Proof. See Appendix A.1. □

Public capital raises output and hence taxable income, contributing to further increases in tax revenue and public investment. This feedback could support different steady-states for economies with identical fundamentals. Proposition 1 rules out this multiplicity when private and public capital jointly exhibit diminishing returns ($\mu + \zeta < 1$) and income effects on labor supply dominate

the incentive from higher wages ($\gamma > 1$). To understand uniqueness, consider candidate steady states with greater employment in production, holding tax productivity and tax-collection employment fixed. When $\mu + \zeta < 1$, the larger tax base and associated capital accumulation raise wages, while consumption rises more than proportionally to wages. With $\gamma > 1$, the resulting decline in the marginal utility of consumption outweighs the higher wage, reducing the utility benefit of an additional hour of work. Meanwhile, greater employment raises the marginal disutility of working. These opposing movements imply that the household's labor-supply condition can hold at only one employment level, which uniquely determines the remaining steady-state quantities.

Outside these conditions, multiplicity can arise. In Appendix A.1 and Figure A.1, we provide an example with two steady-state equilibria.

4 Quantifying Variation in Tax Productivity

We now draw on our cross-country evidence on tax collections relative to tax administration costs and other salient statistics to recover tax productivity in each country. To do so, we parameterize the model in the 98 countries for which we have complete data.

4.1 General Parameters

We begin by assigning values to general parameters that apply to all countries. We deliberately take standard choices wherever they are available. Each period represents one year and thus the discount factor, β , is set at 0.96. The depreciation rates of physical and public capital, δ and δ_G , are taken to be 0.05. Following the evidence of Gollin (2002), we assume that capital's share in the production function is 0.33, consistent with a labor share that is roughly two-thirds across countries of all income levels.

We pick a Frisch elasticity of labor supply, ε , of 0.5, which is somewhere between the smaller values preferred by microeconomic estimates and the larger values favored by most macroeconomic estimates. It is also the long-run value favored in the recent study by Kleven et al. (2025). We normalize the disutility weight on labor supply, α , to 1, since it affects only the level of labor supply, which is not relevant to our analysis. We pick a coefficient of relative risk aversion, γ , of 1.1, which implies that income effects are modestly stronger than substitution effects in labor supply decisions. This is the value Bick et al. (2022) use, for example, to match cross-country differences in hours worked and the decreasing patterns of hours with development.

The remaining two general parameters are somewhat less standard in macroeconomics and, as we

show below, quite important for our quantitative conclusions. These are: the returns to scale in the tax administration sector, η , and the output elasticity of public capital, ζ . We discuss each of these in turn.

Table 3: General Parameters

Description	Parameter	Value
Discount factor	β	0.96
Risk aversion / inverse IES	γ	1.10
Frisch elasticity	ε	0.50
Labor disutility weight	α	1.00
Modern-production private-capital share	μ	0.33
Output elasticity of public capital	ζ	0.20
Tax-administration curvature	η	0.05
Private capital depreciation	δ	0.05
Public capital depreciation	δ_G	0.05

Notes: Country-specific θ is 0.58 for LIC, 0.74 for both LMC and UMC, and 0.84 for HIC. The transfer share is fixed to the prepared excluded-expenditure share.

We argue that η is likely to be far below one, and we set $\eta = 0.05$ in our quantitative analysis based on our reading of the evidence. In the 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). The evidence points to marginal products well below average products, and 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 the self-employed (Kleven, Knudsen, Kreiner, Pedersen and Saez, 2011), larger firms are easier to tax than smaller firms (Pomeranz, 2015; Kleven et al., 2016), and consumers in large, modern retail stores are more easily taxed on their expenditure than consumers in small, traditional stores (Bachas et al., 2024a).⁹

Average products of tax collection come directly from our data on T/X . Direct evidence for the marginal cost of collection is harder to come by, but can be proxied using several recent micro studies. In many low-income countries, property taxes are an important new margin that governments are trying to pursue to increase tax revenues. In Ghana, Dzansi et al. (2025) estimate that the ratio of collections to revenue collector salaries across local governments is 1.56. They find

⁹Many 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. Manuals for practitioners explicitly justify this policy design as driven by administrative cost-benefit considerations: in *The Modern VAT*, Ebrill et al. (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.”

similar numbers in an RCT within one large local government that experimentally increased both collections and expenses to hire tax collectors, which implied a marginal collection per dollar of expenses of 1.62. Our cross-country data for Ghana show that total taxes relative to total tax administration expenses equals 43 (see Figure 1). This implies a ratio of marginal to average of 0.04 ($=1.6/43$).

In a middle-income country, Costa Rica, [Brockmeyer et al. \(2019\)](#) uses an experiment to estimate the impact of increased enforcement on nil filers (registered taxpayers who declare a value of 0 for their taxable activities). The paper finds each additional dollar of enforcement expenditure leads to around \$4 in additional tax revenue. Our cross-country data report average collections per unit of expense in Costa Rica of 82, which implies that the ratio of marginal to average returns is 0.05 ($=4/82$).

Finally, evidence from a high-income country can be found in [Boning et al. \(2025\)](#), which estimates the marginal collections from random audits of individual income taxpayers in the United States. Randomized audits are not the only margin the IRS can use to increase tax collections, but they are likely one of its most important margins. The authors find that each additional dollar spent on audits generates \$13 of tax payments. With an average collection of 260 in the United States in our cross-country data, the ratio of marginal to average returns is 0.05 ($=13/260$). Our choice of a constant $\eta = 0.05$ is consistent with these three pieces of evidence, drawn from low-, middle-, and high-income countries, which point to remarkably similar ratios.

The elasticity of output to public capital, ζ , has received more attention in the literature. However, no clear consensus has emerged on its value—especially in a cross-country setting and for a broad notion of public capital. In our main analysis, we choose $\zeta = 0.2$, but we later test sensitivity to lower and higher values. Our read of the literature is that the wide range of estimates for ζ reflects both difficulties in estimation as well as the natural variation in the actual value of public spending, depending on the project and location. [Aschauer \(1989\)](#) estimated a ζ of 0.39 in his famous study of U.S. infrastructure spending. Others found much lower values for transportation infrastructure, and [Ramey \(2021\)](#) favors 0.05 in her analysis. [Cubas \(2020\)](#) uses national income accounts data to infer a value of 0.09.

The World Bank draws on estimates of ζ to discipline its Long-Term Growth Model (see [Devadas and Pennings, 2018](#)), drawing on a well-known meta-analysis by [Bom and Ligthart \(2014\)](#) based on studies that mostly predate the credibility revolution in economics. Figure A.3 plots the estimated elasticities from the meta-analysis; our benchmark value of 0.20 falls well within their wide range. One potential limitation of these studies for our purposes is that 93 percent of them come from advanced economies. Another is that they are primarily focused on transportation infrastructure. To the extent that public capital represents broader types of public services—like functioning

court systems or police departments—these estimates could significantly underestimate (or overestimate) the true return to public capital.

Several more recent studies in less developed settings are informative about the output elasticity of public capital. However, these studies also do not point to a consensus value 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 close to zero. Both of these studies are again focused on transportation infrastructure. Taken together, they reinforce the need to consider a range of plausible ζ values.

4.2 Cross-Country Parameterization and Predictions

We treat the other parameters as country-specific. One of these is θ , which governs the fraction of public investment that contributes to production. 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 drawn from the World Bank’s Infrastructure Efficiency Index (IEI); see Table 2 in [Devadas and Pennings \(2018\)](#).

Four parameters remain to be set for each country: general productivity, A_M , tax productivity, A_C , the fraction of public revenues that are redistributed as transfers, ξ , and the size of the tax collection sector, L_C . We target these jointly to match four targeted moments in each country: the tax-to-GDP ratio, the T/X ratio from the empirical section above, the transfer share of public expenditures, and tax collection employment as a fraction of total employment.

Figure 2 plots each of the targeted moments in the model and data against GDP per capita. The model generally succeeds in matching the increasing patterns in taxes to GDP (Panel A), taxes per unit of administrative expenditure (Panel B), and the transfer share of public expenditures (Panel C). Some lower-income countries are outliers in taxes to GDP, such as Mozambique and Lesotho, and these are not well captured by the model. These deviations may reflect revenue sources our model omits, such as offshore gas leases in Mozambique and Lesotho’s unique customs union with its richer neighbor, South Africa. Tax collectors as a fraction of total employment (Panel D) overshoots the data for most countries, though the magnitudes are small throughout and well below one percent of all employment in both the model and data. Both are modestly increasing in income

Figure 2: Model's Predictions for Targeted Moments



Notes: This figure plots the model's predicted tax-to-GDP ratios (Panel A), taxes per unit of expenditure on collection (Panel B), share of tax revenues used for transfers rather than public capital (Panel C), and employment share of tax collectors (Panel D). In each graph, the red dots represent the model's predictions, and the blue dots represent their empirical counterparts. Total taxes collected, tax administration expenditures, and the size of the tax collection sector by employment come from our own calculations described in the text. The transfer share comes from the [International Monetary Fund \(2026a\)](#).

per capita.¹⁰


The model also does reasonably well in matching cross-country moments that are not targeted in the calibration. Figure 3 plots the model's predictions for private capital-labor ratios, public capital-labor ratios, investment shares of GDP, and labor supply (relative to the world average). Private capital-labor ratios are readily observable in the data, and we take these directly from the Penn World Tables. Panel A shows that our model gets an almost identical slope in income per capita. This is reassuring, though perhaps not terribly surprising, given the successes of the neo-classical growth model in this dimension. Public capital-labor ratios are more difficult to measure in practice, though they can be proxied by investment data from the IMF's Government Financial Statistics. Panel B plots the IMF's public capital per worker, scaled by our θ values. In practice, the real mapping between nominal public expenditure and real government capital stocks could be far lower or higher, and the choice of only three θ values to cover all low-, middle-, and high-income countries is obviously a crude one. It is still useful to verify that in the model and data, public capital stocks per worker rise nearly linearly in income per capita.

Panel C reports the investment share of GDP at domestic prices, again taken from the Penn World Tables. Both in the data and model these are modestly increasing in GDP per capita, though our model over-predicts the level of investment to GDP on average. Panel D plots average hours worked per adult from Bick et al. (2018), expressed relative to the world average. Both in the model and data, hours exhibit a modest decline in income per capita. Matching the long-run decline in hours is reassuring, particularly in light of the labor-supply responses we report below. It is also expected, given that our calibration of preferences makes income effects modestly dominate substitution effects.

One of the goals of the parameterization was to recover the values of A_C and A_M needed to match the cross-country evidence. Figure 4 plots these for each of the 98 countries. Two patterns can be seen from these two plots. First, A_C exhibits a steeper slope with GDP per capita (Panel A) than does A_M (Panel B). Second, there is more variation conditional on income for A_C than for A_M . We are mostly interested in the first pattern from a development perspective. A regression of $\log A_C$ on $\log$ GDP per capita yields a slope of 0.797 (s.e. 0.074). For A_M , the slope coefficient is 0.344 (s.e. 0.015). For comparison, the slope coefficient of the Penn World Tables' measure of TFP is 0.273 (s.e. 0.018), similar to the slope in A_M . The upshot is that A_C varies more with development than A_M , and our best estimate is that A_C varies around twice as much with income per head as a more general measure of TFP.


¹⁰In the sensitivity analysis below, we explore alternative parameterizations that match either T/Y exactly or T/X exactly. These yield very similar predictions for the long-run effects of raising A_C .

Figure 3: Model's Predictions for Non-Targeted Moments



Notes: This figure plots the model's predicted capital-labor ratios (Panel A), public capital-labor ratios (Panel B), investment shares of GDP (Panel C), and labor supply (Panel D). In each graph, the red dots represent the model's predictions, and the blue dots represent their empirical counterparts. GDP per capita, employment, and private capital stocks come from the Penn World Tables version 11. Public capital stocks come from [International Monetary Fund \(2021\)](#). Labor supply is proxied by hours worked per adult, from [Bick et al. \(2018\)](#).

Figure 4: Tax Productivity and Aggregate Productivity



Notes: This figure plots the model’s predicted tax productivity levels (Panel A) and aggregate productivity terms (Panel B) in each country. Both series are estimated using the model and data described in the text.

4.3 Empirical Investigation of A_C

Our model predicts large differences across countries in A_C . What determines this variation in tax productivity? We leverage a dataset compiled by ISORA, which is a unique source of information on the current practices of tax administrations around the world. We focus on three sets of activities that could determine tax productivity, denoted x : (i) the extent of use of third-party data and digitization; (ii) enforcement processes; (iii) operational processes.¹¹ We use this data to investigate if real activities x are correlated with our model-based measure A_C . We also investigate if these x characteristics predict A_M . If the x factors are associated with A_C but not A_M , this would support the interpretation of A_C as a distinct measure of tax-collection productivity rather than a simple proxy for a country’s general productivity or technological advancement. We use the cross-sectional estimates of A_{Ci} in all countries i and estimate the equation: $A_{Ci} = \beta + \theta \cdot x_i + \varepsilon_i$, with robust standard errors.

In Panel A of Table 4, we estimate simple bivariate regressions of A_C against selected x activities.¹² Columns (1) to (6) focus on data and digitization. In column (1), A_C is positively correlated with the number of data sources where tax authorities use fully computerized systems to process

¹¹Panel A of Table A.1 shows that results are robust to using an alternative data source, the OECD’s Inventory of Tax Technology Initiatives from Organisation for Economic Co-operation and Development (2024), and the classification of technology usage by Okunogbe and Tourek (2024).

¹²Relative to our initial sample ($N = 98$; Section 2.1), Taiwan is dropped due to missing values for other variables.

third-party information, including data from banks, employers, and digital sales platforms. The average tax authority integrates just under 3 such sources (out of a potential of 10) through fully computerized systems.¹³ In column (2), A_C is also positively correlated with pre-filing, whereby tax authorities use third-party data to populate taxpayers' returns before filing. Effective pre-filing requires accurate and comprehensive third-party data, as well as the technological capacity to match it to the correct taxpayer. 57% of countries in our sample use pre-filing; with the US a notable exception, its use is strongly correlated with GDP per capita.

Columns (3) to (6) capture whether the tax authority: (i) uses robotic process automation to automate repetitive, rule-based tasks (47% of countries); (ii) uses data science analytics, such as machine learning and big data, to process tax returns and information forms (84%); (iii) uses whole-of-government identification, whereby a unique identifier links individuals or businesses across agencies (62%)¹⁴; and (iv) uses Artificial Intelligence (61%). These digital technologies vary substantially across countries, and their use is robustly correlated with tax productivity.

Columns (7) to (9) study enforcement. We focus on enforcement processes that can be intensified without directly increasing administrative expenses, as well as activities that are more complex from a technical or organizational perspective and indicative of stronger capacity. In column (7), the ability to garnish wages and other payments owed to taxpayers is positively associated with A_C ; 92% of countries use this tool. In columns (8) and (9), tax productivity is also positively correlated with the ability to offset tax arrears against overpayments or credits on other tax bases (96%) and to collect unpaid taxes through cross-border agreements with foreign tax authorities (67%).

Columns (10) and (11) illustrate compliance-oriented operational practices. Column (10) shows that automated compliance checks without human intervention (adopted in 47% of countries) are associated with higher A_C . Column (11) shows a similar relationship for randomized audits for research and analytical purposes rather than as a deterrent (18%). Such audits can improve operations by identifying compliance gaps, improving risk profiles, estimating tax gaps, and measuring behavioral effects.

Panel A focuses on selected activities covered in ISORA in the three areas of data and digitization, enforcement, and compliance processes. The first four columns of Panel B show that the positive association with A_C holds when we construct indices for all activities covered by ISORA in these

¹³In the case where x_i is initially a continuous variable, we report results based on the standardized x_i .

¹⁴For a tax authority, this can integrate tax files with social security, immigration, business registries, and customs

three areas¹⁵, and when we use a single index for all activities.¹⁶

The next set of columns (5 to 10) in Panel B address the possibility that these positive associations simply reflect the correlation between a country’s overall productivity, A_M , and its general country-level determinants, denoted D_i . We collect proxies for factors that prior work links to overall productivity and technological advancement: legal institutions (property rights), social institutions (trust), digitization (digital penetration), management quality, human capital (primary education), openness (cross-border trade), structural composition (employee share and agricultural share), and financial development (private credit as a share of GDP). Table 4 provides more details on the D_i variables. We include the full set D_i as controls, which reduces the sample size.¹⁷ Columns 5 to 8 show that the associations between tax productivity A_C and x_i hold when controlling for D_i . In column (9), A_C remains correlated with x_i when we directly control for A_M .

Finally, in column (10), the outcome is A_M , and we relate it to both x_i and D_i . This regression shows that A_M is strongly correlated with general productivity determinants D_i , unsurprisingly; at the same time, A_M is uncorrelated with the tax productivity determinants x_i .

The results in Table 4 support the interpretation that A_C is a distinct measure of tax productivity, rather than simply a proxy for overall productivity. These results suggest that countries can increase tax productivity through technological, organizational, and process-related reforms.¹⁸ Although correlational and aggregate, our findings are consistent with rigorous within-country evidence on similar reforms; see Okunogbe and Tourek (2024) and Jensen and Weigel (2026). Thus, while broader factors may affect both A_C and A_M (Besley and Persson, 2011), our results indicate that A_C can also be impacted by productivity-enhancing reforms available to governments.¹⁹ Based on this validation exercise, we now turn to simulating the economic effects of increasing A_C .

¹⁵Panel B of Table A.1 shows that A_C is uncorrelated with the remaining categories covered in ISORA: human resources (HR), customer services, institutional autonomy, and the existence of dedicated units for large taxpayers. The HR section covers topics such as whether there is a multi-year workforce plan in place, whether there are policies for flexible work arrangements, and whether there are specific leadership and talent management programs. The autonomy section covers topics such as whether the tax authority has the discretion to determine its own internal structure, allocate its budget, make and terminate staff appointments, place staff in salary ranges, and establish performance criteria. The taxpayer services section covers activities that reduce compliance costs, enhance communication, and increase accountability and transparency.

¹⁶We standardize activities so that higher values always indicate better practices.

¹⁷Panel C in Table A.1 shows that results are robust to imputing missing values for D_i to estimate the relationship in the full sample of countries with non-missing A_C .

¹⁸The null results in Panel B of Table A.1 merit further discussion. While stronger HR policies and greater autonomy could increase tax productivity, our estimates may reflect either limited effects or insufficient power. Similarly, although organizing enforcement by taxpayer size may raise productivity (see Basri et al. (2021) for evidence from Indonesia), we find no strong cross-country correlation.

¹⁹Indeed, the tax-specific determinants x_i have a larger association with tax productivity than any general productivity determinant D_i (column 5, Panel B, Table 4), while 6 of 9 D_i variables are insignificant. Moreover, x_i remains strongly correlated with GDP per capita after controlling for the D_i variables, and the D_i variables do not themselves significantly predict x_i (results available upon request).

Table 4: Tax Authority Activities and Tax Productivity A_C

Panel A: Selected activities of tax authorities within each category											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Category: 3 rd -party data and digitization											
Computer-based 3 rd -party data integration		0.449*** (0.083)									
1(Pre-filing)		1.021*** (0.173)									
1(Robotics process automation)			0.952*** (0.183)								
1(Data science analytics)				0.944*** (0.188)							
1(Whole-of-government identification)					0.676*** (0.185)						
1(Artificial intelligence)						0.791*** (0.184)					
Category: Enforcement											
1(Garnishment)							1.041*** (0.236)				
1(Offset arrears across domestic tax base)								0.899*** (0.318)			
1(Offset arrears across international tax authorities)									0.460** (0.216)		
Category: Operational processes											
Automated compliance checks										0.544*** (0.112)	
Use of randomized audits for research and risk profiling											0.277* (0.156)
Controls	No	No	No	No	No	No	No	No	No	No	No
Dependent variable	A _C	A _C	A _C	A _C	A _C	A _C	A _C	A _C	A _C	A _C	A _C
Observations	97	97	97	97	97	97	97	97	97	97	97
R ²	0.149	0.256	0.226	0.123	0.108	0.149	0.082	0.032	0.047	0.217	0.031
Panel B: Composite indices combining all activities within each category											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	
3 rd -party data and digitization	0.579*** (0.118)				0.214 (0.130)						
Enforcement		0.390*** (0.104)				0.221** (0.103)					
Operational processes			0.421*** (0.111)				0.207* (0.105)				
Combined index				0.716*** (0.168)				0.348** (0.161)	0.535*** (0.171)	-0.025 (0.128)	
Controls	No	No	No	No	D _i	D _i	D _i	D _i	A _M	D _i	
Dependent variable	A _C	A _C	A _C	A _C	A _C	A _C	A _C	A _C	A _C	A _M	
Observations	97	97	97	97	69	69	69	69	97	69	
R ²	0.182	0.109	0.131	0.219	0.588	0.602	0.594	0.605	0.324	0.781	

Notes: The dependent variables are standardized productivity levels, not logarithms. Panel A reports coefficients for selected individual ISORA activities within three categories: 3rd-party data and digitization; enforcement; and operational processes. All regressors are standardized, except for dummy variables. Panel B reports coefficients for category-level indices constructed using all activities within each category, including those not reported individually in Panel A. Each category index is the unweighted average of the standardized indices for all individual activities in that category. The “combined index” is the unweighted average of the three category indices. In the D_i variables, we: proxy for legal institutions using the measure of quality of the legal system and property rights (source: [Fraser Institute, 2026](#)); proxy for social institutions using the share of respondents that agree with the statement “most people can be trusted” ([World Values Survey Association, 2026](#)); proxy for human capital using the share of adults aged 25+ with at least primary education ([UNESCO Institute for Statistics, 2026](#); [World Bank, 2026b](#)); proxy for trade openness using exports plus imports as a percent of GDP ([International Monetary Fund, 2026b](#)); proxy for structural composition using the employee share of the workforce and agricultural value added as a share of GDP; proxy for financial development using the ratio of private credit to GDP; proxy for digitization using the World Bank’s digital adoption index ([World Bank, 2019](#)); proxy for management quality using firms’ reliance on merit-based professional management rather than personal connections ([Schwab, ed, 2017](#)). All D_i variables are averaged over multiple years between 2014 and 2016 and subsequently standardized. Taiwan is excluded because its ISORA predictors are missing. Robust standard errors are reported. * p<0.10 ** p<0.05 *** p<0.01.

5 Quantifying the Effects of Tax Productivity Shocks

This section studies the quantitative role of tax policy and, in particular, tax productivity shocks. In our model, higher tax productivity raises long-run output and consumption in low-income countries. The output gains far exceed a first-order growth-accounting benchmark (Hulten’s Theorem) based on the tax collection sector size. Finally, we show analytically and numerically how additional tax revenue finances public capital and induces private capital accumulation, multiplying the effects of higher tax productivity.

5.1 Quantitative Counterfactuals

Once the model is calibrated country by country, our main quantitative exercise simulates the long-run general equilibrium effects of raising tax productivity, or A_C . We compute the effects in particular of an increase in A_C of 50 percent, which, we argue, should be considered a transformative change in how efficient tax collection is. We note 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 et al., 2025), better databases of property owners linked to hand-held tablets with GIS capabilities roughly doubled tax collections in Ghana. Although these results do not literally map into changes in A_C , they suggest that large changes in the productivity of tax collectors are not purely hypothetical.


Figure 5 reports the changes in GDP and other key macroeconomic variables in the new steady state resulting from a permanent 50 percent increase in A_C . The red bars report the average percent changes across all low-income countries. GDP increases by 18.42 percent on average in low-income countries, compared with 14.74 percent in middle-income countries and 8.62 percent in high-income countries. Consumption increases by 17.98 percent, 9.58 percent, and 0.65 percent, respectively. Thus, the consumption effects are large and positive in the poorest countries but decreasing in the richest.

Why? To understand the mechanics of the model’s predictions, note that an increase in A_C raises effective tax rates by 50 percent in all economies, since labor in the tax collection sector is fixed: $\bar{\tau} = A_C \bar{L}_C^\eta$. From the government budget, one can see that the higher tax rate raises K_G directly, and further raises it through a higher equilibrium wage; it is at least partly offset as higher taxes reduce labor supply:

$$K_G = \delta_G^{-1} w_M [(1 - \xi) \bar{\tau} (L_M + \bar{L}_C) - \bar{L}_C]. \quad (19)$$

From the production function, $Y_M = A_M (\theta K_G)^\zeta K^\mu L_M^{1-\mu}$, one can see that output is affected directly

Figure 5: Long-Run Effects of Raising Tax Collection Productivity



Notes: This figure reports the effects of a permanent 50-percent increase in tax productivity, A_C , holding tax-collection labor fixed. Bars show unweighted country-mean percentage changes, except the pre- and post-shock tax bars, which show effective tax-rate levels in percent. Colors represent low-income countries (red), middle-income countries (green), and high-income countries (blue), using World Bank country categories.

through increases in public capital, with a short-run output elasticity of public capital equal to ζ . Output is also affected indirectly through changes in private capital and changes in labor supply. From the Euler equation, the rental rate of capital, r , depends only on the rate of time preference and the depreciation rate in steady state, and is hence constant. The marginal product of capital, and capital-output ratios, are thus also constant in steady state: $r = \mu Y_M / K$. Any changes in output are thus the same in percentage terms as the changes in private capital. The wage rate and marginal product of labor increase in output and decrease in labor: $w_M = (1 - \mu) Y_M / L_M$.

From the labor-leisure household optimality condition, $L_M^{1/\varepsilon} = C^{-\gamma} (1 - \bar{\tau}) w_M$, higher wage rates raise labor supply while higher tax rates lower it. Consumption depends on after-tax wages and labor supply: $c = [(1 - \bar{\tau}) w_M / L_M^{1/\varepsilon}]^{1/\gamma}$.

With these equations in mind, one can see more clearly how the low-income countries benefit more from increases in tax productivity in these counterfactual simulations. In low-income countries, the gains in public capital are 79.98 percent. These drive an 18.42 percent increase in private capital, which is the same as the GDP increase. Wages increase by 19.17 percent, benefiting workers. The starting tax rate is so low that a 50 percent increase in taxes takes the low-income economies from a tax rate of 1.22 percent to 1.83 percent. After-tax wages rise by 18.43 percent, which certainly still leaves workers better off. Labor supply falls by just 0.63 percent. The resulting change in consumption is 17.98 percent.

The calculation is different in middle-income and high-income countries. Public capital increases substantially in those countries too, at 73.20 and 63.78 percent, respectively. Private capital stocks rise, but by less, at 14.74 and 8.62 percent, with wages rising significantly in both cases. The downside is that tax rates rise from 16.9 percent to 25.28 percent in middle-income countries and from an already robust 29.86 percent to 44.78 percent in the high-income ones. This leads to declines in after-tax wages in the richest countries. This explains the more modest 9.58 percent increase in consumption in middle-income nations and the 0.65 percent increase in high-income countries.

In reality, middle-income and rich countries would likely reduce the sizes of their tax collection sectors in response to an increase in tax productivity. This would in turn keep effective tax rates closer to their original levels. In that case, the consumption and labor-supply changes would be closer to zero than these simulations show. The point is still that lower-income countries respond quite differently to tax productivity changes than do richer economies. These endogenous changes are discussed further below, particularly in the extension to Ramsey taxation.

Table 5 puts the long-run GDP changes for low-income countries in perspective by considering Hulten's approximation as a simple benchmark. Hulten's approximation says that the aggregate

Table 5: Model Predictions versus Hulten Accounting in Low-Income Economies

Prediction method	Δ GDP (%)
Model output	18.42
Hulten wage-bill term: $0.5w_M L_C/Y$	0.01
Hulten tax-revenue term: $0.5T/Y$	0.41
Public-capital contribution: $\zeta \% \Delta K_G$	16.00

Notes: All four rows use the same 10 solved low-income countries. Hulten terms use baseline shares. Overbars denote unweighted country means. The public-capital row holds other production inputs fixed; it is not the full equilibrium response.

impact of a 50 percent TFP increase in a sector is roughly equal to 50 percent times the sector’s Domar weight (defined as its sales relative to GDP). The tax collection sector does not have sales per se, so we consider its wage bill share and also total tax collections to GDP as weights.

In both cases, the model’s GDP gain of 18.42 percent far exceeds the simple Hulten approximation. Using the wage bill share, the A_C gains should boost the aggregate economy by a minuscule 0.01 percent. With the tax-to-GDP ratio as a weight, the Hulten approach predicts a 0.41 percent GDP increase. The model’s prediction is still an order of magnitude larger. The fourth row isolates the direct public-capital channel: a 79.98 percent increase in K_G scaled by the output elasticity $\zeta = 0.2$ yields a 16 percent GDP gain, most of the model’s total. The rest comes from the induced rise in private capital and change in labor supply.

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 public and private investment illustrated above in Figure 5. Hulten’s theorem presumes efficient allocation without externalities; here public capital acts as a production “externality” financed by distortionary taxation, so the first-order approximation misses the reallocation toward public and private capital.

For additional perspective, Table 6 reports the effects of other counterfactual changes one could imagine in low-income countries. The first column reproduces the effects of a fifty percent increase in A_C (reported as a fraction of baseline, for expositional purposes). The second column simulates raising the efficiency of public expenditure from 0.58 up to 0.84, the level computed for high-income countries by the World Bank. This would result in a 16 percent increase in GDP, both public and private capital stocks, and consumption. The tax rate and T/X ratio would not change in any important way. In all other aspects, this experiment looks remarkably like the higher A_C scenario. The takeaway is that low-income countries can, in principle, achieve the same gains from improving their tax collection capacity by investing their current tax proceeds more effectively.

Low-income countries also do not need better tax productivity to increase tax revenues and public

Table 6: Effects of Other Counterfactuals on Low-Income Countries

Variable	$A_C \times 1.5$	$\theta = 0.84$	$L_C \times 1.5$	$A_M \times 1.5$
GDP	1.18	1.16	1.00	2.28
Private capital	1.18	1.16	1.00	2.28
Public capital	1.80	1.16	1.00	2.28
Consumption	1.18	1.16	1.00	2.28
Transfers	1.78	1.16	1.02	2.28
After-tax wage	1.18	1.17	1.00	2.34
Modern labor L_M	0.99	0.99	1.00	0.97
T/X	1.49	1.00	0.68	0.97
Tax rate τ	1.50	1.00	1.02	1.00

Notes: Each entry is the unweighted mean of country-specific counterfactual/baseline ratios. Collection labor is fixed except in its own shock.

capital: they can simply hire more tax collectors using the current systems. The third column of Table 6 simulates raising the number of tax collectors by fifty percent. Unfortunately, the model predicts virtually no gains from such a change. Tax rates rise only modestly due to the strong decreasing returns to scale in that activity. Intuitively, this would be like sending more tax collectors into areas with self-employed family farms and businesses, or, more generally, small firms. In the model, public capital increases by around 0.1 percent and GDP increases by even less. The tax-to-expenditure ratio falls significantly to 68 percent of the previous level, as expenditures rise significantly while tax collections increase only very modestly.

The final column of Table 6 simulates raising A_M by fifty percent. This is no “policy counterfactual” to be sure: few concrete policy options identified by economists correspond to such massive TFP increases. It is still useful to simulate this scenario as a frame of reference to the similar-sized increases in tax productivity. Countries make enormous gains from such a change, with GDP per capita increasing by a factor of 2.28, along with public and private capital stocks, reflecting neoclassical forces. Labor supply falls (only slightly more than when raising A_C) due to income effects. Tax rates do not change, and the T/X ratio changes only modestly, reflecting the earlier intuition that richer households make for juicier targets for tax collection, but also more expensive tax collectors.

5.2 Where Do The Large Gains from Tax Productivity Shocks Come From?

To see where our result comes from, it helps to start from the simplest version of our model and add each ingredient progressively. Consider the static version of our economy with no private capital. Output is $Y = A_M(\theta K_G)^\zeta L_M$, the tax rate is $\tau = A_C L_C^\eta$, and the government turns all of its revenue

into public capital, $K_G = \tau Y$. Start with inelastic labor supply, so that L_M is fixed. Solving out public capital gives $Y = [A_M \theta^\zeta \tau^\zeta L_M]^{1/(1-\zeta)}$, and therefore

$$\frac{d \log Y}{d \log A_C} = \frac{\zeta}{1 - \zeta}. \quad (20)$$

Hulten's theorem ([Hulten, 1978](#)) would instead assign the tax authority's Domar weight, λ_C . Because the authority sells nothing, that weight is its cost share, $\lambda_C = w L_C / Y$, which is below half of one percent everywhere in our sample. At our benchmark $\zeta = 0.20$, taken from the meta-analysis of [Bom and Ligthart \(2014\)](#), expression (20) equals 0.25, some two to three orders of magnitude larger. In this setting, the elasticity of output to tax productivity depends on ζ only, and there is amplification when $\zeta > 0$. Raising τ sets off the loop familiar from [Barro \(1990\)](#). More revenue buys more public capital, which raises output, which raises revenue at the same tax rate, which buys more public capital again. Summing that loop, $\zeta + \zeta^2 + \zeta^3 + \dots = \zeta / (1 - \zeta)$. A GDP accounting that uses as Domar weight the wage share or tax share does not account for this loop. [Baqae and Farhi \(2019, 2020\)](#) show that the Hulten theorem is an envelope result that holds at an efficient allocation, where prices are the multipliers on the resource constraints, and reallocation has no first-order value. Here public capital is unpriced, and the tax is a wedge, so neither condition holds.

In general, labor supply responds, since a higher tax rate changes how much households choose to work. Differentiating Y with respect to A_C now gives

$$\frac{d \log Y}{d \log A_C} = \frac{\zeta}{1 - \zeta} + \frac{1}{1 - \zeta} \frac{\partial \log L_M}{\partial \log A_C}. \quad (21)$$

The first term is the mechanical channel of (20). The second states that if households work less as a result of taxation, labor falls and output falls with it, amplified by the same factor $1/(1 - \zeta)$. We now show how large this indirect term is and when it varies with the level of development.

Solving the household's problem gives the labor reallocation effect as the sum of three forces. Write $\varphi \equiv (1 - \xi)\tau$ for the share of output that becomes public capital after transfers, and define

$$\mathcal{S} \equiv \frac{\tau}{1 - \tau}, \quad \mathcal{I} \equiv \gamma \frac{\varphi}{1 - \varphi}, \quad \mathcal{W} \equiv (1 - \gamma) \frac{\zeta}{1 - \zeta}, \quad D \equiv \frac{s}{\varepsilon} + \frac{\gamma - \zeta}{1 - \zeta},$$

where $s \equiv L_M / L$ is the share of labor in production. The term $\mathcal{S}$ is the substitution effect. A more productive tax authority raises the effective tax rate, the after-tax wage falls, and households work less. The term $\mathcal{I}$ is the income effect from paying for the public capital. Output diverted into public investment leaves less to consume, the marginal utility of consumption is higher, and households

work more. It is scaled by γ because γ governs how fast marginal utility rises as consumption falls. The term $\mathcal{W}$ is the wealth effect from being richer. Public capital raises output, and with $\gamma > 1$ households take part of that gain as leisure. The denominator D measures how inelastic labor supply is, so that a smaller Frisch elasticity ε or a larger γ shrinks every response.

When there are no transfers, all revenue is invested, $\varphi = \tau$, and the income effect is $\gamma\mathcal{S}$, so that it is only partially offset by the substitution effect and we obtain:

$$\frac{d \log Y}{d \log A_C} = \frac{\zeta}{1 - \zeta} + \frac{1}{1 - \zeta} \cdot \frac{(\gamma - 1)\mathcal{S} + \mathcal{W}}{D}. \quad (22)$$

The net indirect term is a net income effect $(\gamma - 1)\mathcal{S}$, which rises with the tax rate, and a constant wealth effect. Therefore, without transfers, the output elasticity to tax productivity increases as countries become richer and their taxes higher. In a low-tax economy, the extra public capital dominates, consumption rises, and households take part of the gain as leisure, which subtracts from output. In a high-tax economy, the tax bill dominates, consumption is squeezed, and households work more, which adds to it. However, the indirect term is quantitatively small without transfers, and we obtain output elasticities of 0.2413, 0.2523, and 0.2602 for low-, middle-, and high-income countries, respectively.

Once we account for transfers, the output elasticity with respect to tax productivity decreases with development. Once a share ξ of revenue is rebated rather than invested, households are still taxed at τ but public capital is financed only at $\varphi = (1 - \xi)\tau < \tau$. The substitution effect is unchanged, because it depends on wages that are taxed, while the income effect shrinks, because of transfers. The income and substitution effects no longer offset each other, and we obtain the following proposition.

Proposition 2. *For a static economy without private capital, with transfers, and elastic labor supply, the elasticity of output to tax productivity is*

$$\frac{d \log Y}{d \log A_C} = \frac{\zeta}{1 - \zeta} + \frac{1}{1 - \zeta} \cdot \frac{\mathcal{I} - \mathcal{S} + \mathcal{W}}{D}. \quad (23)$$

Setting $\xi = 0$ give the formula (22). Since the transfer share rises with income in our data, the income effect decreases in rich countries. At our group calibrations $\mathcal{I} - \mathcal{S}$ falls from -0.01 in low-income countries to -0.13 in middle-income and -0.28 in high-income ones, labor supply falls by 0.01, 0.05 and 0.10 log points respectively, and (23) takes the values 0.24, 0.19 and 0.13. The elasticity of output to tax productivity now declines with development, and consumption turns negative at the top.

Our full model adds private capital, which accumulates in proportion to output and so raises the multiplier in the first term from $1/(1 - \zeta)$ to $1/(1 - \mu - \zeta)$, a long-run margin emphasized by [Baqae and Malmberg \(2025\)](#). This raises the elasticity from 0.25 to roughly 0.43 but changes nothing about the comparison with Hulten. Appendix [A.5](#) derives all the expressions and proves Proposition 2.

5.2.1 The Role of Decreasing Returns to Tax Collection

Our results above do not depend on the marginal returns to tax collection η . Because $\tau = A_C L_C^\eta$ and L_C is held at its calibrated value, a one percent improvement in tax productivity raises the tax rate by one percent whatever the returns to tax collection are. Once the government is allowed to choose how many collectors to employ, η becomes very important for the elasticity of output to tax productivity, and we obtain the following proposition.

Proposition 3. *For a static economy without private capital, with transfers, elastic labor supply, and optimal tax administration L_C , the elasticity of output to tax productivity is*

$$\frac{d \log Y}{d \log A_C} = \Lambda + (\eta \Lambda - \kappa) \frac{d \log L_C}{d \log A_C}, \quad \kappa \equiv \frac{1}{1 - \zeta} \cdot \frac{(1 - s)/\varepsilon}{D}, \quad (24)$$

where Λ denotes the fixed- L_C elasticity of (23) and κ is the output cost of withdrawing one worker from production.

The final term is determined by the government's first-order condition, which trades η log points of tax rate against the wage bill of the tax collectors required to raise them. Different η now affect the elasticity of output to tax productivity.

Table 7: Effect of η on Elasticity of Output to Tax Productivity

$d \log Y / d \log A_C$	$\eta = 0.05$	$\eta = 0.20$	$\eta = 0.50$	$\eta = 1$
Low income	0.236	0.228	0.112	0.008
Middle income	0.166	0.071	0.014	0.004
High income	0.083	0.029	0.009	0.003

Notes: Elasticity of output to tax productivity in the static economy without private capital, with elastic labor supply and transfers, when the government chooses collection employment L_C to maximize household welfare. Each income group is evaluated at its average tax rate, transfer share, and collection employment share, with $\gamma = 1.1$, $\varepsilon = 0.5$, and $\zeta = 0.20$. Entries are numerical derivatives taken given the re-optimized L_C .

The elasticity of output to tax productivity decreases with η . At $\eta = 0.05$ collectors are very unproductive at producing τ , so the tax rate is effectively pinned by A_C . The government barely adjusts its payroll in response to a change in A_C , ninety-nine percent of the improvement passes through into a higher tax rate, and the low-income elasticity is 0.236 against 0.237 when L_C is fixed. Therefore, endogenizing collection has essentially no effect when $\eta = 0.05$. At $\eta = 1$, the tax rate is linear in L_C , so a log point of productivity substitutes for an equal log point of tax collectors. The government holds the tax rate at its target and saves the improvement as a smaller payroll instead, $d \log L_C / d \log A_C$ approaches -1 , and (24) collapses to κ : output rises only by the value of the collectors who are released back into production.

That limit where $\eta = 1$ is also where Hulten's theorem is restored. Intuitively, at the optimum the government hires collectors until the last one is worth their pay, so the tax authority's cost share values collection at its marginal product. At the same time, an improvement in A_C raises the product of every collector, that is, the average. Since η is the ratio of the marginal to the average product, the two coincide only at $\eta = 1$. The effects of tax productivity shocks beyond the Hulten theorem are therefore not an artifact of holding the tax authority's size fixed. They extend to endogenous collection, but survive only because returns to tax collection are sharply decreasing.

We can now return to the static setting with no private capital and inelastic labor, but with L_C chosen optimally, and recover our Barro (1990)-type result in equation (20).

With labor inelastic, the disutility of work is a constant, so the government simply maximizes consumption $C = (1 - \varphi)Y$. Differentiating $\log C$ with respect to $\log L_C$ gives its first-order condition,

$$\eta(\Lambda_0 - W_c) = \frac{\lambda_C}{1 - \zeta}, \quad \text{with } \Lambda_0 \equiv \frac{\zeta}{1 - \zeta}, \quad W_c \equiv \frac{\varphi}{1 - \varphi}, \quad (25)$$

which equates the value of what a collector produces to what a collector costs. The left side is η log points of tax rate, each worth $\Lambda_0 - W_c$ in consumption, since a higher tax rate buys public capital but takes output to pay for it. The right side is the output forgone by moving that worker out of production, which is the Domar weight scaled by public capital multiplier. The left side falls in L_C and the right side rises, so the optimum is unique, and it always satisfies $\varphi^* < \zeta$: the government never taxes past the point where public capital stops paying for itself. Differentiating (25) through gives the elasticity.

Proposition 4. *For a static economy without private capital, with transfers, inelastic labor supply, and optimal tax administration L_C , the elasticity of output to tax productivity is*

$$\frac{d \log Y}{d \log A_C} = \frac{\zeta}{1 - \zeta} - \frac{\eta W_c R}{\eta R + (\Lambda_0 - W_c)/s}, \quad R \equiv W_c(1 + W_c) = \frac{\varphi}{(1 - \varphi)^2}. \quad (26)$$

As $\eta \rightarrow 0$ the second term vanishes and we recover the Barro (1990)-type result (20), $\zeta/(1 - \zeta)$. As η grows, the expression tends to $\Lambda_0 - W_c$, which itself goes to zero. Table 8 reports the elasticity together with the optimal investment rate φ^* it implies.

Table 8: Effect of η with Inelastic Labor Supply and Optimal Tax Administration

	$\eta = 0.05$	$\eta = 0.20$	$\eta = 0.50$	$\eta = 1$
<i>Elasticity of output to tax productivity</i>				
Low income	0.250	0.247	0.148	0.013
Middle income	0.234	0.136	0.027	0.007
High income	0.209	0.094	0.022	0.007
<i>Optimal public investment rate φ^*, against $\zeta = 0.20$</i>				
Low income	0.027	0.050	0.150	0.195
Middle income	0.141	0.172	0.193	0.198
High income	0.168	0.184	0.194	0.198

Notes: Same economy as Table 7 but with inelastic labor supply. The upper panel reports the elasticity of output to tax productivity from equation (26). The lower panel reports the implied optimal public investment rate $\varphi^* = (1 - \xi)\tau^*$, which always lies below $\zeta = 0.20$, the rate that maximizes consumption.

The second panel of the table is the explanation for the first. As η rises, the government drives the public investment rate toward $\varphi^* = \zeta$, which is the Barro (1990) optimum at which the marginal social value of public capital is zero. Once it is there, a more productive tax authority will not tax more. The government holds φ at ζ and employs fewer collectors instead, and the only gain is the second-order one of releasing them back into production. How close the government can get to that point is exactly what η governs, since η determines how much tax rate a collector delivers relative to the output it costs.

At our benchmark $\eta = 0.05$, the government is far from the optimal investment rate. The optimal investment rate in low-income countries is 0.027 against a target of 0.20, so the economy remains far short of the efficient public capital stock, additional revenue is still valuable at the margin, and the elasticity is 0.25, the same as when L_C is held fixed. At $\eta = 1$, optimal public capital investment rises towards 0.20, and the elasticity of output to tax productivity decreases toward the Hulten benchmark.

In summary, it is the sharply decreasing returns to collection η that keep poor countries far enough from the efficient allocation for tax productivity to matter.

5.3 Efficient Benchmark

Are increases in public capital desirable from an efficiency perspective? To get insight into this question, we follow [Ramey \(2021\)](#) and ask how the competitive outcomes above compare to the efficient allocation that would have been chosen by a benevolent social planner. Following tradition, we shut down markets and allow the planner to direct inputs to production as she chooses, and to allocate output to consumption, investment in private capital, or investment in public capital. In other words, we eliminate the restriction that households supply labor in markets and that the government has to use up resources in order to take goods from the household and turn them into public investment.

We consider a social planner whose tax collection is costless. This social planner maximizes the utility of the representative household subject to the aggregate resource constraint. We can compare the public capital-output ratios of the social planner versus the competitive equilibrium solution in steady-state.

Proposition 5. *In steady-state, the efficient public capital-output ratio is*

$$\frac{K_G^{SP}}{Y^{SP}} = \frac{\zeta}{\beta^{-1} - 1 + \delta_G}.$$

The competitive equilibrium delivers entirely different values of the public capital-output ratio:


$$\frac{K_G^{CE}}{Y^{CE}} = \frac{(1 - \mu)}{\delta_G} \cdot \frac{(1 - \xi)\tau(L_M + L_C) - L_C}{L_M}. \quad (27)$$

Proof. See Appendix [A.2](#). □

The efficient public capital-output ratio depends critically, and quite intuitively, on the parameter ζ . The planner will want a higher stock of public capital relative to output the higher is ζ . The rate of time preference and depreciation rate also factor, and for our calibrated parameters the efficient public capital-output ratio is around two. Interestingly, this applies to all economies: neither A_M nor A_C features in optimal public capital stocks relative to output.

The efficient public capital-output ratio also depends on the labor share in the production function, effective tax rates (which depend on A_C), the fraction of tax revenues invested rather than transferred, the total number of tax collectors, and the total labor supply. The output elasticity of public capital is not present at all. These equations highlight how this neoclassical model is not wired to reach an efficient steady state.

Figure 6: Public Capital: Efficient Benchmark Relative to Market Allocation



Notes: This figure plots the efficient public capital to output ratio relative to the market’s ratio.

We solved the social planner’s steady state for each country given the same parameter values in the competitive equilibrium. Figure 6 plots the ratio of the planner’s capital-output ratio to the one present in the competitive equilibrium. A value of one, in other words, means that the public capital-output ratios are the same under the planner and market outcomes.

We find that in low-income countries, the public capital-output ratios are significantly higher under the planner’s outcome than in the market allocation. On average, the planner has a much higher ratio than that of the market, suggesting underprovision of public capital. Middle- and high-income countries are closer to the efficient level, though there is significant variation country to country. On average, countries above around \$8,000 per capita income per year have a ratio around one. This analysis points in the direction of low public capital stocks in less developed countries being an inefficient outcome relative to this idealized scenario when output could be costlessly redirected toward investments in “public” and “private” types of capital. To the extent that low-income countries could reduce the costs of collecting taxes, in other words, the resulting increases in capital stocks would move them closer to the efficient benchmark.

5.4 Sensitivity Analysis

The output elasticity of public capital ζ , and the returns to scale in tax collection η , are two parameters of our analysis that are less standard and matter for our quantitative results. Panel A of Table 9 recalibrates the model country by country for $\zeta \in \{0.10, 0.20, 0.30\}$, holding $\eta = 0.05$, and for $\eta \in \{0.10, 0.20, 1.00\}$, holding $\zeta = 0.20$. Panel B explores alternative parameterizations that match either T/Y exactly or T/X exactly. Throughout the table, the government chooses collection employment optimally before and after each shock, as in Section 6.1.

The estimated variation in tax productivity does not depend on ζ . The slope of $\log A_C$ on $\log$ GDP per capita is 0.8 in all three columns, because A_C is identified from collections per unit of expenditure and the employment share of the tax authority. Neither of these two sufficient statistics involves the production function. However, ζ changes how cross-country income differences are explained by public capital relative to aggregate productivity. A larger ζ attributes more of the income gap to the larger public capital stocks of rich countries, so the slope of $\log A_M$ falls from 0.53 at $\zeta = 0.10$ to 0.16 at $\zeta = 0.30$.

The magnitude of the gains from tax productivity, by contrast, increases with ζ . Raising A_C by fifty percent raises GDP in low-income countries by 7.36 percent at $\zeta = 0.10$, 18.77 percent at our benchmark, and 38.32 percent at $\zeta = 0.30$, with similar changes in consumption and after-tax wages. The magnitude of these gains is consistent with our explanation of the magnifying role of ζ for elasticity of output to tax productivity in Section 5.2.²⁰

The returns to scale in tax collection, η , matter in two places. The first is the measurement of tax productivity. Since $\tau = A_C L_C^\eta$, the model recovers tax productivity as $\log A_C = \log \tau - \eta \log L_C$, and richer countries employ more collectors. A larger η therefore attributes more of their higher tax rates to the size of their tax administrations and less to their productivity. The slope of $\log A_C$ on $\log$ GDP per capita falls from 0.80 at our benchmark to 0.37 under constant returns, $\eta = 1$. Even in that extreme case, tax productivity rises with income at least as steeply as aggregate productivity, whose slope is 0.34. Second, under constant returns to tax collection, the gains from tax productivity shocks are much smaller and closer to those predicted by the Hulten theorem. GDP rises by only 1.1 percent, consumption by 0.6 percent, and after-tax wages by 0.5 percent.

Our baseline calibration targets both collections per unit of expenditure, T/X , and the tax-to-GDP ratio, T/Y . Panel B of Table 9 recalibrates every country to match only one of them, holding the remaining parameters at their baseline values. In both cases, raising A_C by fifty percent raises GDP substantially, and most in low-income countries. Targeting T/Y alone matches the higher

²⁰For instance, we find that in the dynamic model the first term of the elasticity of output to tax productivity is $\zeta/(1 - \mu - \zeta)$ and increasing in ζ .

Table 9: Parameter Sensitivity with Optimal Collection

Panel A. Parameter sensitivity						
Statistic	Baseline	$\zeta = 0.10$	$\zeta = 0.30$	$\eta = 0.10$	$\eta = 0.20$	$\eta = 1.00$
OLS slope of $\log A_M$	0.34	0.53	0.16	0.34	0.34	0.34
OLS slope of $\log A_C$	0.80	0.80	0.80	0.77	0.73	0.37
<i>Tax productivity: $A_C \times 1.5$</i>						
Δ GDP (%)	18.77	7.36	38.32	19.84	22.33	1.06
Δ consumption (%)	18.31	6.94	37.78	19.32	21.58	0.64
Δ after-tax wage (%)	18.77	6.99	39.02	19.79	22.00	0.52
<i>Modern productivity: $A_M \times 1.5$</i>						
Δ GDP (%)	127.73	98.24	180.69	127.58	127.24	127.66
Δ consumption (%)	127.73	98.24	180.70	127.58	127.25	127.74
Δ after-tax wage (%)	134.06	102.82	190.52	133.91	133.56	134.10

Panel B. Fiscal moment target sensitivity: $A_C \times 1.5$				
Income group	Statistic	Baseline (both)	T/X only	T/Y only
Low income	Δ GDP (%)	18.77	18.87	13.78
	Δ consumption (%)	18.31	18.71	6.21
	Δ public capital (%)	81.97	82.10	66.30
	Δ after-tax wage (%)	18.77	19.32	2.03
	Δ tax rate (pp)	0.66	0.23	9.39
Middle income	Δ GDP (%)	14.68	17.12	11.40
	Δ consumption (%)	9.70	15.02	5.41
	Δ public capital (%)	71.89	78.00	59.00
	Δ after-tax wage (%)	5.53	13.64	-0.14
	Δ tax rate (pp)	8.03	3.64	9.45
High income	Δ GDP (%)	8.55	12.57	7.63
	Δ consumption (%)	3.87	8.40	3.22
	Δ public capital (%)	47.52	63.49	42.74
	Δ after-tax wage (%)	-2.03	3.82	-2.39
	Δ tax rate (pp)	8.66	7.71	8.10

Notes: Collection employment L_C is chosen to maximize steady-state utility before and after each shock. Panel A: each column recalibrates the model, changing only the indicated parameter from the baseline ($\zeta = 0.20$, $\eta = 0.05$). Effects are unweighted means of percent changes across the 10 low-income countries. The effects of tax productivity gains on GDP increase with ζ but decline, close to Hulten, when $\eta = 1$. Panel B: the model is recalibrated to both fiscal moments (baseline) or to only T/X or the tax-to-GDP ratio. Effects are unweighted means across 10 low-, 46 middle-, and 42 high-income countries. In both cases, raising A_C by fifty percent raises GDP substantially, and most in low-income countries.

tax-to-GDP ratios of low-income countries in the data but does not qualitatively change our result.

6 Extensions

Our baseline makes four simplifying assumptions about the tax administration and the economy: (1) the size of the tax authority is exogenously calibrated rather than optimally chosen by the government, (2) all output is taxable and benefits from public capital, (3) collection uses only labor, and (4) collectors earn the private-sector wage. This section relaxes each of these assumptions in turn. Section 6.1 lets the government choose the number of collectors, which is where the returns to collection η become important, and also studies Ramsey taxation with government borrowing. Section 6.2 adds an untaxed traditional sector. Section 6.3 adds capital to the collection technology. And Section 6.4 allows collectors to earn a public-sector wage premium.

6.1 Endogenizing Tax Collection

The baseline calibration matches the size of the tax-collection sector, L_C , to its observed employment share. This treats the government's choice of collection capacity as given by data rather than as a policy decision. Instead, we let the government choose L_C while holding all other country-specific parameters at their calibrated values. The employment of collection affects the tax rate through $\tau(L_C) = A_C L_C^\eta$. Hiring more collection personnel can therefore increase the tax rate and hence tax revenue. The tax revenue can then be available for transfers and investment in public capital. As more labor is put into tax collection, the disutility of labor increases. The government chooses collection employment by balancing these benefits and costs.

For country i , holding all other variables and parameters fixed, we choose $L_{C,i}^*$ to maximize one-period household utility in the steady state subject to steady-state equilibrium conditions. Note that we do not recalibrate the remaining parameters or rematch the empirical moments. However, due to changes in L_C , a new equilibrium steady state is generated for each country and results in different consumption, employment, capital stock, output, and fiscal variables. The exercise thus identifies a conditional steady-state optimum rather than the outcome of a new calibration. It also abstracts from the transition between the calibrated and optimal steady states.

Table 10 reports the country-level results by income group. The ratio L_C/L_C^* provides the most direct measure of the tax collection personnel gap. If this ratio is below one, the calibrated collection employment falls short of the welfare-maximizing level. In contrast, a ratio above one means that calibrated employment exceeds the welfare-maximizing level. The mean ratio is less than one for

low- and middle-income countries and above one for high-income countries. These ratios suggest that low- and middle-income countries have insufficient employment in the tax collection sector.

The associated equilibrium responses clarify the underlying trade-off. In low- and middle-income countries, moving toward the optimum increases the tax rate and public capital. In these economies, the fiscal gains from additional collection capacity outweigh the labor and resource costs of employing more collectors and hence boost GDP and consumption. In high-income countries, consumption and GDP also rise at the optimum, but public capital is approximately unchanged. For this group, welfare gains come from reducing collection employment and its resource cost, even though the government accumulates little public capital. Maximizing household welfare therefore need not maximize either tax-collection capacity or the public-capital stock.

Table 10: Optimal Collection Labor by Income Group

Income group	L_C/L_C^*	L_C/L (%)	Δ Utility	Δ Consumption (%)	Δ GDP (%)	ΔK_G (%)
Low income	0.78	0.03	0.00	0.12	0.12	0.48
Middle income	0.53	0.16	0.01	0.38	0.64	3.31
High income	2.37	0.25	0.01	0.76	0.38	0.02

Notes: Each entry is the unweighted mean of country-specific calculations. Utility changes are levels of one-period utility; consumption, GDP and public-capital changes are country-specific percentage changes.


Figure 7 visualizes the corresponding country-level evidence. It compares the calibrated employment share, L_C/L , with the optimal share L_C^*/L^* , at each country's GDP per capita. The figure shows the divergence between calibrated and optimal collection employment share. On average, countries with low GDP per capita have a lower calibrated employment share than the optimum.

Ramsey Taxation. The exercise above finds the tax authority to maximize steady-state utility and ignores the transition. We now let the government choose the whole sequence of collection inputs, $\{L_{C,t}\}$, to maximize households' lifetime welfare, taking into account how households respond. The government's accounts are like the baseline model. Collected taxes, $T_t = \tau_t w_t L$, pay for transfers, $TR_t = \xi T_t$, for tax collector wages, $X_t = w_t L_{C,t}$, and for public investment, so that public capital evolves exactly as in Section 3,

$$K_{G,t+1} = (1 - \delta_G)K_{G,t} + w_t [(1 - \xi)\tau_t L - L_{C,t}] \equiv (1 - \delta_G)K_{G,t} + \Phi_t w_t, \quad (28)$$

where Φ_t is revenue net of transfers and collection costs, in units of the wage, and the transfer share ξ is held at its calibrated value. We make two changes to the baseline. First, households supply a fixed amount of labor, $L = L_{M,t} + L_{C,t}$. Section 5.2 shows that labor-supply responses matter little for the output effects of tax productivity, and fixing labor isolates collection costs as the only

Figure 7: Optimal Tax Collection Share



Notes: Tax-collection employment as a percent of total employment, by country, against GDP per capita (PPP, 2021 USD, log scale). Red: calibrated share L_C/L . Blue: welfare-maximizing share L_C^*/L^* . Vertical lines connect the two values for each country, and curves are smoothed fits.

source of a marginal cost of public funds.²¹ Second, the government has a general account that can borrow at the gross rate R_t and levy a wedge ω_t on the gross return to private capital, including a levy on initial wealth. This account only shapes the return to private capital: public investment, transfers, and collection must still be paid for out of collected taxes, as in the baseline. Following the primal approach (Lucas and Stokey, 1983; Chari and Kehoe, 1999), the government chooses allocations directly. Because the levy on initial wealth keeps the implementability condition from binding, its problem reduces to maximizing $\sum_t \beta^t u(C_t)$ subject to the resource constraint and (28) (Appendix A.6).

Proposition 6. *At an interior steady state of the Ramsey problem, collection and public capital satisfy*

$$\frac{1 - \beta(1 - \delta_G)}{\mathcal{M}} = \underbrace{\frac{F_{K_G} + 1 - \delta_G}{R} - 1}_{\text{return on public relative to private capital}} + \underbrace{\beta \frac{\Phi F_{L_M K_G}}{\mathcal{M}}}_{\text{tax base gain net of personnel pay}}, \quad (29)$$

²¹With elastic labor supply, the household's labor-supply condition, $\alpha L_t^{1/\varepsilon} = C_t^{-\gamma}(1 - \tau_t)w_t$, becomes an additional constraint on the government, because τ_t is tied to $L_{C,t}$ rather than chosen freely. Its multiplier adds the marginal excess burden of the labor tax to the collection markup in (32).

and private capital satisfies

$$1 = \beta(F_K + 1 - \delta) + \beta \frac{\Phi F_{LM}K}{\mathcal{M}}, \quad (30)$$

where

$$\mathcal{M} \equiv (1 - \xi) \tau_{LC} L - 1 - \Phi \frac{F_{LM}L_M}{F_{LM}}$$

is the revenue that one more collector adds to public investment, net of the collector's own wage and in units of the wage, $\tau_{LC} \equiv \partial \tau / \partial L_C$.

The right-hand side of (29) collects the returns to one more unit of public capital. The first term is its net return relative to private capital: a unit of public capital yields $F_{K_G} + 1 - \delta_G$ next period, while the same resources lent to the government or by arbitrage invested in private capital yield R . The second term arises because public capital raises the modern-sector wage, which enlarges the tax base and with it the revenue Φw available for public investment; because collectors are paid the market wage, Φ already nets out their higher pay. The left-hand side is the cost of financing that unit through collection. Its numerator, $1 - \beta(1 - \delta_G)$, is the discounted user cost of public capital. Its denominator, $\mathcal{M}$, is what the marginal collector contributes to public investment. With our Cobb–Douglas technology,

$$\mathcal{M} = (1 - \xi) \eta \frac{T}{X} - 1 + \frac{\mu}{1 - \mu} s_G, \quad s_G \equiv \frac{\delta_G K_G}{Y}, \quad (31)$$

where T/X is the collection ratio in (10), the statistic we measure in Section 2.1. A collector raises $\eta T/X$ in revenue per unit of pay. The non-transferred share $1 - \xi$ goes to public investment, which must also pay the collector's wage, hence the -1 . The last term, which is small, is the extra revenue from the higher wage of the workers who remain in production. Using $\beta R = 1$, (29) can be rewritten as

$$\underbrace{F_{K_G} + \frac{\Phi F_{LM}K_G}{\mathcal{M}}}_{\text{social marginal product}} = \underbrace{(\beta^{-1} - 1 + \delta_G)}_{\text{user cost}} \underbrace{\left(1 + \frac{1}{\mathcal{M}}\right)}_{\text{collection markup}}. \quad (32)$$

The social marginal product of public capital, which adds its effect on the tax base to its direct effect on output, must equal its user cost marked up by $1/\mathcal{M}$: a marginal cost of public funds that arises purely from the cost of collection (Keen and Slemrod, 2017). Because public funds are valuable, an interior optimum has $\mathcal{M} > 0$: the government never expands collection past the point at which the marginal collector stops covering their own wage. Moreover, $(1 - \xi)\eta T/X$ is proportional to $L_C^{\eta-1}$, so with $\eta = 0.05$ doubling the tax authority nearly halves the gross marginal return, and hiring raises the markup much faster than it raises revenue. Higher tax productivity, by contrast, raises $\mathcal{M}$ at every size of the tax authority. When collection is costless, $\mathcal{M} \rightarrow \infty$, the markup and the tax base term vanish, and (32) reduces to the planner's condition of Proposition 5.

With costly collection, and under our Cobb–Douglas technology, the Ramsey public capital-output ratio is

$$\frac{K_G^R}{Y^R} = \frac{\zeta}{\beta^{-1} - 1 + \delta_G} \cdot \frac{\mathcal{M}}{\mathcal{M} + 1 - s_G^{SP}}, \quad s_G^{SP} \equiv \frac{\zeta \delta_G}{\beta^{-1} - 1 + \delta_G}, \quad (33)$$

the efficient ratio scaled by a factor that rises from zero to one with $\mathcal{M}$, where $s_G^{SP} \approx 0.11$ is the efficient public investment rate.

We quantify the Ramsey steady-state in the Appendix, Table A.3. In low-income countries, the Ramsey government hires more collectors, about half again as many as in the baseline, but public capital is essentially unchanged, at about 5 percent of the efficient ratio. The reason is that the marginal collector barely covers their own wage. Inverting (33) gives $\mathcal{M} = s_G(1 - s_G^{SP})/(s_G^{SP} - s_G)$, and with public investment at about 0.6 percent of output, $\mathcal{M} \approx 0.05$ at the Ramsey steady state, so the marginal cost of public funds exceeds 20. The micro evidence points the same way: the marginal return of 1.6 that Dzansi et al. (2025) measure in Ghana, combined with the average transfer share of low-income countries, implies $\mathcal{M} \approx 0.2$, against about 5.7 for the marginal return of 13 that Boning et al. (2025) measure in the United States. Middle-income countries also hire more collectors, about a third more, and hold public capital at about 60 percent of the efficient ratio, while high-income countries cut collection employment by about two-thirds and hold public capital at about 89 percent of the efficient ratio, against 93 percent in the baseline. These directions match the steady-state exercise above. Appendix A.6 sets out the full Ramsey problem and proves Proposition 6.

6.2 Extending the Model with a Traditional Sector

The baseline model has only one sector whose labor income is taxable and whose output rises with public capital. This abstraction may be unsuitable for low-income economies, where a portion of economic activity, including home production, small-scale agriculture, and informal business, is difficult to tax. Moreover, public capital need not raise output equally across sectors. We therefore introduce an untaxed traditional sector to assess whether gains from higher tax productivity are maintained, under the two considerations above.

We define a traditional sector that only uses labor, without private or public capital. We calibrate its employment share using two alternative measures: total agricultural employment and a constructed measure of unmechanized agricultural employment (World Bank, 2026a). The first provides a broad proxy for traditional activity and the second excludes the estimated mechanized portion of agriculture, which need not resemble the untaxed, labor-only sector in the model. These alternatives change the employment target, not the model’s economic structure.

Accounting for traditional production reduces the estimated gains from higher tax productivity in low-income countries. With tax-collection employment held fixed, a 50 percent increase in tax productivity raises GDP by 6.2 percent under the all-agriculture measure, compared with 18.42 percent in the baseline model. Consumption rises by 5.0 percent. The extension therefore preserves positive gains of the tax productivity shock in low-income countries.


The traditional sector produces $Y_T = A_T L_T$ using labor alone, without private or public capital. Traditional income is untaxed, so households earn the full marginal product of their labor, $w_T = A_T$. Aggregate output is the sum of modern and traditional production, $Y = Y_M + Y_T$. We retain the baseline technologies for modern production and tax collection, but allow households to supply labor to traditional production as well as modern employment. Household utility in one period becomes

$$U(C, L_M + L_C, L_T) = \frac{C^{1-\gamma}}{1-\gamma} - \alpha_M \frac{(L_M + L_C)^{1+1/\varepsilon}}{1+1/\varepsilon} - \alpha_T \frac{L_T^{1+1/\varepsilon}}{1+1/\varepsilon}.$$

The wage is $w_T = A_T$, and since the sector is untaxed, households earn the full marginal product. We fix $\alpha_M = 1$ and $\alpha_T = 1$, and calibrate A_T alongside the existing parameters (L_C, A_M, A_C, ξ) to match each country's agricultural employment share.

Figure 8 shows the calibrated productivities for all three sectors.

Figure 8: Model-Implied Tax Productivities across Sectors



Notes: Calibrated productivities in the model with an untaxed traditional sector, with traditional-sector productivity calibrated to each country's agricultural employment share. Panels A, B, and C plot tax productivity A_C , modern-sector productivity A_M , and traditional-sector productivity A_T , each relative to the world average, against GDP per capita.

Table 11 reports the counterfactuals by income groups.

Table 12 shows the prediction of a 50% A_C shock to lower-income countries under Hulten's Theorem. Under the calculations of Hulten's Theorem, the tax collection sector is small, measured by its wage input $w_M L_C / Y$ or its product T / Y . Consequently, the predicted gain from the change in tax sector productivity is minimal.

Table 11: Policy Counterfactuals: All-Agriculture Calibration

Shock	Income group	GDP	Consumption	K_G	After-tax wage	T/X	$\Delta\tau$ (pp)
$A_C \times 1.5$	Low income	1.06	1.05	2.02	1.22	1.61	1.32
$A_C \times 1.5$	Middle income	1.13	1.07	1.73	1.03	1.46	9.57
$A_C \times 1.5$	High income	1.08	1.00	1.63	0.89	1.40	15.14
$A_M \times 1.5$	Low income	1.50	1.42	3.92	2.74	1.40	-0.00
$A_M \times 1.5$	Middle income	2.22	2.19	2.38	2.37	1.00	-0.00
$A_M \times 1.5$	High income	2.28	2.28	2.28	2.34	0.97	-0.00
$L_C \times 1.5$	Low income	1.00	1.00	1.00	1.00	0.68	0.05
$L_C \times 1.5$	Middle income	1.00	1.00	1.01	1.00	0.68	0.39
$L_C \times 1.5$	High income	1.00	1.00	1.01	0.99	0.68	0.66
$\theta = 0.84$	Low income	1.06	1.05	1.30	1.21	1.07	-0.00
$\theta = 0.84$	Middle income	1.05	1.05	1.06	1.06	1.00	-0.00
$\theta = 0.84$	High income	1.00	1.00	1.00	1.00	1.00	-0.00

Notes: GDP, consumption, public capital, after-tax wages and T/X are mean country counterfactual/baseline factors. The final column is an actual percentage-point change: $100(\tau' - \tau)$, not a fraction-valued difference. The baseline calibration has 97 countries. Solved low/middle/high-income samples by shock: $A_C \times 1.5$: 10/46/39; $A_M \times 1.5$: 10/46/41; $L_C \times 1.5$: 10/46/41; $\theta = 0.84$: 10/46/41. Transfer shares remain fixed. The $A_C \times 1.5$ counterfactual is inadmissible for Denmark, Sweden (2 cases), because the implied tax rate violates $0 \leq \tau' < 1$. These cases are excluded from the shock means, not assigned zero effects.

Table 12: Hulten Accounting versus Model Predictions for a 50% A_C Shock

Prediction method	Δ GDP (%)
Model output	6.22
Hulten wage-bill term: $0.5w_M L_C/Y$	0.01
Hulten tax-revenue term: $0.5T/Y$	0.31
Public-capital contribution: $\zeta \overline{\% \Delta K_G} \times \overline{Y_M/Y}$	5.67

Notes: All four rows use the same 10 solved low-income countries. Hulten terms use baseline shares. Overbars denote unweighted country means. The public-capital row holds other production inputs fixed; it is not the full equilibrium response. The modern-output share Y_M/Y is evaluated at baseline; the last row multiplies the two country means.

Treating the agricultural sector as a nontaxable traditional sector may be untenable for many advanced economies, where agriculture is mostly mechanized and pays various taxes. An alternative measure of the traditional sector is constructed as the unmechanized agricultural share (see Appendix A.4).

6.3 Capital in Tax Collection

A model with a labor-only tax collection function may assign higher collection productivity to tax administrations whose collectors are better equipped. We include collection capital to assess whether the positive relationship between calibrated tax collection productivity and GDP per capita depends on excluding this input. The positive relationship persists in the extended model.

Tax collection now combines capital and labor:

$$\begin{aligned}\tau &= A_C K_C^{\eta\psi} L_C^{\eta(1-\psi)} \\ &= \underbrace{\left[A_C \left(\frac{K_C}{L_C} \right)^{\eta\psi} \right]}_{\text{productivity inferred omitting capital}} L_C^{\eta}.\end{aligned}$$

Here, K_C denotes collection capital, and ψ allocates the total input elasticity η between capital and labor. Holding the collection labor fixed, the term in brackets is the productivity inferred by a labor-only model. It combines capital per collector with efficiency conditional on inputs.

The expenditure data establish that tax administration uses substantial resources beyond salaries. Non-wage expenditure accounts for roughly one-quarter to one-third of operating expenditure across income groups (Table 13). Combined operating and capital ICT expenditure averages between 8.5% and 15.05% of operating expenditure, with the highest mean in the high-income group. These expenditure ratios motivate the capital extension, but they do not directly measure capital stocks or capital per collector.

Table 13: Tax-Administration Expenditure Shares by Income Group

Income group	Non-wage share (%)	ICT share (%)
Low income	26.83	11.49
Middle income	33.32	8.50
High income	27.43	15.05

Notes: Both denominators are operating expenditure. Non-wage shares are $(O - W)/O$; ICT includes operating and capital ICT and requires both observations. Equal-country means use the exact calibration reference year, without filling missing components. Non-wage low/middle/high-income sample counts are 5/33/37; ICT counts are 1/24/32. Observed low-income ICT countries: MWI. The main model uses the non-wage share, not the ICT sensitivity.

Administration expenditure includes capital replacement:

$$X = wL_C + \delta K_C.$$


The government budget finances X , and the resource constraint becomes

$$C + \delta(K + K_C) + \delta_G K_G = Y.$$

To calibrate collection capital, we impose $\delta_C K_C / X$ equals the income-group mean of the non-wage share at the calibrated equilibrium. We recalibrate the model using the same moment targets, objective weights, and remaining fixed parameters as in the baseline. Collection capital and collection labor remain fixed when solving the competitive equilibrium.

Figure 9 shows the relationship between tax collection productivity A_C and GDP per capita. Panels A, B, and C compare this relationship under different specifications of the tax collection function. Introducing collection capital using the non-wage expenditure share lowers the log-log slope from 0.80 in the labor-only model to 0.65. The ICT-share sensitivity gives a slope of 0.74. The relationship remains positive in all three specifications. Therefore, the cross-country pattern of collection productivity cannot be explained by variation in the use of capital.

Figure 9: Tax Collection Productivity with Capital and Labor Inputs



Notes: Calibrated tax productivity A_C relative to the world mean. Panel A uses the labor-only collection technology. Panels B and C use $\tau = A_C K_C^\psi L_C^{\eta(1-\psi)}$, with collection capital calibrated so that $\delta_C K_C / X$ equals the income-group mean of the non-wage share (Panel B) or of the ICT share (Panel C) of tax-administration expenditure.

6.4 Accounting for Public-Sector Wage Differences

Our baseline pays tax collectors the private-sector wage. Under that assumption, a country with low collections per unit of expenditure, T/X , must have low tax productivity, A_C . Public-sector wage premia tend to be larger in lower-income countries, so the equal-wage assumption may understate

their tax productivity and overstate how steeply A_C rises with income. This section asks how much of the cross-country gradient in A_C survives once we account for differences in public-sector pay.

Let $\pi_i = w_{C,i}/w_{M,i}$ denote the ratio of the wage paid to tax collectors to the private-sector wage in country i . The baseline imposes $\pi_i = 1$ in each country. A value of π_i above one captures a public-sector wage premium. Collection expenditure becomes $X_i = \pi_i w_{M,i} L_{C,i}$ and, because collectors' earnings are taxed, collections become $T_i = \tau_i w_{M,i} (L_{M,i} + \pi_i L_{C,i})$. The ratio of tax collection to expenditures is then

$$\frac{T_i}{X_i} = \frac{\tau_i (L_{M,i} + \pi_i L_{C,i})}{\pi_i L_{C,i}}.$$

For given values of T/X and of collection employment, the implied tax rate, and with it tax productivity $A_{C,i} = \tau_i L_{C,i}^{-\eta}$, rises approximately in proportion to π_i . The wage ratio also affects public investment, because collectors are paid out of revenue. The steady-state public-capital budget becomes

$$\delta_G K_{G,i} = (1 - \xi_i) \tau_i w_{M,i} (L_{M,i} + \pi_i L_{C,i}) - \pi_i w_{M,i} L_{C,i},$$

so a higher π_i leaves less revenue for public capital.

We measure π_i with the estimates of [Abdallah et al. \(2023\)](#), who compare the wages of public- and private-sector workers with similar observable characteristics. Their estimates cover general-government employees rather than tax administrations specifically, so we use them as a proxy for the relative cost of collection labor. Direct estimates are available for 67 of the 98 countries in our sample. For the remaining 31, we impute π_i with a Gaussian Nadaraya–Watson regression of the observed wage ratios, in levels, on log GDP per capita. We then recalibrate every country with π_i fixed at its observed or imputed value, targeting the same moments and holding all other parameters at their baseline values.

Figure 10 compares the inferred tax productivities from accounting for public sector wage premia with the benchmark calibration. The new calibration raises A_C most in lower-income countries, where public-sector premia are largest, and changes it little at the top of the income distribution. The relationship between tax productivity and income flattens but remains steep. The slope of $\log A_C$ on $\log$ GDP per capita falls from 0.80 to 0.72, still about twice the slope of aggregate productivity in the baseline, 0.34. Because the adjustment shifts $\log A_C$ by roughly $\log \pi_i$, this flattening approximately measures how fast public-sector premia decline with income. Differences in public-sector pay therefore account for about a tenth of the cross-country gradient in tax productivity.

Figure 10: Tax Collection Productivity With and Without Public Sector Wage Adjustment.



Notes: Calibrated tax productivity A_C (log scale) against GDP per capita (PPP, 2021 USD, log scale) in the baseline model. Hollow circles: baseline calibration, in which collectors earn the private-sector wage. Filled circles: recalibration in which collectors earn π_i times the private-sector wage. Lines are OLS fits; the log–log slope falls from 0.80 without the adjustment to 0.72 with it.

7 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 a new macro-historical dataset covering countries across a wide range of income levels and spanning multiple decades, we document a robust positive association between tax collections per expenditure on collection and GDP per capita across countries, with a qualitatively similar but less precisely estimated relationship within countries over the long run.

We interpret these empirical patterns through a quantitative growth model that incorporates tax administration into neoclassical and Ramsey frameworks. 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. Quantitative findings indicate that improvements in tax productivity could substantially impact aggregate outcomes, and significantly more than suggested by the tax administration sector’s rather modest share of total employment. Low-income countries today, according to our model’s calculations, likely have inefficiently low levels of public capital to output due in large part to their low levels of tax productivity.

References

- Abdallah, Chadi, David Coady, and La-Bhus Fah Jirasavetakul**, “Public-Private Wage Differentials and Interactions Across Countries and Time,” IMF Working Paper WP/23/64, International Monetary Fund, Washington, DC March 2023.
- Aschauer, David Alan**, “Is Public Expenditure Productive?,” *Journal of Monetary Economics*, 1989, 23 (2), 177–200.
- Asian Development Bank, Inter-American Center of Tax Administrations, International Monetary Fund, Intra-European Organisation of Tax Administrations, and Organisation for Economic Co-operation and Development**, “International Survey on Revenue Administration (ISORA),” IMF Data, dataset 2026.
- Bachas, Pierre, Lucie Gadenne, and Anders Jensen**, “Informality, Consumption Taxes, and Redistribution,” *Review of Economic Studies*, 2024, 91 (5).
- , **Matthew Fisher-Post, Anders Jensen, and Gabriel Zucman**, “Globalization, Capital Taxation, and Development: Evidence from a Macro-Historical Database,” 2024. Unpublished Working Paper, Harvard University.
- Balan, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan L. Weigel**, “Local Elites as State Capacity,” *American Economic Review*, 2022, 112 (3).
- Baqae, David Rezza and Emmanuel Farhi**, “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 2019, 87 (4), 1155–1203.
- **and** —, “Productivity and Misallocation in General Equilibrium,” *Quarterly Journal of Economics*, 2020, 135 (1), 105–163.
- **and Hannes Malmberg**, “Long-Run Comparative Statics,” NBER Working Paper 33504, National Bureau of Economic Research 2025.
- Barro, Robert J.**, “Government Spending in a Simple Model of Endogeneous Growth,” *Journal of Political Economy*, 1990, 98 (5).
- Basri, Chatib, Mayara Felix, Rema Hanna, and Ben Olken**, “Tax Administration versus Tax Rates: Evidence from Corporate Taxation in Indonesia,” *American Economic Review*, 2021, 111 (12), 3827–71.
- Baxter, Marianne and Robert G. King**, “Fiscal Policy in General Equilibrium,” *American Economic Review*, 1993, 83 (3), 315–334.
- Besley, Tim, Ethan Ilzetzi, and Torsten Persson**, “Weak States and Steady States,” *American Economic Journal: Macroeconomics*, 2013, 5 (4), 203–235.
- Besley, Timothy and Torsten Persson**, *Pillars of Prosperity*, Princeton University Press, 2011.

- Bhandari, Anmol, David Evans, Yuki Yao, and Ellen R. McGrattan**, “Business Income Under-reporting and Public Finance,” 2024. Federal Reserve Bank of Minneapolis Working Paper.
- Bick, Alexander, Nicola Fuchs-Schündeln, and David Lagakos**, “How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications,” *American Economic Review*, 2018, 108 (1), 170–199.
- , —, —, and **Hitoshi Tsujiyama**, “Structural Change in Labor Supply and Cross-Country Differences in Hours Worked,” *Journal of Monetary Economics*, 2022, 130, 68–85.
- Bom, Pedro R.D. and Jenny E. Ligthart**, “What Have We Learned From Three Decades of Research on the Productivity of Public Capital?,” *Journal of Economic Surveys*, December 2014, 28 (5), 889–916.
- Boning, William C, Nathaniel Hendren, Ben Sprung-Keyser, and Ellen Stuart**, “A Welfare Analysis of Tax Audits Across the Income Distribution,” *Quarterly Journal of Economics*, 2025, 140 (1).
- Boppart, Timo and Per Krusell**, “Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective,” *Journal of Political Economy*, 2020, 128 (1), 118–157.
- , **Patrick Kiernan, Per Krusell, and Hannes Malmberg**, “The Macroeconomics of Intensive Agriculture,” 2023. NBER Working Paper 31101.
- Brockmeyer, Anne, Alejandro Estefan, Karina Arras, and Juan Carlos Suárez Serrato**, “Taxing Property: Theory and Evidence from Mexico,” *Working Paper*, 2023.
- , **Spencer Smith, Marco Hernandez, and Stewart Kettle**, “Casting a Wider Tax Net: Experimental Evidence from Costa Rica,” *American Economic Journal: Economic Policy*, 2019, 11 (3).
- Caselli, Francesco**, “Accounting for Cross-Country Income Differences,” in P. Aghion and S. Durlauf, eds., *Handbook of Economic Growth*, 2005.
- Chamley, Christophe**, “Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives,” *Econometrica*, 1986, (3), 607–622.
- Chari, V. V. and Patrick J. Kehoe**, “Optimal Fiscal and Monetary Policy,” in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics, Volume 1C*, Amsterdam: Elsevier, 1999, pp. 1671–1745.
- Christians, Allison**, “Drawing the boundaries of tax justice,” *Working Paper*, 2013.
- Colmer, Jonathan M., David Lagakos, and Martin Shu**, “Is the Electricity Sector a Weak Link in Development?,” 2024. NBER Working Paper 32041.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger**, “Taxing Capital? Not a Bad Idea after All!,” *American Economic Review*, 2009, 99 (1), 25–48.

- Cubas, German**, “Public Capital and Economic Development,” *Economic Journal*, 2020, 130 (632), 2354–2381.
- D’Arcy, Michelle, Marina Nistotskaya, and Ola Olsson**, “Cadasters and Economic Growth: A Long-Run Cross-Country Panel,” *Journal of Political Economy*, 2024, 132 (11).
- Devadas, Sharmila and Steven Pennings**, “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 October 2018.
- Dincecco, Mark and Gabriel Katz**, “State Capacity and Long-Run Economic Performance,” *Economic Journal*, 2016, 126 (590).
- Donaldson, Dave and Richard Hornbeck**, “Railroads and American Economic Growth: A “Market Access” Approach,” *Quarterly Journal of Economics*, 2016, 131 (2), 799–858.
- Dzansi, James, Anders Jensen, David Lagakos, and Henry Telli**, “Technology and Tax Capacity: Evidence from Local Governments in Ghana,” NBER Working Paper 29923, National Bureau of Economic Research 2025.
- Ebrill, Liam P., Michael Keen, Jean-Paul Bodin, and Victoria Summers**, *The Modern VAT*, Washington, DC: International Monetary Fund, 2001.
- EU**, “Tax Revenue Mobilization in Developing Countries,” *Technical Report*, 2015.
- Faber, Benjamin**, “Trade Integration, Market Size, and Industrialization: Evidence from China’s National Trunk Highway System,” *Review of Economic Studies*, 2014, 81 (3), 1046–1070.
- Farhi, Emmanuel and Iván Werning**, “Fiscal Multipliers: Liquidity Traps and Currency Unions,” in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics, Volume 2*, Amsterdam: Elsevier, 2016, pp. 2417–2492.
- Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer**, “The Next Generation of the Penn World Table,” *American Economic Review*, 2015, 105 (10), 3150–3182.
- Fraser Institute**, “Economic Freedom of the World: Dataset,” Online database 2026.
- Gennaioli, Nicola and Hans-Joachim Voth**, “State Capacity and Military Conflict,” *Review of Economic Studies*, 2015, 82 (4).
- Glomm, Gerhard and B. Ravikumar**, “Public Versus Private Investment in Human Capital: Endogenous Growth and Income Inequality,” *Journal of Political Economy*, 1992, 100 (4), 818–834.
- Gollin, Douglas**, “Getting Income Shares Right,” *Journal of Political Economy*, April 2002, 110 (2), 458–474.
- Gordon, Roger and Wei Li**, “Tax Structures in Developing Countries: Many Puzzles and a Possible Explanation,” *Journal of Public Economics*, August 2009, 93 (7-8).

- Groningen Growth and Development Centre**, “Penn World Table, Version 11.0,” Dataset 2025.
- Herrendorf, Berthold and Ákos Valentinyi**, “Which Sectors Make Poor Countries So Unproductive?,” *Journal of the European Economic Association*, 2012, 10 (2), 323–341.
- Hjort, Jonas and Lin Tian**, “The Economic Impact of Internet Connectivity in Developing Countries,” *Annual Review of Economics*, 2025, 17, 99–124.
- Hsieh, Chang-Tai and Peter J. Klenow**, “Relative Prices and Relative Prosperity,” *American Economic Review*, 2007, 97 (3), 562–585.
- Hulten, Charles R.**, “Growth Accounting with Intermediate Inputs,” *Review of Economic Studies*, 1978, 45 (3), 511–518.
- IMF**, “Current Challenges in Revenue Mobilization: Improving Tax Compliance,” *Staff Report*, 2015.
- International Monetary Fund**, “Investment and Capital Stock Dataset (ICSD): May 2021 Update,” Fiscal Affairs Department, dataset 2021.
- , “Government Finance Statistics (GFS),” IMF Data, database 2026.
- , “World Economic Outlook Database,” Online database 2026.
- Jensen, Anders**, “Employment Structure and the Rise of the Modern Tax System,” *American Economic Review*, 2022, 112 (1), 213–234.
- **and Jonathan Weigel**, “No Taxation without Administration: Bringing the State Back into the Public Finance of Developing Countries,” *Journal of Economic Literature*, 2026, 64 (1).
- Jones, Larry E., Rodolfo E. Manuelli, and Peter E. Rossi**, “Optimal Taxation in Models of Endogenous Growth,” *Journal of Political Economy*, 1993, 101 (3), 485–517.
- Judd, Kenneth L.**, “Redistributive Taxation in a Simple Perfect Foresight Model,” *Journal of Public Economics*, 1985, 28 (1), 59–83.
- Keen, Michael and Joel Slemrod**, “Optimal Tax Administration,” *Journal of Public Economics*, 2017, 152, 133–142.
- Khan, Adnan, Asim Khwaja, and Benjamin Olken**, “Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors,” *Quarterly Journal of Economics*, 10 2015, 131 (1).
- Kleven, Henrik, Claus Kreiner, Kristian Larsen, and Jakob Søgaaard**, “Micro versus Macro Labor Supply Elasticities: The Role of Dynamic Returns to Effort,” *American Economic Review*, 2025, 115 (9), 2849–90.
- Kleven, Henrik Jacobsen, Claus Thustrup Kreiner, and Emmanuel Saez**, “Why can Modern Governments Tax so Much?,” *Economica*, 2016, 83 (330).

- , **Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez**, “Unwilling or Unable to Cheat?,” *Econometrica*, 2011, 79 (3).
- Leeper, Eric M., Todd B. Walker, and Shu-Chun S. Yang**, “Government Investment and Fiscal Stimulus,” *Journal of Monetary Economics*, 2010, 57 (8), 1000–1012.
- Lucas, Robert E. and Nancy L. Stokey**, “Optimal Fiscal and Monetary Policy in an Economy without Capital,” *Journal of Monetary Economics*, 1983, 12 (1), 55–93.
- Mayshar, Joram, Omer Moav, and Luigi Pascali**, “The Origin of the State: Land Productivity or Appropriability?,” *Journal of Political Economy*, 2022, 130 (4).
- Naritomi, Joana**, “Consumers as Tax Auditors,” *American Economic Review*, 2019, 109 (9), 3031–72.
- OECD**, “Tax Administration 2024: Comparative Information on OECD and Other Advanced and Emerging Economies,” *Annual Report*, 2024.
- Okunogbe, Oyebola and Gabriel Tourek**, “How Can Lower-Income Countries Collect More Taxes?,” *Journal of Economic Perspectives*, 2024, 38 (1).
- Organisation for Economic Co-operation and Development**, “Inventory of Tax Technology Initiatives (ITTI): 2024 Update,” Dataset, developed with partner organisations 2024.
- Pomeranz, Dina**, “No Taxation without Information,” *American Economic Review*, 2015, 105 (8).
- Pritchett, Lant**, “The Tyranny of Concepts: CUDIE (Cumulated, Depreciated, Investment Effort) Is Not Capital,” *Journal of Economic Growth*, December 2000, 5 (4), 361–384.
- Ramey, Valerie A.**, “Can Government Purchases Stimulate the Economy?,” *Journal of Economic Literature*, 2011, 49 (3), 673–685.
- , “The Macroeconomic Consequences of Infrastructure Investment,” in Edward L. Glaeser and James M. Poterba, eds., *Economic Analysis and Infrastructure Investment*, Cambridge, MA: National Bureau of Economic Research, 2021, pp. 219–268.
- **and Neville Francis**, “A Century of Work and Leisure,” *American Economic Journal: Macroeconomics*, 2009, 1 (2), 189–224.
- Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu**, “Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis,” *Journal of Monetary Economics*, 2008, 55, 234–250.
- Schwab, Klaus, ed.**, *The Global Competitiveness Report 2017–2018*, Geneva: World Economic Forum, 2017.
- Sheahan, Megan and Christopher B. Barrett**, “Understanding the Agricultural Input Landscape in Sub-Saharan Africa: Recent Plot, Household, and Community-Level Evidence,” Policy Research Working Paper 7014, World Bank, Washington, DC August 2014.

- Straub, Ludwig and Iván Werning**, “Positive Long-Run Capital Taxation: Chamley-Judd Revisited,” *American Economic Review*, 2020, *110* (1), 86–119.
- Sugin, Linda**, “Payroll Taxes, Mythology, and Fairness,” *Harvard Journal on Legislation*, 2014, *51*.
- Tilly, Charles**, *Coercion, Capital, and European States, AD 990–1990*, Blackwell, 1990.
- UNESCO Institute for Statistics**, “Educational Attainment, at Least Completed Primary, Population 25+ Years, Total (%) (Cumulative),” Distributed through World Bank World Development Indicators 2026.
- World Bank**, *Information and Communications for Development 2018: Data-Driven Development*, Washington, DC: World Bank, 2019.
- , “Employment in Agriculture (% of Total Employment) (Modeled ILO Estimate),” World Development Indicators, indicator SL.AGR.EMPL.ZS 2026.
- , “World Development Indicators,” Database 2026.
- World Values Survey Association**, “World Values Survey: Data and Documentation,” Online data archive 2026.

A Online Appendix

A.1 Existence and Uniqueness of the Steady State Equilibrium

We now show that the steady state for the economy can be reduced to a single equation in which the only variable is labor L_M . Write the problem as $\mathcal{R}(L_M) = 0$. We use the endpoint property to establish existence and monotonicity to ensure uniqueness. Additionally, we provide an example of multiple steady-state equilibria.

The steady state for the economy can be characterized by $(C, K, K_G, L_M, Y, w, r)$ such that

$$\begin{aligned} 1 &= \beta(1 + r - \delta), \\ \alpha(L_M + L_C)^{1/\varepsilon} &= C^{-\gamma}(1 - \tau)w, \\ \delta_G K_G &= w \left[(1 - \xi)\tau(L_M + L_C) - L_C \right], \\ C + \delta K + \delta_G K_G &= Y, \\ Y &= A(\theta K_G)^\zeta K^\mu L_M^{1-\mu}, \\ r &= \mu A(\theta K_G)^\zeta K^{\mu-1} L_M^{1-\mu}, \\ w &= (1 - \mu)A(\theta K_G)^\zeta K^\mu L_M^{-\mu}. \end{aligned}$$

where $\tau = A_C(L_C)^\eta$ is known since L_C is fixed.

Proof of Proposition 1. Positive public capital requires

$$L_M > \frac{1 - (1 - \xi)\tau}{(1 - \xi)\tau} L_C = \underline{L}_M.$$

The rental rate is $r = \beta^{-1} - 1 + \delta$.

We can write the wage w as a function of L_M . Factor pricing gives $\frac{K}{L_M} = \frac{\mu w}{(1-\mu)r}$. Substitute this and the government-budget expression for K_G into the wage equation. The result is:

$$w(L_M) = \left[(1 - \mu)^{1-\mu} \left(\frac{\mu}{r} \right)^\mu A \theta^\zeta \left(\frac{(1 - \xi)\tau(L_M + L_C) - L_C}{\delta_G} \right)^\zeta \right]^{\frac{1}{1-\mu-\zeta}}$$

The wage is strictly increasing in modern labor:

$$\frac{d \log w(L_M)}{d L_M} = \frac{\zeta}{1 - \mu - \zeta} \times \frac{(1 - \xi)\tau}{(1 - \xi)\tau(L_M + L_C) - L_C} > 0.$$

Next, the resource constraint implies that

$$\begin{aligned}
C(L_M) &= Y - \delta K - \delta_G K_G \\
&= \frac{wL_M}{1-\mu} - \frac{\delta\mu wL_M}{(1-\mu)r} - w[(1-\xi)\tau(L_M + L_C) - L_C] \\
&= w(L_M) \left[\frac{1-\delta\mu/r}{1-\mu} - (1-\xi)\tau \right] L_M + w(L_M) [1 - (1-\xi)\tau] L_C
\end{aligned}$$

This is consumption available after replacing depreciated private capital and financing public-capital investment.

Moreover, because $r = \beta^{-1} - 1 + \delta > \delta$, $\frac{1-\delta\mu/r}{1-\mu} > 1 > (1-\xi)\tau$. Consequently, $C(L_M)/w(L_M)$ is positive and strictly increasing:

$$\frac{d}{dL_M} \left(\frac{C(L_M)}{w(L_M)} \right) = \frac{1-\delta\mu/r}{1-\mu} - (1-\xi)\tau > 0.$$

Finally, the labor supply condition gives the desired scalar equation

$$\mathcal{R}(L_M) = \underbrace{[\log \alpha + \varepsilon^{-1} \log(L_M + L_C)]}_{\text{log marginal disutility of working}} - \underbrace{[\log(1-\tau) + \log w(L_M) - \gamma \log C(L_M)]}_{\text{log marginal utility benefit of after-tax earnings}}. \quad (34)$$

The first term, log marginal disutility of working, is increasing in L_M . The second term, log marginal utility benefit of after-tax earnings, is decreasing in L_M . Write

$$\log(1-\tau) + \log w(L_M) - \gamma \log C(L_M) = \log(1-\tau) + (1-\gamma) \log w(L_M) - \gamma \log \left(\frac{C(L_M)}{w(L_M)} \right).$$

Note that both $\log w(L_M)$ and $C(L_M)/w(L_M)$ are increasing in L_M . Since $\gamma > 1$, its derivative is therefore

$$\frac{d}{dL_M} [\log(1-\tau) + \log w(L_M) - \gamma \log C(L_M)] = (1-\gamma) \frac{d \log w(L_M)}{dL_M} - \gamma \frac{d}{dL_M} \log \left(\frac{C(L_M)}{w(L_M)} \right) < 0.$$

Thus $\mathcal{R}(L_M)$ is continuous and strictly increasing throughout the interval $(\underline{L}_M, \infty)$.


As L_M approaches the lower feasible endpoint, $w(L_M)$ approaches zero, while $C(L_M)/w(L_M)$ approaches a finite, strictly positive value. Since $\gamma > 1$, log marginal benefit therefore approaches $+\infty$, whereas log marginal disutility remains finite. Hence $\mathcal{R}(L_M)$ approaches $-\infty$.

As $L_M \rightarrow \infty$, both $w(L_M)$ and $C(L_M)/w(L_M)$ approach infinity. Log marginal benefit therefore approaches $-\infty$, while log marginal disutility approaches $+\infty$. Hence $\mathcal{R}(L_M)$ approaches $+\infty$.

By continuity and the limits' behavior, $\mathcal{R}(L_M) = 0$ has a solution; by strict monotonicity, that solution is unique. The remaining quantities C, K, K_G, Y, w, r are uniquely recovered from the equations above. Therefore, conditional on fixed L_C , there exists exactly one steady-state equilibrium. $\square$

To illustrate multiplicity, we choose $\gamma = 2$, $\mu = 0.5$, $\zeta = 1$, and $\alpha = \varepsilon = \theta = L_C = 1$. These values violate the conditions for a single equilibrium, placing the economy in a region where two steady states are possible. We set $\beta = 0.8$, $\delta = 0.25$, $\delta_G = 0.5$, $\tau = 0.5$, and $\xi = 0$ to keep the remaining relationships simple. Finally, productivity $A = 7$ makes the residual's minimum negative. Figure A.1 shows two positive steady-state equilibria at $L_M = 3$ and $L_M \approx 6.65$.

Figure A.1: Residual Function $\mathcal{R}(L_M)$ and Two Steady-State Equilibria



Notes: Steady states are zeros of the residual $\mathcal{R}(L_M)$: the log marginal disutility of work minus the log marginal utility of after-tax earnings, as a function of modern-sector labor L_M . Parameters are $\gamma = 2$, $\mu = 0.5$, $\zeta = 1$, $\alpha = \varepsilon = \theta = L_C = 1$, $\beta = 0.8$, $\delta = 0.25$, $\delta_G = 0.5$, $\tau = 0.5$, $\xi = 0$, and $A = 7$. These values violate the conditions for a unique steady state, and the residual has two zeros, at $L_M = 3$ and $L_M \approx 6.65$.

A.2 Efficient Social Planner Benchmark

Formally, the planner chooses $\{C_t, L_t, Y_t, K_{t+1}, K_{G,t+1}, L_{M,t}, L_{C,t}\}_{t=0}^{\infty}$ to maximize

$$\sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\gamma}}{1-\gamma} - \alpha_M \frac{L_t^{1+1/\varepsilon}}{1+1/\varepsilon} \right]$$

subject to the aggregate resource constraint:

$$Y_t + (1 - \delta)K_t + (1 - \delta_G)K_{G,t} - C_t - K_{t+1} - K_{G,t+1} \geq 0$$

where $L_t = L_{M,t} + L_{C,t}$ and $Y_t = A_M(\theta K_{G,t})^\zeta K_t^\mu L_{M,t}^{1-\mu}$ each period and given K_0 and $K_{G,0}$.

Proof of Proposition 5. The planner's optimality conditions are:

$$\begin{aligned} C_t^{-\gamma} &= \beta C_{t+1}^{-\gamma} \left(\mu \frac{Y_{t+1}}{K_{t+1}} + 1 - \delta \right), \\ C_t^{-\gamma} &= \beta C_{t+1}^{-\gamma} \left(\zeta \frac{Y_{t+1}}{K_{G,t+1}} + 1 - \delta_G \right), \\ L_t^{1/\varepsilon} &= C_t^{-\gamma} (1 - \mu) \frac{Y_t}{L_{M,t}}, \end{aligned}$$

which govern optimal investment in private capital, public capital, and deployment of labor inputs. The public capital optimality condition gives the social planner's efficient public capital-output ratio:

$$\frac{K_G^{SP}}{Y^{SP}} = \frac{\zeta}{\beta^{-1} - 1 + \delta_G}.$$

For the competitive equilibrium, recall that, in steady state, the depreciation of public capital equals the net tax revenue:

$$\delta_G K_G = w \left[(1 - \xi) \tau (L_M + L_C) - L_C \right]$$


where $w = (1 - \mu) \frac{Y}{L_M}$. Rearranging, we obtain the competitive equilibrium public capital-output ratio:

$$\frac{K_G^{CE}}{Y^{CE}} = \frac{(1 - \mu)}{\delta_G} \cdot \frac{(1 - \xi) \tau (L_M + L_C) - L_C}{L_M}.$$

□


A.3 Appendix Figures and Tables

Figure A.2: Statutory Tax Rates by GDP per Capita



Note: This figure plots predicted average statutory rates for personal income taxes (PIT), corporate income taxes (CIT), value-added taxes (VAT), and tariffs by GDP per capita. Source: [Bachas et al. \(2024b\)](#).

Figure A.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 \(2014\)](#).

A.4 Unmechanized Agriculture

Agricultural employment data pose a measurement challenge: in rich countries, agriculture is mechanized and taxed, making it a poor proxy for the untaxed traditional sector we aim to model. We therefore construct an alternative target based on unmechanized agricultural employment:

$$\text{Unmechanized Agricultural Employment}_i = \text{Agricultural Employment Share}_i \times (1 - \text{Mechanization}_i)$$

We model mechanization as linear in log GDP per capita, set the four-country average for Malawi, Nigeria, Tanzania, and Uganda to 31.9% according to [Sheahan and Barrett \(2014\)](#), and set the richest country to 100%. We then recalibrate A_T to match this constructed moment.

Figure A.4 shows the recalibrated productivities. Tax productivity is increasing in GDP per capita.

Figure A.4: Productivities with Traditional Sector Targeted to Unmechanized Agricultural Sector



Notes: As in Figure 8, but with traditional-sector productivity calibrated to unmechanized agricultural employment: the agricultural employment share times one minus a mechanization rate that is linear in log GDP per capita, equal to 31.9 percent for the average of Malawi, Nigeria, Tanzania, and Uganda ([Sheahan and Barrett, 2014](#)) and to 100 percent for the richest country.

Table A.2 reports the policy counterfactuals under the unmechanized calibration.

A.5 The Elasticity of Output to Tax Productivity

A.5.1 The Environment

Everything is static. There is no private capital, no depreciation, and no discounting. A fixed measure of workers is split between goods production and tax collection, $L = L_M + L_C$, with L_C set by the government. Output is made from labor and public capital,

$$Y = A_M (\theta K_G)^\zeta L_M, \quad 0 < \zeta < 1, \quad (35)$$

Table A.1: Robustness Checks: Determinants of Tax Productivity A_C

Panel A: Alternative tax-technology measures						
	(1)	(2)	(3)	(4)	(5)	(6)
Combined index	0.445*** (0.137)				0.294** (0.131)	-0.039 (0.076)
3 rd -party data and detection		0.391*** (0.134)				
Identification capacity			0.434** (0.164)			
Collection				0.110 (0.137)		
Controls	No	No	No	No	D_i	D_i
Dependent variable	A_C	A_C	A_C	A_C	A_C	A_M
Observations	59	59	59	59	59	59
R-squared	0.144	0.129	0.108	0.011	0.508	0.744

Panel B: Other ISORA dimensions				
	(1)	(2)	(3)	(4)
Taxpayer services	0.182 (0.136)			
Human resources		0.009 (0.042)		
Institutional autonomy			0.040 (0.043)	
Large-taxpayer unit				-0.027 (0.034)
Controls	No	No	No	No
Dependent variable	A_C	A_C	A_C	A_C
Observations	96	96	96	96
R-squared	0.015	0.001	0.005	0.005

Panel C: Imputed economy-wide controls					
	(1)	(2)	(3)	(4)	(5)
Combined index	0.336*** (0.118)				
3 rd -party data and digitization		0.239** (0.100)			
Enforcement			0.182** (0.075)		
Operational processes				0.199** (0.076)	-0.035 (0.067)
Controls	D_i	D_i	D_i	D_i	D_i
Dependent variable	A_C	A_C	A_C	A_C	A_M
Observations	97	97	97	97	97
R-squared	0.647	0.634	0.634	0.636	0.802

Notes: The specification, dependent variables and D_i controls are the same as in Table 4. Columns (2) to (4) of Panel A reports coefficients for the three technology dimensions used in [Okunogbe and Tourek \(2024\)](#), based on the OECD's Inventory of Tax Technology Initiatives. Column (1) uses the unweighted average across the three categories as the combined index regressor. Column (5) augments column (1) with D_i controls. Column (6) is identical to column (5), except the outcome is A_M . Panel B studies the remaining four ISORA categories that were not used in Table 4. All regressors are indices; within each category, each variable is standardized and the regressor for that index is the unweighted average over all variables. In Panel C, missing values of each D_i variable are imputed using predicted values from a regression of that variable on log GDP per capita and region dummies, estimated on the sample with non-missing values of D_i . Robust standard errors are reported. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table A.2: Policy Counterfactuals: Unmechanized-Agriculture Calibration

Shock	Income group	GDP	Consumption	K_G	After-tax wage	T/X	$\Delta\tau$ (pp)
$A_C \times 1.5$	Low income	1.13	1.11	1.89	1.18	1.54	2.14
$A_C \times 1.5$	Middle income	1.14	1.09	1.73	1.04	1.46	9.05
$A_C \times 1.5$	High income	1.08	1.00	1.64	0.89	1.40	15.09
$A_M \times 1.5$	Low income	1.95	1.87	2.92	2.51	1.15	-0.00
$A_M \times 1.5$	Middle income	2.27	2.27	2.28	2.34	0.98	0.00
$A_M \times 1.5$	High income	2.28	2.28	2.28	2.34	0.97	-0.00
$L_C \times 1.5$	Low income	1.00	1.00	1.01	1.00	0.68	0.09
$L_C \times 1.5$	Middle income	1.00	1.00	1.01	1.00	0.68	0.37
$L_C \times 1.5$	High income	1.00	1.00	1.01	0.99	0.68	0.66
$\theta = 0.84$	Low income	1.12	1.11	1.23	1.19	1.03	-0.00
$\theta = 0.84$	Middle income	1.05	1.05	1.05	1.05	1.00	0.00
$\theta = 0.84$	High income	1.00	1.00	1.00	1.00	1.00	-0.00

Notes: GDP, consumption, public capital, after-tax wages and T/X are mean country counterfactual/baseline factors. The final column is an actual percentage-point change: $100(\tau' - \tau)$, not a fraction-valued difference. The baseline calibration has 97 countries. Solved low/middle/high-income samples by shock: $A_C \times 1.5$: 10/46/39; $A_M \times 1.5$: 10/46/41; $L_C \times 1.5$: 10/46/41; $\theta = 0.84$: 10/46/41. Transfer shares remain fixed. The $A_C \times 1.5$ counterfactual is inadmissible for Denmark, Sweden (2 cases), because the implied tax rate violates $0 \leq \tau' < 1$. These cases are excluded from the shock means, not assigned zero effects.

so there are constant returns in the private input L_M and returns to scale of $1 + \zeta$ once public capital is counted. Firms take K_G as given, so $w = \partial Y / \partial L_M = Y / L_M$ and the wage bill exhausts output.

Tax collectors produce the tax rate,

$$\tau = A_C L_C^\eta, \quad 0 < \eta < 1, \quad \tau < 1, \quad (36)$$

revenue is $T = \tau Y$, and the government splits it between transfers and public capital:

$$TR = \xi T, \quad K_G = (1 - \xi) T = (1 - \xi) \tau Y. \quad (37)$$

Transfers are lump sum. The household owns its labor income, pays the tax, and receives the transfer, so its budget is $C = (1 - \tau)wL_M + TR$; since $wL_M = Y$ this is

$$C = (1 - \varphi)Y, \quad \text{where} \quad \varphi \equiv (1 - \xi) \tau. \quad (38)$$

Consumption is output minus public investment, as it must be: transfers move resources around but do not use any up.

τ is the effective tax rate that the household faces at the margin when it decides how much to work.

$\varphi = (1 - \xi)\tau$ is the investment rate: the share of output that becomes public capital.

The Domar weight. The tax authority has no sales share. Valuing its output at the cost of its inputs, as the national accounts value any non-market producer, its Domar weight is its wage bill share,

$$\lambda_C \equiv \frac{wL_C}{Y} = \frac{L_C}{L_M}. \quad (39)$$

This is the number we would use when applying Hulten's theorem. In the low-income group $L_C/L \approx 0.072$ percent, so $\lambda_C \approx 0.00072$.

Aggregate output. Substituting (37) into (35) and solving the fixed point in Y ,

$$Y = \Omega L_M^{\frac{1}{1-\zeta}}, \quad \Omega \equiv [A_M \theta^\zeta \varphi^\zeta]^{\frac{1}{1-\zeta}}, \quad w = \frac{Y}{L_M} = \Omega L_M^{\frac{\zeta}{1-\zeta}}. \quad (40)$$

A solution exists and is unique whenever $\zeta < 1$.

Throughout, write

$$\Lambda_0 \equiv \frac{\zeta}{1-\zeta}, \quad W \equiv \frac{\tau}{1-\tau}, \quad W_c \equiv \frac{\varphi}{1-\varphi}, \quad s \equiv \frac{L_M}{L}. \quad (41)$$

At the benchmark $\zeta = 0.20$, so $\Lambda_0 = 0.25$.

A.5.2 Inelastic Labor

Fix L , so $L_M = L - L_C$ is fixed as well. Only Ω varies with A_C .

Proposition 7. *With labor inelastic,*

$$\frac{d \log Y}{d \log A_C} = \frac{\zeta}{1-\zeta} = \Lambda_0, \quad (42)$$

independent of ξ , of λ_C , of η , and of the level of τ .

Proof. Take logs of (40): $(1 - \zeta) \log Y = \log A_M + \zeta \log \theta + \zeta \log \varphi + \log L_M$. With ξ fixed, $d \log \varphi = d \log \tau = d \log A_C$ by (36), and $d \log L_M = 0$. $\square$

At $\zeta = 0.20$ the elasticity is 0.25, against a Domar weight of $\lambda_C = 0.00072$: a factor of 347. Raising A_C by fifty percent raises output by $1.5^{0.25} - 1 = 10.7$ percent, where Hulten's theorem predicts $1.5^{\lambda_C} - 1 = 0.03$ percent.

Hulten's theorem is an envelope result (Hulten, 1978; Baqaee and Farhi, 2019): at an efficient allocation prices are the multipliers on the resource constraints, so the reallocation induced by a

productivity shock has no first-order value and only the direct effect, valued at sales, survives. Two hypotheses fail here. First, the tax authority has no sales. Second, public capital is unpriced: its marginal product is $\zeta Y/K_G$, but firms take K_G as given, and no one is paid for supplying it. So the allocation is inefficient, and the discarded reallocation terms matter.

Consumption and the Barro threshold. From (38),

$$\frac{d \log C}{d \log A_C} = \underbrace{\Lambda_0}_{\text{more public capital}} - \underbrace{W_c}_{\text{more investment out of output}}, \quad (43)$$

positive if and only if $\varphi < \zeta$. The consumption-maximizing policy is $\varphi^* = \zeta$, the classic result of Barro (1990): invest in the public input up to the point where its output elasticity equals the share of output taken to pay for it.

For our income group calibrations, we have:

	τ	ξ	$\varphi = (1 - \xi)\tau$	$d \log Y / d \log A_C$	$d \log C / d \log A_C$
Low income	0.032	0.260	0.024	0.250	+0.226
Middle income	0.236	0.415	0.138	0.250	+0.090
High income	0.335	0.488	0.171	0.250	+0.043

A.5.3 Elastic Labor Supply

Now let the household choose how much to work. The government still conscripts L_C , so an extra unit of household labor goes into production: $L = L_M + L_C$ with L_C fixed, hence $dL = dL_M$.

Only one thing changes in (40): L_M is no longer a constant. Differentiating $\log Y = \log \Omega + \frac{1}{1-\zeta} \log L_M$,

$$\frac{d \log Y}{d \log A_C} = \underbrace{\frac{\zeta}{1-\zeta}}_{\text{the public-capital multiplier}} + \underbrace{\frac{1}{1-\zeta} \frac{d \log L_M}{d \log A_C}}_{\text{the labor-supply term}} \quad (44)$$

The two terms are intuitive. The first is the output elasticity with inelastic labor supply. The second captures that if the household works more or less, output changes— amplified by the same multiplier $1/(1 - \zeta)$.

We now sign and size the $d \log L_M / d \log A_C$ term.

A.5.4 Solving for the Labor Supply Response

The household has preferences $C^{1-\gamma}/(1-\gamma) - \alpha L^{1+1/\varepsilon}/(1+1/\varepsilon)$ and chooses L taking the tax rate τ , public capital K_G , and the transfer TR as given. It does not internalize that its own labor raises output, hence revenue, hence the transfer it receives, and it does not internalize that its labor raises public capital. Its first-order condition is therefore

$$\alpha L^{1/\varepsilon} = C^{-\gamma}(1-\tau)w. \quad (45)$$

Using $C = (1-\varphi)\Omega L_M^{1/(1-\zeta)}$ and $w = \Omega L_M^{\zeta/(1-\zeta)}$ from (40) and (38),

$$\alpha L^{1/\varepsilon} = (1-\varphi)^{-\gamma}(1-\tau)\Omega^{1-\gamma}L_M^{\frac{\zeta-\gamma}{1-\zeta}}. \quad (46)$$

This is one equation in L_M . Note that τ and φ now enter separately: τ through the after-tax wage, φ through consumption.

We now log-differentiate. Hold L_C and ξ fixed, so $d\log\varphi = d\log\tau = d\log A_C \equiv \hat{a}_C$, and recall $d\log\Omega = \Lambda_0\hat{a}_C$, $d\log(1-\tau) = -W\hat{a}_C$, $d\log(1-\varphi) = -W_c\hat{a}_C$. On the left, $d\log L = s d\log L_M$ because $dL = dL_M$. Then

$$\frac{s}{\varepsilon}\hat{L}_M = [\gamma W_c - W + (1-\gamma)\Lambda_0]\hat{a}_C + \frac{\zeta-\gamma}{1-\zeta}\hat{L}_M, \quad (47)$$

and collecting terms gives the result.

Proposition 8. *With elastic labor and transfers,*

$$\frac{d\log L_M}{d\log A_C} = \frac{\gamma W_c - W + (1-\gamma)\Lambda_0}{D}, \quad D \equiv \frac{s}{\varepsilon} + \frac{\gamma-\zeta}{1-\zeta} > 0, \quad (48)$$

and therefore

$$\frac{d\log Y}{d\log A_C} = \frac{\zeta}{1-\zeta} + \frac{1}{1-\zeta} \cdot \frac{\gamma W_c - W + (1-\gamma)\Lambda_0}{\frac{s}{\varepsilon} + \frac{\gamma-\zeta}{1-\zeta}}. \quad (49)$$

The three terms are the three forces on labor supply. First, $-W = -\tau/(1-\tau)$ is the substitution effect. A better tax authority raises the statutory rate, the after-tax wage falls, and the household works less. Second, $+\gamma W_c = \gamma\varphi/(1-\varphi)$ is the income effect from paying for the public capital. More output is diverted into public investment, consumption is lower than it otherwise would be, the marginal utility of consumption is higher, and the household works more. It is scaled by γ because γ governs how fast marginal utility rises as consumption falls. Third, $(1-\gamma)\Lambda_0$ is the

wealth effect from being richer. Public capital raises output by Λ_0 ; with $\gamma > 1$ income effects dominate substitution effects, this is negative, and the household takes some of the gain as leisure.

The denominator D captures how hard it is to change labor: a smaller Frisch elasticity ε or a larger γ shrinks the response.

Why transfers change the sign. With $\xi = 0$ we have $W_c = W$ and (48) collapses to $(1 - \gamma)(\Lambda_0 - W)/D$, which is positive in rich countries: consumption falls, and with $\gamma > 1$ the household responds by working more. With $\xi > 0$ we have $W_c < W$, so the substitution term $-W$ is no longer offset by an equally large income effect, and labor falls. This is what makes the labor response negative, and with ξ increasing with income, the effect is strongest in rich countries.

The no-transfer benchmark. In the case $\xi = 0$, $\varphi = \tau$ and the elasticity in (49) reduces, at $s \approx 1$, to the simple $0.24 + 0.04W$, so it rises mechanically with the tax rate:

$\xi = 0$	τ	W	$d \log L / d \log A_C$	$d \log Y / d \log A_C$	$d \log C / d \log A_C$
Low income	0.032	0.033	-0.007	0.241	+0.208
Middle income	0.236	0.309	+0.002	0.252	-0.056
High income	0.335	0.503	+0.008	0.260	-0.243

Labor supply rises in response to a tax productivity shock in middle- and high-income countries, and the output elasticity is larger there than in poor countries. The coefficient 0.04 is $(\gamma - 1)/[(1 - \xi)D]$, which is small only because γ is barely above one; the entire no-transfer labor channel is worth less than two log points even at the highest tax rate in the sample.

A.5.5 The Dynamic Model with Inelastic Labor Supply

Our full model is dynamic, has private capital, and depreciates both capital stocks. Holding labor supply fixed — which, as Section A.5.4 shows, costs only a few log points — the same argument goes through in closed form.

The Euler equation pins the marginal product of private capital at $r = \mu Y/K = \beta^{-1} - 1 + \delta$, so the capital-output ratio is a constant and $d \log K = d \log Y$: private capital accumulates in proportion to output. The government budget is $\delta_G K_G = w[(1 - \xi)\tau(L_M + L_C) - L_C]$ with $w = (1 - \mu)Y/L_M$, so public capital is also proportional to output, and with labor fixed $d \log K_G = d \log Y + d \log \Phi$, where $\Phi \equiv (1 - \xi)\tau L - L_C$ is revenue net of the cost of collecting it, measured in labor units.

Differentiating the production function $Y = A_M(\theta K_G)^\zeta K^\mu L_M^{1-\mu}$ and substituting both,

$$\frac{d \log Y}{d \log A_C} = \frac{\zeta \Psi}{1 - \mu - \zeta}, \quad \Psi \equiv \frac{(1 - \xi) \tau L}{(1 - \xi) \tau L - L_C} = 1 + \frac{\lambda_C}{s_G} \geq 1, \quad (50)$$

where $s_G \equiv \delta_G K_G / Y$ is the share of output going to public investment. An interior steady state requires $\mu + \zeta < 1$, which at $\mu = 0.33$ and $\zeta = 0.20$ is comfortably satisfied.

Expression (50) differs from the static (42) in exactly two ways. Private capital adds μ to the denominator, because it is a second stock that expands in proportion to output and so joins the same feedback loop; this alone raises the multiplier from $\zeta / (1 - \zeta) = 0.25$ to $\zeta / (1 - \mu - \zeta) = 0.426$. And $\Psi > 1$ corrects for the fact that collectors must now be paid out of revenue, so only revenue net of collection costs buys public capital. Note where the Domar weight finally appears: inside Ψ , as the ratio λ_C / s_G of the collection wage bill to public investment. It enters as a small correction to the multiplier, not as the multiplier itself.

At our income group calibrations:

	τ	ξ	Ψ	$d \log Y / d \log A_C$
Low income	0.032	0.260	1.031	0.439
Middle income	0.236	0.415	1.016	0.432
High income	0.335	0.488	1.016	0.432

The correction Ψ never reaches four percent, so to a good approximation the elasticity is simply $\zeta / (1 - \mu - \zeta) = 0.43$ — which is the familiar steady-state elasticity of output to the tax rate in a Barro (1990) economy with capital accumulation, a long-run margin emphasized by Baqaee and Malmberg (2025).

That the elasticity is essentially the same in all three rows with labor held fixed.

A.5.6 Endogenous Collection Labor in the Dynamic Model

Section A.5.5 holds L_C at its calibrated value. Suppose the government instead chooses L_C to maximize household welfare, moving before the household as in the extension. Output then responds to A_C both directly and through the induced change in the size of the tax authority, so that

$$\frac{d \log Y}{d \log A_C} = \Lambda + (\eta \Lambda - \kappa) \frac{d \log L_C}{d \log A_C}, \quad \kappa \equiv \frac{1}{1 - \zeta} \cdot \frac{(1 - s)/\varepsilon}{D}, \quad (51)$$

where Λ is the fixed- L_C elasticity of (49) and κ is the output cost of withdrawing a worker from production.

To derive (51), let both A_C and L_C vary, and write $\hat{a}_C \equiv d \log A_C$ and $\hat{L}_C \equiv d \log L_C$. The tax rate and the investment rate now move by $d \log \tau = d \log \varphi = \hat{a}_C + \eta \hat{L}_C$, and total labor by $d \log L = s \hat{L}_M + (1-s) \hat{L}_C$, since $dL = dL_M + dL_C$. Log-differentiating the household's condition (46) as in Section A.5.4, with these two changes, gives

$$\frac{s}{\varepsilon} \hat{L}_M + \frac{1-s}{\varepsilon} \hat{L}_C = N(\hat{a}_C + \eta \hat{L}_C) + \frac{\zeta - \gamma}{1 - \zeta} \hat{L}_M, \quad N \equiv \gamma W_c - W + (1 - \gamma) \Lambda_0,$$

so that

$$\hat{L}_M = \frac{N}{D}(\hat{a}_C + \eta \hat{L}_C) - \frac{1-s}{\varepsilon D} \hat{L}_C. \quad (52)$$

The first term is the labor response of Proposition 8 to a change in the tax rate, whatever its source. The second is new: at a given tax rate, more collectors raise total labor and hence the marginal disutility of work, and households respond by supplying less labor to production.

By (40), $d \log Y = \Lambda_0 d \log \varphi + \hat{L}_M / (1 - \zeta)$. Substituting (52),

$$\begin{aligned} d \log Y &= \left(\Lambda_0 + \frac{N}{(1 - \zeta) D} \right) (\hat{a}_C + \eta \hat{L}_C) - \frac{1}{1 - \zeta} \cdot \frac{1-s}{\varepsilon D} \hat{L}_C \\ &= \Lambda \hat{a}_C + (\eta \Lambda - \kappa) \hat{L}_C, \end{aligned}$$

where the term in parentheses is the fixed- L_C elasticity Λ of (49), evaluated at the government's choice of L_C . Dividing by $\hat{a}_C$ gives (51). The coefficient on $\hat{L}_C$ is the output effect of hiring collectors. A one percent increase in L_C raises the tax rate by η percent, which has the same effect on output as an η percent increase in A_C , hence the term $\eta \Lambda$; from this we subtract κ , the output lost as the additional collectors crowd out production labor.

We still need to determine $d \log L_C / d \log A_C$, which comes from the government's problem. The government chooses L_C to maximize $U = C^{1-\gamma} / (1 - \gamma) - \alpha L^{1+1/\varepsilon} / (1 + 1/\varepsilon)$, taking the household's response as given, so its first-order condition is $C^{1-\gamma} \partial \log C / \partial \log L_C = \alpha L^{1+1/\varepsilon} \partial \log L / \partial \log L_C$. Using (45) to replace $\alpha L^{1/\varepsilon}$ by $C^{-\gamma} (1 - \tau) w$, together with $C = (1 - \varphi) Y$ and $w = Y / L_M$, this becomes

$$\frac{\partial \log C}{\partial \log L_C} = \frac{1 - \tau}{(1 - \varphi) s} \cdot \frac{\partial \log L}{\partial \log L_C}, \quad (53)$$

where, from (38), (52), and the expression for $d \log Y$ above,

$$\frac{\partial \log C}{\partial \log L_C} = \eta (\Lambda - W_c) - \kappa, \quad \frac{\partial \log L}{\partial \log L_C} = \frac{s (\eta N - (1 - s) / \varepsilon)}{D} + 1 - s.$$

The government hires collectors until the consumption they generate equals the additional work they require, valued at the household's marginal rate of substitution between consumption and labor, the after-tax wage $(1 - \tau)w$.

A.5.7 Inelastic Labor Supply with Endogenous Collection

The same exercise with labor inelastic isolates the mechanism, because the disutility of work is then a constant, and the government maximizes $C = (1 - \varphi)Y$. Its first-order condition in L_C is

$$\eta(\Lambda_0 - W_c) = \frac{\lambda_C}{1 - \zeta}, \quad (54)$$

which equates η log points of tax rate, each worth $\Lambda_0 - W_c$ in consumption by (43), to the output cost of the collector who delivers them. The left side falls in L_C and the right side rises, so the optimum is unique and always satisfies $\varphi^* < \zeta$. Differentiating (54) through,

$$\frac{d \log Y}{d \log A_C} = \frac{\zeta}{1 - \zeta} - \frac{\eta W_c R}{\eta R + (\Lambda_0 - W_c)/s}, \quad R \equiv W_c(1 + W_c) = \frac{\varphi}{(1 - \varphi)^2}, \quad (55)$$

which returns $\zeta/(1 - \zeta)$ as $\eta \rightarrow 0$ and tends to $\Lambda_0 - W_c$ as η grows.

$d \log Y / d \log A_C$	$\eta = 0.05$	$\eta = 0.20$	$\eta = 0.50$	$\eta = 1$
Low income	0.250	0.247	0.148	0.013
Middle income	0.234	0.136	0.027	0.007
High income	0.209	0.094	0.022	0.007

Each entry solves (54) exactly and reproduces a numerical derivative through the re-optimized L_C to nine decimal places. The mechanism is visible in the optimal investment rate, which in low-income countries runs 0.027, 0.050, 0.150 and 0.195 across the four columns: as η rises the government drives φ^* toward $\zeta = 0.20$, the Barro optimum of (43). At that point the marginal social value of public capital is zero by construction, so a more productive tax authority buys nothing in output and permits a smaller payroll. How close the government can get is exactly what η governs. At $\eta = 0.05$ it gets nowhere — $\varphi^* = 0.027$ against a target of 0.20 — the economy remains far short of the efficient public capital stock, additional revenue is still valuable, and the elasticity is 0.250 against 0.25 with L_C fixed.

A.6 The Ramsey Problem

This appendix sets out the Ramsey extension of Section 6.1. It shows that the government's problem reduces to a choice of allocations subject to the resource constraint and the law of motion of public capital, proves Proposition 6, derives the closed forms and the income-group results reported in the main text, and shows how the optimality conditions change when collection is financed outside the public-investment budget.

Environment. Technologies are as in the baseline. Output is $Y_t = F(K_t, K_{G,t}, L_{M,t})$, with constant returns in $(K_t, L_{M,t})$, and the tax rate is produced by collectors, $\tau_t = \tau(L_{C,t}) = A_C L_{C,t}^\eta$. Firms take $K_{G,t}$ as given and pay $r_t = F_K(t)$ and $w_t = F_{L_M}(t)$, where $F_x(t)$ denotes a partial derivative evaluated at $(K_t, K_{G,t}, L_{M,t})$. A representative household supplies a fixed amount of labor, $L = L_{M,t} + L_{C,t}$, and is paid the market wage in production and in collection, so it is willing to supply any split of L between the two. It solves

$$\begin{aligned} \max_{\{C_t, K_{t+1}, b_{t+1}\}} \quad & \sum_{t=0}^{\infty} \beta^t u(C_t) \\ \text{s.t.} \quad & C_t + K_{t+1} + b_{t+1} = (1 - \omega_t)(1 + r_t - \delta)K_t + R_t b_t + (1 - \tau_t)w_t L + TR_t, \end{aligned} \quad (56)$$

given K_0 and b_0 , where b_t is government debt held by the household, R_t is its gross return, and ω_t is a wedge on the gross return to private capital; ω_0 is a levy on initial wealth.

The government has two accounts. The public-investment account is the budget constraint of the baseline model. Collected taxes, $T_t = \tau_t w_t L$, finance transfers, $TR_t = \xi T_t$, collection spending, $X_t = w_t L_{C,t}$, and public investment:

$$K_{G,t+1} + TR_t + X_t = T_t + (1 - \delta_G)K_{G,t} \iff K_{G,t+1} = (1 - \delta_G)K_{G,t} + \Phi(L_{C,t})w_t, \quad (57)$$

with $\Phi(L_C) \equiv (1 - \xi)\tau(L_C)L - L_C$. This is the law of motion of public capital in Section 3 and in (28). The general account borrows and sets the capital wedge,

$$R_t b_t = b_{t+1} + \omega_t(1 + r_t - \delta)K_t, \quad (58)$$

and, by assumption, cannot transfer resources to the public-investment account: public investment, transfers, and collection are financed only out of collected taxes. Adding (56), (57), and (58), and using $F = r_t K_t + w_t L_{M,t}$, gives the resource constraint

$$C_t + K_{t+1} + K_{G,t+1} = F(K_t, K_{G,t}, L - L_{C,t}) + (1 - \delta)K_t + (1 - \delta_G)K_{G,t}. \quad (59)$$

Collectors' pay does not appear in (59): their wages are household income, and the resource cost of collection is the output that collectors would otherwise produce. A competitive equilibrium is an allocation $\{C_t, K_{t+1}, K_{G,t+1}, L_{C,t}\}$, prices $\{r_t, w_t, R_t\}$, and policies $\{\omega_t, b_{t+1}\}$ such that the household solves (56), firms maximize profits, and (57)–(59) hold.

Household optimality. The household's first-order conditions are

$$u'(C_t) = \beta u'(C_{t+1})(1 - \omega_{t+1})(1 + r_{t+1} - \delta), \quad u'(C_t) = \beta u'(C_{t+1})R_{t+1}, \quad (60)$$

together with the transversality conditions $\lim_{t \rightarrow \infty} \beta^t u'(C_t) K_{t+1} = \lim_{t \rightarrow \infty} \beta^t u'(C_t) b_{t+1} = 0$. The two Euler equations imply the arbitrage condition $R_{t+1} = (1 - \omega_{t+1})(1 + r_{t+1} - \delta)$: the return on government debt equals the return households earn on private capital. In a steady state, consumption is constant and $\beta R = 1$.

Lemma 1 (Implementability). *In any competitive equilibrium,*

$$\sum_{t=0}^{\infty} \beta^t u'(C_t) [C_t - (1 - \tau_t)w_t L - TR_t] = u'(C_0) [(1 - \omega_0)(1 + r_0 - \delta)K_0 + R_0 b_0]. \quad (61)$$

Proof. Multiply the period- t budget constraint in (56) by $\beta^t u'(C_t)$ and sum from $t = 0$ to T . By (60),

$$\begin{aligned} \beta^t u'(C_t) K_{t+1} &= \beta^{t+1} u'(C_{t+1})(1 - \omega_{t+1})(1 + r_{t+1} - \delta) K_{t+1}, \\ \beta^t u'(C_t) b_{t+1} &= \beta^{t+1} u'(C_{t+1}) R_{t+1} b_{t+1}, \end{aligned}$$

so the asset terms telescope and

$$\begin{aligned} \sum_{t=0}^T \beta^t u'(C_t) [C_t - (1 - \tau_t)w_t L - TR_t] \\ = u'(C_0) [(1 - \omega_0)(1 + r_0 - \delta)K_0 + R_0 b_0] - \beta^T u'(C_T) (K_{T+1} + b_{T+1}). \end{aligned}$$

Letting $T \rightarrow \infty$, the last term vanishes by the transversality conditions. $\square$

Lemma 2 (Reduction). *An allocation $\{C_t, K_{t+1}, K_{G,t+1}, L_{C,t}\}$ that converges to a steady state is part of a competitive equilibrium if and only if it satisfies the resource constraint (59) and the law of motion (57) with $w_t = F_{L_M}(t)$, provided that the levy on initial wealth it requires is feasible.*

Proof. Necessity is immediate. For sufficiency, set $r_t = F_K(t)$, $w_t = F_{L_M}(t)$, $\tau_t = A_C L_{C,t}^\eta$, $T_t = \tau_t w_t L$, and $TR_t = \xi T_t$, so that the public-investment account holds by (57). Choose ω_{t+1} so that the

household's capital Euler equation in (60) holds along the given consumption path, and set R_{t+1} by the arbitrage condition. Choose ω_0 so that (61) holds, and define b_{t+1} recursively from (56). Because (61) holds and the allocation converges, the transversality conditions are satisfied, so the household's plan is optimal. The general account (58) then holds by Walras' law, since (59) is the sum of the three budget constraints. $\square$

The government's problem. By Lemma 2, the implementability condition does not constrain the government, which solves

$$\max_{\{C_t, K_{t+1}, K_{G,t+1}, L_{C,t}\}} \sum_{t=0}^{\infty} \beta^t u(C_t)$$

subject to (59) and (57), with $w_t = F_{LM}(K_t, K_{G,t}, L - L_{C,t})$, for all t , given K_0 and $K_{G,0}$. The levy on initial wealth does not let the government escape the cost of collection, because it cannot finance public investment.

Proof of Proposition 6. Attach multipliers $\beta^t \chi_t$ to (59) and $\beta^t v_t$ to (57), and let $q_t \equiv v_t / \chi_t$ be the value of a unit of funds in the public-investment account relative to a unit of output. The first-order conditions for C_t , $L_{C,t}$, $K_{G,t+1}$, and K_{t+1} are

$$u'(C_t) = \chi_t, \tag{62}$$

$$\chi_t F_{LM}(t) = v_t \left[\Phi'(L_{C,t}) F_{LM}(t) - \Phi(L_{C,t}) F_{LM L_M}(t) \right], \tag{63}$$

$$\chi_t + v_t = \beta \chi_{t+1} \left[F_{K_G}(t+1) + 1 - \delta_G \right] + \beta v_{t+1} \left[1 - \delta_G + \Phi(L_{C,t+1}) F_{LM K_G}(t+1) \right], \tag{64}$$

$$\chi_t = \beta \chi_{t+1} \left[F_K(t+1) + 1 - \delta \right] + \beta v_{t+1} \Phi(L_{C,t+1}) F_{LM K}(t+1). \tag{65}$$

One more collector moves a worker out of production, which lowers output by F_{LM} and raises the wage by $-F_{LM L_M}$; hence the two terms in (63). Dividing (63) by $\chi_t F_{LM}(t)$,

$$q_t = \frac{1}{\mathcal{M}_t}, \quad \mathcal{M}_t \equiv \Phi'(L_{C,t}) - \Phi(L_{C,t}) \frac{F_{LM L_M}(t)}{F_{LM}(t)} = (1 - \xi) \tau_{L_C}(t) L - 1 - \Phi(L_{C,t}) \frac{F_{LM L_M}(t)}{F_{LM}(t)}. \tag{66}$$

At an interior steady state, allocations and multipliers are constant. Dividing (64) by χ and rearranging,

$$q[1 - \beta(1 - \delta_G)] = \beta(F_{K_G} + 1 - \delta_G) - 1 + \beta q \Phi F_{LM K_G}.$$

Substituting $q = 1/\mathcal{M}$ and using $\beta R = 1$, so that $\beta(F_{K_G} + 1 - \delta_G) - 1 = (F_{K_G} + 1 - \delta_G)/R - 1$, gives (29). Dividing (65) by χ and substituting $q = 1/\mathcal{M}$ gives (30). $\square$

The Ramsey rule and closed forms. Dividing (29) by β and using $[1 - \beta(1 - \delta_G)]/\beta = \beta^{-1} - 1 + \delta_G$ gives (32). With $F = A_M(\theta K_G)^\xi K^\mu L_M^{1-\mu}$, we have $F_{K_G} = \xi Y/K_G$, $F_{L_M} = (1 - \mu)Y/L_M$, $F_{L_M L_M} = -\mu F_{L_M}/L_M$, $F_{L_M K_G} = \xi F_{L_M}/K_G$, and $F_{L_M K} = \mu F_{L_M}/K$. Because $\tau_{L_C} = \eta \tau/L_C$ and, from (10), $T/X = \tau L/L_C$, the first two terms of $\mathcal{M}$ equal $(1 - \xi)\eta T/X - 1$. In a steady state, (57) gives $\delta_G K_G = \Phi F_{L_M}$, so $s_G \equiv \delta_G K_G/Y = (1 - \mu)\Phi/L_M$, and the last term of $\mathcal{M}$ equals $\mu\Phi/L_M = \mu s_G/(1 - \mu)$, which gives (31). The same identity gives $\Phi F_{L_M K_G} = \xi \delta_G$, so (32) becomes

$$\xi \frac{Y}{K_G} + \frac{\xi \delta_G}{\mathcal{M}} = (\beta^{-1} - 1 + \delta_G) \left(1 + \frac{1}{\mathcal{M}}\right),$$

and solving for K_G/Y gives (33). Multiplying (33) by δ_G and solving for $\mathcal{M}$ gives the inversion used in the main text,

$$\mathcal{M} = \frac{s_G(1 - s_G^{SP})}{s_G^{SP} - s_G}.$$

Similarly, $\Phi F_{L_M K} = \mu s_G Y/K$, so (30) becomes $\beta^{-1} - 1 + \delta = \mu(Y/K)(1 + s_G/\mathcal{M})$, or

$$\frac{K^R}{Y^R} = \frac{\mu}{\beta^{-1} - 1 + \delta} \left(1 + \frac{s_G}{\mathcal{M}}\right).$$

The household's Euler equation requires $(1 - \omega)(1 + F_K - \delta) = \beta^{-1}$, while the government's condition puts F_K below $\beta^{-1} - 1 + \delta$, so the implementing wedge is negative: the government subsidizes private capital. As collection becomes costless, $\mathcal{M} \rightarrow \infty$, the subsidy vanishes, and K_G/Y converges to the efficient ratio of Proposition 5.

Income-group results. Table A.3 reports the steady states discussed in the main text. For each income group, we set τ , L_C/L , and ξ to their averages in the baseline calibration and back out A_C from $\tau = A_C L_C^\eta$. The Ramsey steady state then solves two equations in L_C and K_G/Y : the steady-state law of motion, $s_G = (1 - \mu)\Phi(L_C)/(L - L_C)$, and (33), with $\mathcal{M}$ given by (31). Neither equation involves A_M or θ .

Financing collection outside the public-investment budget. Suppose instead that collection spending is paid from the general account, $X_t + R_t b_t = b_{t+1} + \omega_t(1 + r_t - \delta)K_t$, while all non-transferred revenue goes to public investment, $K_{G,t+1} = (1 - \delta_G)K_{G,t} + (1 - \xi)T_t$. To keep the comparison clean, let collectors be hired from outside household labor, so that $L_{M,t} = L$ and $T_t = \tau_t w_t L$, and collection is a resource cost: $C_t + K_{t+1} + K_{G,t+1} + w_t L_{C,t} = F(K_t, K_{G,t}, L) + (1 - \delta)K_t + (1 - \delta_G)K_{G,t}$. The same steps give $q = 1/[(1 - \xi)\tau_{L_C} L] = 1/[(1 - \xi)\eta T/X]$ and, at an interior

Table A.3: Ramsey Steady States by Income Group

	Low income	Middle income	High income
<i>Baseline calibration (group averages)</i>			
Tax rate τ (%)	1.22	16.86	29.86
Collection employment L_C/L (%)	0.030	0.160	0.250
Transfer share ξ	0.260	0.415	0.488
Collections per expenditure T/X	40.7	105.4	119.4
Net marginal revenue $\mathcal{M}$	0.51	2.11	2.11
K_G/Y relative to efficient	0.054	0.597	0.926
<i>Ramsey steady state</i>			
Collection employment L_C/L (%)	0.044	0.216	0.091
Tax rate τ (%)	1.24	17.11	28.39
Collections per expenditure T/X	28.3	79.3	312.6
Net marginal revenue $\mathcal{M}$	0.05	1.35	7.05
Collection markup $1 + 1/\mathcal{M}$	20.6	1.74	1.14
K_G/Y relative to efficient	0.054	0.603	0.888

Notes: The top panel evaluates the baseline allocation at the income-group averages of τ , L_C/L , and ξ , so T/X and $\mathcal{M}$ differ from averages of country-level values. The bottom panel solves the steady-state law of motion of public capital and (33) jointly for L_C and K_G/Y , holding A_C and ξ fixed, with $\beta = 0.96$, $\mu = 0.33$, $\zeta = 0.20$, $\eta = 0.05$, and $\delta = \delta_G = 0.05$. The efficient public capital-output ratio is $\zeta/(\beta^{-1} - 1 + \delta_G) \approx 2.18$.

steady state,

$$\frac{1 - \beta(1 - \delta_G)}{(1 - \xi)\tau_{L_C}L} = \frac{F_{K_G} + 1 - \delta_G}{R} - 1 + \beta \frac{1 - \eta}{\eta} F_{L_M K_G} L_C,$$

$$1 = \beta(F_K + 1 - \delta) + \beta \frac{1 - \eta}{\eta} F_{L_M K} L_C,$$

so the collection markup becomes $1 + 1/[(1 - \xi)\eta T/X]$. In the tax base terms, the larger tax base is valued at q , while the higher pay of collectors is now a private resource cost, and the difference is $q(1 - \xi)\tau L - L_C = L_C(1 - \eta)/\eta$. Relative to Proposition 6, the marginal collector no longer has to be paid out of the revenue they raise, which removes the -1 from the denominator. This matters most where marginal returns are low. At a gross marginal return of 1.6 and no transfers, the collection markup is about 1.6 in this variant but about 2.7 under the baseline accounting, in which taxes pay for their own collection. We use the baseline accounting throughout the paper.