NBER WORKING PAPER SERIES

EMPLOYMENT STRUCTURE AND THE RISE OF THE MODERN TAX SYSTEM

Anders Jensen

Working Paper 25502
http://www.nber.org/papers/w25502

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2019

I thank Oriana Bandiera, Tim Besley, Henrik Kleven and Emmanuel Saez for relentless support on this project. I am grateful to David Autor, Alan Auerbach, Pierre Bachas, Michael Best, Lucie Gadenne, Francois Gerard, Michael Keen, Asim Khwaja, Wojciech Kopczuk, Camille Landais, Joana Naritomi, Paul Niehaus, Ben Olken, Torsten Persson, Jim Poterba, Dina Pomeranz, Roger Gordon, Joel Slemrod, Johannes Spinnewijn, Stefanie Stantcheva, John Van Reenen, Hans-Joachim Voth, Danny Yagan, Gabriel Zucman, as well as seminar participants at Chicago Harris, Columbia, Harvard, HKS, IIES, MIT, Northwestern, Stanford, Toulouse School of Economics, UC Berkeley, UC San Diego, UCLA, University of Michigan, University of Zurich, Warwick, Yale, and ZEW, for helpful comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Employment Structure and the Rise of the Modern Tax System
Anders Jensen
NBER Working Paper No. 25502
January 2019
JEL No. D31,H11,H21,H24,H26,H71,O15,O22

ABSTRACT

This paper studies how the transition from self-employment to employee-jobs over the long run of development explains growth in income tax capacity. I construct a new database which covers 100 household surveys across countries at different income levels and 140 years of historical data within the US (1870-2010). Using these data, I first establish four new stylized facts: 1) within country, the share of employees increases over the income distribution, and increases at all levels of income as a country develops; 2) the income tax exemption threshold moves down the income distribution as a country develops, tracking employee growth; 3) the employee share above the tax exemption threshold is maximized and remains constantly high; 4) movements in the tax exemption threshold account for the observed variation in tax collection across development. These findings are consistent with a model where a high employee share is a necessary condition for effective taxation and where the rise in income covered by information trails through increases in employee shares drives expansion of the income tax base. To provide a causal estimate of the impact of employee share on the exemption threshold, I study a state-led US development program implemented in the 1950s-60s which shifted up the level of employee share. The identification strategy exploits within-state changes in court-litigation status which generates quasi-experimental variation in the effective implementation date of the program. I find that the exogenous increase in employee share is associated with an expansion of the state income tax base and an increase in state income tax revenue.

Anders Jensen
Harvard Kennedy School
79 JFK Street
Cambridge, MA 02138
and NBER
anders_jensen@hks.harvard.edu

An online appendix is available at http://www.nber.org/data-appendix/w25502

1 Introduction

Tax capacity grows as economies develop. This is true across today's developing countries, and historically within today's advanced countries. In this paper, I show how the transition from self-employment to employee-jobs explains growth in tax capacity. This increase in employee share is a defining characteristic of changes to employment structure over the long run of development. Micro evidence shows that transitions into employee-jobs are associated with improved compliance at the individual level by creating third-party information trails (Kleven et al., 2011), but this evidence has little to say about state tax capacity over the long run. Macro evidence provides correlations between employee shares and tax take (Besley & Persson, 2014; Kleven, Kreiner & Saez, 2016), but this evidence lacks clearly identified empirical channels. To build a bridge between the micro and macro contributions, I propose a research design which combines descriptive evidence and quasi-experimental evidence. In this design, I empirically identify a new channel through which employee share impacts tax capacity along the development path. To implement the design, I construct a novel data-set with micro-data for 100 countries at all levels of development and 140 years within the US (1870-2010).

The novel channel explains decreases in the income tax exemption threshold through increases in employee share that occur gradually further down the income distribution. To motivate the channel, Panel A of Figure 1 shows four countries at increasing levels of development [India, Peru, South Africa, US]. Within each country, it plots employee shares of employment across deciles of the income distribution and the location of the exemption threshold above which earned income becomes liable for taxes. In India, the exemption threshold is located in the top percentiles, the only part where employee share is high. As countries reach higher levels of development, the threshold gradually moves down as the employee-share goes up in deciles further down the income distribution. This close co-movement is also observed within the US over time (Panel B, Fig 1). This paper explains the co-movements as the impact of increases in employee share on the exemption threshold.

I measure the size of the tax base as the share of the income distribution that lies above the exemption threshold. Figure 1 shows that the income tax base in the US is 45 times larger than in India (Panel A). This large variation contrasts with the absence of variation in the income tax rate over development (Figure 5), and suggests a potentially important role for the tax base in explaining the significant variation observed in tax collection across development. In fact, I find that after controlling for the size of the base, there is no remaining association between tax collection and development. In addition to

the channel identified in this paper, this general novel finding suggests that the tax base may be an important channel through which other fundamental drivers of tax capacity impact tax collection.[1]

My employment classification is based on whether the work-type generates derivative information trails relevant for income tax enforcement. The existence of information trails, and the reporting of such information directly to the tax authority by third-parties, has been found to dramatically improve tax compliance in both developed and developing countries, and for both direct taxes (Kleven et al., 2011; Best, 2014) and indirect taxes (Pomeranz, 2015; Naritomi, 2018). Theoretical work (Kleven, Kreiner & Saez, 2016; Gordon & Li, 2009) has provided micro-foundations for the enforcement success of information trails and third-party reporting. I classify as self-employed all work-types that do not generate such derivative information trails for individual income taxation: own account workers, family workers, unincorporated household enterprises, casual wage-laborers, and employers.

I construct a micro database on employment structure and individual income using nationally representative household surveys from 100 countries and within the US (1870-2010). Importantly, the surveys capture all work-types that cover the production boundaries of the U.N. System of National Accounts and all sources of income that can be subject to income tax. The selection of surveys ensures all data is comparable across countries and within countries over time, with detailed dis-aggregation of employment structure over the full income distribution of the economically active workforce.

In a first part of the paper, I use the micro-database to provide four new stylized facts on employment and tax structures along the development path. Stylized fact #1 shows that within country the employee-share of the economically active workforce is increasing through the income distribution, and over development the employee-share is increasing gradually further down the income distribution. Stylized fact #2 shows that the exemption threshold gradually moves down the income distribution, tracking increases in employee-share locally to its left. Stylized fact #3 documents that the exemption threshold is systematically located such that it maximizes the employee-share above the threshold, producing a tax base that is constantly saturated in its employee-composition. Stylized facts #1-#3 are remarkably similar in both trends and levels when I compare US historical profiles to profiles of current countries at similar levels of development. Stylized fact #4 shows a close association between size of tax base and level of income tax collection (relative to GDP), at all levels of development. Residual collection, after controlling for size of base, is uncorrelated with development.

[1] The location of the exemption threshold directly addresses the extensive margin of compliance of moving workers into the tax base. As such, the importance of the base is consistent with a development policy literature which emphasizes the extensive margin of compliance as an important driver of tax capacity (Keen, 2012).

These stylized facts are consistent with employee-share proxying for enforceable income and increases in employee-shares further down the income distribution causing tax collection to grow through expansions of the enforceable tax base. I show that these facts also hold within Mexico over the long-run (1960-2010), and are not confounded by redistributive or sectoral targeting. In a second part of the paper, I complement this descriptive evidence with an identified estimate of the impact of employee-share on tax base and collection. I study the impacts of a previously unexploited historical US state-led development program, the Industrial Development Bonds (IDB) program, implemented in the 1950-1960s. IDB used revenue-neutral debt-funding to expand industry in areas characterized by under-employment among self-employed, acting as a positive level-shifter in employee-share.

The identification strategy exploits the institutional details of IDB implementation. In particular, the unusual funding scheme meant that the state House had to vote in a statute or constitutional provision that exempted IDB from the 'public purpose' provision, whereby a local government may not enter the debt market for an otherwise private purpose. But given the lack of historical precedent, this amendment, and by extension the IDB program, was considered legally uncertain until the state supreme court would uphold its constitutionality in a leading court case. In 40% of IDB states, this legal uncertainty persisted for more than 10 years. I show that the upholding decision effectively acted as a binding constraint on IDB issuance. I use changes in court litigation status as the source of identifying variation. I compare changes in outcomes before and after the upholding event to changes before and after the vote-in event within the same state. My estimation is helped by the significant cross-state variation in timing of upholding, and in the length of the vote-uphold lag. I provide two pieces of evidence to support the identifying assumption that, conditional on vote-in, the timing of the upholding event is uncorrelated with the error term. First, I find that the only determinants of the time-lag are time-invariant, state-specific historical characteristics which previous studies have found to impact the speed of supreme court litigation. Second, I confirm that the regression results are driven by sharp upon-impact changes around the upholding event, with a clear break from a stable pre-trend.

In a regression model that limits the estimation to short-run impacts, I find that the upholding decision led to a significant transition from self-employment into employee-jobs, and an expansion of the state income tax base and collection. I provide several pieces of evidence to support the interpretation that the upholding event impacted the public finance outcomes directly through the employee-share. First, I graphically show that the expansion of the base occurs with a short and precise lag to the increase in the employee-share. Second, I find that the increase in employee-share occurred in deciles

below and locally to the left of the state exemption threshold. Third, I find no effect on collection in a sub-group where there was an impact on employee-share but no impact on the tax base. Fourth, I find no direct impacts of upholding on confounding determinants of tax base and collection.

I shed light on the underlying mechanism in a heterogeneity analysis which exploits the staggered implementation of co-operation agreements between individual state tax administrations and IRS. I show that there is no employee-share impact on tax base expansions in the subset of IDB states that effectively fully relied on the IRS for state tax enforcement purposes by the time of the upholding event. This is consistent with the interpretation that fiscal administrative costs matter for base expansion decisions, and that this cost depends on the enforceability differences between self-employed and employees on the incremental base. This mechanism result increases the external validity of the US identified estimate for the development context, where administrative costs are an important determinant of tax policy (IMF, 2015). I further show that the levels of administrative costs are comparable between US state tax administrations in the IDB period and national tax administrations in developing countries. External validity is also improved by the finding that the upholding event led to a gradual shift leftward of the employee-share profile in the income-distribution and a local decrease in the threshold, leaving the employee-share on the tax base saturated. These identified changes closely replicate the descriptive changes documented in stylized facts #1-#3 along the development path.

In a final section of the paper, I model the reform decision to incrementally lower the exemption threshold. I provide formulae for the optimal location of the exemption threshold under different objective functions, thereby contributing to the optimal tax literature which has hitherto focused on tax rates. The full set of results are consistent with a simple revenue maximizing objective function and an administrative enforcement cost of base expansion that depends negatively on the employee-share on the incremental base. The theory implies that the elasticity of the base with respect to employee-share below the threshold is a meaningful summary measure for the impact of employment structure. I calculate this elasticity using the identified estimate in the US states section, and bring it to the observed cross-country variation in employment structure. The estimated elasticity can account for 29.9% of the observed variation in size of base between low-income and high-income countries.

The following section discusses related literature. Section III describes the micro database and provides new stylized facts. Section IV provides an identified estimate of employee share on tax base expansions. Section V provides a model to rationalize the results and quantify the employee-share channel in the tax-base variation across countries. Section VI discusses implications and concludes.

2 Related literature

This paper is related to the micro and macro studies on information trails as a determinant of individual compliance and state tax capacity. Kleven, Kreiner & Saez (2016) show theoretically that collusive behavior between employees and the employer is hard to sustain when there exist business records, making third-party information reporting by firms a powerful tool of tax enforcement. Gordon & Li (2009) show how information reporting by financial institutions can also improve tax enforcement. These models are supported by empirical studies in both developed and developing countries that show positive tax enforcement effects of information reporting. Kleven et al. (2011) use Danish random audits to show that increases in information coverage associated with employee jobs dramatically improves income tax enforcement. Best (2014) shows that third-party reporting also limits income tax evasion in Pakistan. In Chile, Pomeranz (2015) shows that randomized audit threats have less impact on transactions that are subject to double reporting from both buyers and sellers, indicating its evasion deterrence effect. Bachas & Soto (2018) and Best et al. (2015) show, respectively in Costa Rica and Pakistan, that taxes based on turnover can be a useful alternative to corporate profit taxation because sales are easier to observe than profits. Naritomi (2018) uses a Brazilian reform to find positive enforcement effects of an increased availability of third-party information. Brockmeyer & Hernandez (2018) find positive revenue impacts from a reform in Costa Rica which increased the withholding rate while holding the reporting requirement constant. Carillo, Pomeranz and Singhal (2017) show limits to third-party information effectiveness when taxpayers can adjust on non verifiable margins. Tazhitdinova (2018) finds in the US context that even basic self-reporting requirements are effective at reducing income tax evasion. This paper provides micro evidence on changes in information trails related to employment structure and its impacts on tax structure along the full development path.

The paper also relates to the literature on the determinants of government growth over development (Besley and Persson, 2011; 2014).[2] Demand side determinants include 'Wagner's law' whereby public goods have a income elasticity above one (see e.g. Musgrave, 1966); and, democratization and increased political power of the poor (Acemoglu & Robinson, 2000). This paper is more closely related to supply side studies that show how changes in economic structure impact the capacity to supply tax revenue (Bird & Oldman, 1964; Hinrichs, 1966; Kleven, Kreiner & Saez, 2016). I contribute by providing descriptive and identified evidence on a new tax policy channel, the income tax base, through which economic structure affects the capacity to raise taxes.

[2] See also Cage and Gadenne (2015) for a comprehensive study of revenues trends in developing and developed countries.

The paper is related to the literature on changes in employment structure over development. Current evidence focuses on the cross-country stylized fact that self-employment declines over development (Banerjee & Duflo, 2007; Gollin, 2008 La Porta & Shleifer, 2014) This paper provides new evidence on the decline of self employment over development in larger cross-country and longer time series samples and at dis-aggregated levels over a country's income distribution. Using a previously unexploited US development program, I also contribute with identified evidence on the impact of a place-based program on the transition out of self-employment. My self-employment category is closely related to the informality concepts used in ILO (2009) and La Porta & Shleifer (2014).[3] Thus, this paper complements the macro literature on informality and development, with micro based evidence on changes to informality over the income distribution along the development path. This paper is more generally related to studies of informality's impacts on efficiency and equity of public finances, including Gerard & Gonzaga (2018) and Olken & Singhal (2011). Finally, this paper's methodology relates to studies of macro economic changes using newly constructed micro evidence (Gollin, Lagakos & Waugh, 2013; Bicks, Fuchs-Schundeln & Lagakos, 2018; Feng, Lagakos, & Rauch, 2018).

Finally, this paper relates to long-run studies of tax systems and development in US states. Baicker, Clemens & Singhal (2012) document the growth of state budgets in the postwar period. Akcigit, Grigsby, Nicholas and Stantcheva (2018) investigate the impacts of personal and corporate income taxes on innovation over the 20th century. Serrato & Zidar (2018) study trends in state corporate tax systems and their impacts on state tax collections, while Serrato & Zidar (2016) estimate the incidence of changes to state corporate tax rates and apportionment rules. Dincecco & Troiano (2015) study the revenue impacts of the introduction of the state individual income tax, and Troiano (2017) focuses on the revenue and audit impacts of cooperation agreements between state administrations and the IRS. This paper provides novel evidence on the determinants of state income tax base expansions and its revenue impacts. Place-based economic development programs represent a significant share of current US state and local government expenditure (Kline & Moretti, 2014). This paper contributes by providing evidence on a previously unstudied historical place-based, state-led development program. The IDB program sought to attract capital to specifically designated rural areas for the limited purpose of industrial expansion, using revenue-neutral funding schemes. Relative to previously studied US development programs, including the Tennessee Valley Authority (Kline & Moretti, 2013), IDB was distinct in its channel of impact, its narrowness of scope, and its strong geographical targeting.

[3]My classification of job-types as self-employed is also empirically close to the 'hard-to-tax' groups, which was first defined by Musgrave (1981).

3 Descriptive evidence on employment structure and tax structure along the development path

I first describe the survey data underlying the cross-country and US historical analysis. Next, I outline the definitions and the methodology used to construct the key variables. I then provide novel stylized facts on employment and tax structure along the development path.

3.1 Data

The micro database is constructed from nationally representative household surveys. The key advantage of using household surveys, as opposed to firm surveys, is that our measure of employment is not restricted to activities which generate a wage, but also include self-employed as well as family workers. Relative to administrative data, the household surveys present the advantage of observing income along the full income distribution, independently of the income tax exemption threshold.

The database is limited to nationally representative household surveys which cover all types of employment and provide continuous measures of total income. I was able to collect such data in 100 countries with a population of at least one million. The different parts of the worldwide income distribution are almost equally represented, with 20 % from low-income countries; 28% from lower-middle income countries; 21% from upper-middle income countries; and, 31% from high-income countries.[4] For 44 countries, I was able to draw on harmonized data-sets, comprised of Luxembourg Income Study (LIS, 36 countries), Economic Research Forum Micro Data Catalog (ERF; 6 countries), and International Public-Use Micro-data Project (IPUMS, 2 countries). For the remaining 56 countries, which are exclusively in the low-income and middle-income groups, I sourced the data directly from the country's national statistical agency. The database contains 93 living conditions surveys, 5 labor force surveys, and 2 censuses. When multiple years were available, I choose the year closest to 2014; exact years, data-sources, and sample sizes for all households are given in Appendix A.1.

I construct the within country dimension of the database by combining new and previously used micro sources from the US. I collect data between 1950 and 2010 from Census micro-data extracted via IPUMS USA. Before 1950, the Census did not record work type and continuous measures of in-

[4] I calculate the income classification of a country in the year for which I have survey data using the historical World Bank Classification data.

come at the individual level. I use the 1935-36 Study of Consumer Purchases, which was considered the precursor to the Census methodology of data on income at the individual level. 300,000 households were interviewed based on sampling units chosen to represent the "demographic, regional, and economic characteristics of the United States" (ICPSR, 2009). Importantly, both the work type and income categories in the 1935 survey are consistent with the later Census-based definitions. All national surveys carried out in the late 19th century and the end of the 1920s focused on sampling the work and living conditions of employed wage-earners. To construct a historical pre-1900 profile of employment structure, I use data resulting from a collective effort between Williamson & Lindert (2016) and IPUMS USA. Williamson & Lindert use local tax assessments and occupational directories for 'registered occupations' and local censuses for 'unregistered occupations'. Labor force counts using the 1 percent US Census sample were provided specifically for the data-project by IPUMS USA. I use the Williamson & Lindert computations of gross earned income, which include wage income, farm income and non-farm business income. However, unlike the surveys from 1935 onward which contain harmonized employee and self-employed variables, the 1870 data required building types of employment categories. I use a text search algorithm which exploits the highly detailed work titles from the enumerator instructions to the 1870 Census in order to construct self-employed and employee categories. Further details on the US historical data are in Appendix A.1.

3.2 Definitions and methodology

The aim of this section is to construct a distributional profile of employment structure in the individual gross income distribution, and relate distributional employment changes to changes in the size of the income tax base across countries and over time. In all surveys, I follow the same procedure to construct the income distribution. I limit the sample to the economically active population, following the definition of employment from U.N. System of National Accounts (SNA). This definition has been used in recent cross-country studies of hours worked (Bicks et al., 2018) and unemployment (Feng et al., 2018). To calculate individual gross income, I use a broad set of categories that are included in the determination of individual income tax liability in countries around the world: employment income, self-employment income, capital income, and miscellaneous income.[5] I exclude two important sources

[5]Not all countries tax all these sources of income on the individual income tax. I nevertheless maintain this broad measure to ensure comparability of the income distribution across all countries. As an alternative, I can construct the income

of income: social transfers, and non-monetary income. Social transfers are excluded because they are not subject to income tax. In contrast, certain non-monetary payments in a work-setting can be subjected to income tax, but I exclude this income source simply because I cannot consistently measure it across all surveys.[6] In 7 countries, I cannot calculate gross individual income with precision. In those cases, I proxy for gross income using expenditure.[7]

I code employment as self-employed versus employee on the basis of whether the work-type generates derivative information trails that can be used for income tax enforcement. Third-party coverage of employee earnings through employer reporting and withholding has been found in both developed and developing countries to significantly improve income tax enforcement (Kleven et al., 2011; Best, 2014).[8] In developing countries where third-party coverage is limited, such information includes paper trails generated by contractual arrangements such as written labor contracts. This information-trails definition of employees is conceptually consistent with the contractual definition of employees in Banerjee & Newman (1993). Based on this definition, I classify as self-employed: own account workers, casual wage-laborers, contributing family workers, employers, and unincorporated household enterprises. This classification is very similar to ILO's classification of informal labor (2002), with the exception of domestic workers in private households, which I classify as self-employed and ILO as employee. Note that both classifications are related to 'ex ante' criteria with respect to the tax system and which vary with development. They include economic risk of the job holder, authority within the establishment, organizational scale, and complexity of production. This is conceptually distinct from the classification of the underground economy (LaPorta & Shleifer, 2008), which is based on 'ex post' tax registration and/or evasion status. As such, the observed variation in employment structure may be related to structural, non-tax changes along the development path.[9]

My classification is internationally comparable in surveys across countries and within the US over

distribution based on the (more narrow) earned income concept, which excludes capital income and miscellaneous income. Results based on this alternative income definition are qualitatively and quantitatively similar.

[6]The only exception is for respondents that are economically active, but report no monetary income - per example, family workers in agriculture that fully consume their own output. For this limited sample of respondents, in a limited number of countries, I impute an income based on answers to questions about the market-value of the own-consumed output. For more details on the importance of non-monetary income for the main results, please see Appendix A.1.

[7]This is mainly due to the inability to measure self employment income in the agricultural sector. See Appendix A.1 for further details. Note that in these countries, I do not attempt to convert expenditure into gross income for the purpose of locating the income tax exemption threshold in the income distribution.

[8]Kleven, Kreiner, and Saez (2016) develop a three-tiered agency model, with a government, a firm, and workers, to micro-found the success of third-party reporting.

[9]Naturally, the ex ante criteria are plausibly endogenous to the tax system. The ex ante versus ex post distinction serves the point that my employment classification is not mechanically and directly derived from a tax-based classification, unlike the underground economy classification.

time. In all surveys I can clearly distinguish employees from employers and from both family and non-family workers in household enterprises. As such, my classification is not confounded by a distinction between firm of size 1 versus larger firms. In countries where seasonal wage-work is arguably most important, I can distinguish such workers from contract-based employees. Since casual wage-work is most prevalent at lower levels of development, any mis-classification would lead me to under-estimate the true transition out of self-employment along the development path.[10]

In all surveys, I code employment type based on the primary job in which the respondent spent most hours working during a reference period. Many individuals have many jobs at the same time (Banerjee & Duflo 2007), but this will affect the representativeness of my estimates only to the extent that these jobs fall in different categories in my classification. An individual who contributes on the family farm while being an own-account worker within the same reference period would be classified as 'self-employed' in either job. In surveys where the reference period is not yearly, there may be bias in the measure of employment structure if the employment type in the reference period is not representative of the entire year. This is likely to be important in developing countries, where there is strong seasonality in job type. This introduces bias to the extent that the jobs at different periods of the year fall in different employment structure categories. An individual that is a casual wage laborer during the harvest season and is an own-account worker for the remainder of the year would be classified as self-employed in any period of the year.

The income tax base is defined as the share of the economically active population whose gross income lies above the income tax exemption threshold. I define the exemption threshold as the minimum level of gross income above which individuals become liable to pay some strictly positive income tax.[11] I consider a non-married filer without any dependants. Without any further assumptions on filing behavior, the exemption threshold for this filer equates to the standard deduction (or allowance). This relief is most often universally and automatically available to all taxpayers as a fixed amount, and is unrelated to any expenditures incurred by the taxpayer. I abstract from other features of the tax system that vary across countries, such as the existence of individual credits, which can allow the

[10] I can always distinguish between working for someone else for pay versus for in-kind payment, and I exclude the latter from the employee category. As such, casual wage-laborers that receive in-kind payment are systematically classified as self-employed. This leaves the group of casual workers that are not paid in-kind as the group that I potentially mis-classify as employee, whenever the survey answers do not provide sufficient precision about the nature of the employee-work. I discuss measurement-induced bias in Appendix Section A.6.

[11] A very limited number of countries, many of which are located in Eastern Europe, have recently moved to the broadest possible income tax base, where the exemption threshold is 0. The income tax system in all the countries in my sample feature a strictly positive exemption threshold in the year of the survey data. A few countries feature a 0-th bracket, where the filer still has to file an income tax return but the tax rate is 0; in this case, I define as the exemption threshold the value of income above the 1st bracket where there is a strictly positive rate.

filer to further reduce her final tax liability. I do this because it allows me to calculate the exemption threshold in a simple way without having to make any behavioral assumptions about filing and credit take-up (Benzarti, 2018). In addition, the simplicity of this threshold and its common existence means that I can create a harmonized measure of the size of the income tax base that is internationally comparable across space and time. Finally, since the standard deduction constitutes a lower bound on the true exemption threshold, my estimates will potentially understate the narrowness of the income tax base in developing countries.[12] To increase international comparability, I code values of the exemption threshold in all cross-country surveys from a unique source, namely the country tax reports of the IBFD Tax Research Platform.[13] I code the value of the threshold in the year of the survey. For the US series, I use the historical IRS series on the personal exemptions between 1913 and 2012.

I measure the distributional employment profiles by calculating the share of employees in the economically active workforce in deciles of a country's gross income distribution, using individual weights. I locate the exemption threshold in every gross individual income distribution, and calculate the income tax base as sum of percentiles above the exemption threshold. My employment classification is consistent in the aggregate with the U.N. SNA, which means that all activities identified within the SNA production boundary are captured in the employment estimates in every income decile in every country. As such, the size of the income tax base can be interpreted as the share of the economically active population that is subject to income tax. I use expenditure-side real GDP at chained PPPs in 2011 US$ from the same year as the country survey data.

The selection of surveys and the construction of variables has been done to ensure that comparability across countries and over time is as high as possible. All of my surveys are nationally representative and have detailed and consistent work-type classifications and continuous income categories which allow me to measure the employee-share of the economically active workforce over deciles of the gross income distribution. My measure of the income tax base is based on the standard deduction, which universally features in all tax systems in my sample and which can be calculated without any additional assumptions about tax filing behavior.

[12] In the Appendix (Section A.6), I argue that this choice of measurement leads me to understate the importance of the tax base as a determinant of tax collection.

[13] IBFD is the primary data-source for the country-specific tax summaries produced by several international consultancies (including Deloitte, KPMG, and PwC).

3.3 Results

In this sub-section, I present new stylized facts on changes to employment structure and tax structure along the development path. To summarize findings in the cross-country section, I partition the sample of 100 surveys into ten development groups of 10 countries. In each group and decile of the gross income distribution, I construct the unweighted average of employee-shares in all countries in that group.[14] I complement the cross-country evidence with time-series evidence from the US (1870-2010), with a focus on understanding how the US looked like historically when it was less-developed. The time-series spans almost as large an income range as the cross-country data: in 1870, the earliest year of historical data, the US had a similar level of per capita income to India in the survey year.[15]

Stylized fact #1: Within country employee share increases over the income distribution, and at all levels of income as a country develops

Figure 2 displays the employee distributional profiles in the 10 cross-country development groups. The first stylized fact documents that within each country, the employee-share increases with income, and that across development, the upward-sloping employee-profile moves leftward in the gross income distribution. In other words, in each successive development group, the increase in employee share is concentrated in deciles gradually further down the income distribution. In the poorest countries in the sample, such as Democratic Republic of the Congo, Niger, and Malawi, the employee-share is effectively nil apart from a steep increase in the top percentiles of the income distribution. During the transition from low to lower-middle income countries, such as India, Pakistan, Zambia and Bolivia, the growth in employee share is concentrated in the top three deciles. As the transition to upper-middle income countries takes place, to countries including China, Brazil, South Africa, and Turkey, the growth in employee-share occurs mainly in the four middle deciles (from the 4th to the 7th decile). Finally, when going from upper-middle to high-income countries including the US, UK, Germany, and Japan, the remaining growth in employee share is concentrated in the bottom three deciles.

In Figure 3, I plot the US profiles between 1870 and 2000. In 1870, on the eve of its second industrial revolution, the employee-share in the US was concentrated in the top decile. Over time, the employee-profile gradually moves leftward in the income distribution. For every US historical profile, I construct a 'synthetic' cross-country profile, based on the average of profiles of countries with per capita income

[14] As an alternative, I can construct weighted profiles, using the country's population share in total population of the development group as the country-specific weight. The results based on this weighting are very similar.

[15] Shortly after the year of the survey, 2004, India graduated to the lower-middle income group. As such, the US historical data spans development levels between the lower-middle income group and the high-income group.

within 10% of the US income in the historical year.[16] I insert these synthetic profiles in Figure 3. Interestingly, the patterns look very similar, both qualitatively and quantitatively, when comparing a historical US profile to a cross-country profile at the same level of per capita income.

Stylized fact #2: Over development, the tax exemption threshold moves down the income distribution in co-movement with increases in employee share

The second stylized fact documents a strong co-movement over development between increases in employee share occurring gradually further down the income-distribution, and decreases in the exemption threshold. This is illustrated using four countries from the cross-section and four points in time within the US in Figure 1. At lowest levels of development, the exemption threshold is systematically located in the top percentiles of the income distribution, generating a very narrow income tax base between 1% and 5% of the workforce. Over development, we observe a close co-movement between decreases in the location of the exemption threshold and growth in employee share locally to its left. Given the gradual shifts in employee-share profile (fact #1), this co-movement generates a gradual increase in the size of the income tax base over development, as documented in Panel B of Figure 4. The income tax base is very narrow in low-income countries, taxing only 1-10% of the economically active population. The base expands in middle-income countries to 30-50%, and reaches 90-95% in high income countries. Again, the patterns are similar when comparing the expansion of the income tax base across countries to the historical expansion within the US over time.

Stylized fact #3: Location of tax exemption threshold systematically maximizes the employee-share on the income tax base

When comparing the poorest to the richest countries in the cross-section, panel B of Figure 4 implies that the income tax base expands by over 4500% over the course of development. The third stylized fact shows that despite this large variation in size, the composition of income tax base remains constantly maximized in employee-share over development. To show this, I calculate in all surveys the percentile location of the exemption threshold which would maximize the employee-share on the base, and the associated employee-share. I then calculate the employee-share on the real tax base, as the employee-share above the actual exemption threshold. Finally, I divide the actual employee-share by the maximized share, and plot the ratio, expressed as a percent, against GDP per capita in Panel C of Figure 4. This ratio would be equal to 100 percent if a country locates its exemption threshold in the ex-

[16] Per example, the US in 1870 had similar per capita income to India and Pakistan in their respective survey-years. The US in 1935 was similar to Indonesia and Peru, and the US in 1950 was similar to Mexico and Uruguay.

act percentile which maximizes the employee-share on the income tax base. I find that on average, the actual employee composition of the tax base achieves 95% of the maximized employee-composition. What is more, this ratio is constant across all levels of development, both across countries and within the US over time.[17] This fact suggests a constant policy-rule regarding the employee-composition on the tax base, despite large variation in both overall employee-composition (fact #1) and size of the base (fact #2) over development.

Stylized fact #4: Changes in income tax base accounts for variation in income tax collection across development

The final stylized fact shows that variation in the size of the income tax base can account for the bulk of the observed variation in income tax collection (relative to GDP) across development levels. In Panel A of Figure 5, I plot income tax collection, size of the income tax base, and the (top) marginal statutory income tax rate, against GDP per capita, in the cross-country section. Since I am introducing a novel tax policy instrument, I benchmark it against the marginal rate which has hitherto been the main proxy for government statutory tax policy in studies explaining variation in tax collection across development (Gordon & Li, 2009; Besley & Persson, 2014). Panel A reveals a close association between the size of the base and the level of income tax collection: both variables slowly increase between low-income and lower-middle income countries, but witness strong growth during the transition from upper-middle to high-income countries. In contrast, there is essentially no variation in the top marginal tax rate across development. In Panel B, I directly study the association between each of the two statutory policy instruments and tax collection, while controlling for the other instrument. This produces a strongly positive and linear relationship between the size of the base and tax collection, after controlling for the rate. In countries with the most narrow tax base (less than 10% percent), tax collection is between 1% and 3% of GDP. Countries with a relatively more expanded base (between 40 and 60%) collect on average 8% of of GDP in income taxes. Finally, countries with the broadest base (above 90%) have a collection level between 13% and 18%. On the other hand, once I control for the size of the base, I find no direct association between the marginal tax rate and the level of collection. Finally, I show that, once I remove the association between the base and the level of (income tax/GDP), there is no remaining association between tax collection and development. In Panel C, I directly plot

[17] The extent to which this ratio can deviate from 100% depends on the steepness of employee-share gradient in gross income. Indeed, if the employee-share was uniform across the income distribution, then any arbitrary location of the threshold would generate a ratio of 100%. As such, this test is informative of policy-rules in settings except for high-income countries, where the employee-share is high and uniform across the income distribution (Figure 2 and 3).

the residual of (income tax/GDP), after a regression on the income tax base, against GDP per capita. The residual displays a significant amount of variance, suggesting that there is meaningful variation in collection left to explain. But the association between collection and development is now flat. As a novel stylized fact, this result suggests that drivers of tax capacity over development, including the employment structure channel studied in this paper, operate in large part through the tax base.

3.4 Robustness

In this sub-section, I provide robustness checks for the stylized facts. I first show that the stylized facts also hold within Mexico over 50 years between 1960 and 2010. This time-series covers the lower levels of per capita income in the cross-country data. In the Mexican time-series, I uncover the same stylized facts: the employee-share profile is upward-sloping and gradually moves leftward in the gross income distribution; the exemption threshold gradually moves down the distribution and expands the size of the income tax base; the employee-composition on the tax base is constantly maximized.[18]

In a second robustness check, I consider the possibility that the exemption threshold targets characteristics which are correlated with employment structure. I code the value of the national poverty line and the minimum wage for all countries in my sample in the years where I have survey data.[19] These are thresholds set by the government and serve as the basis for most redistributive and anti-poverty efforts. I use the IBFD summaries to show that in less than 5% of countries in the sample is the exemption threshold explicitly defined in the tax code as either equal to, or a multiplicative of, the national poverty line or the minimum wage. I further show in the appendix that there does not either appear to be any implicit relationship between these two sets of thresholds, as countries with similar levels of income and size of tax base display enormous variation in the ratio of the exemption threshold to the poverty line or the minimum wage. This suggests that the exemption threshold is not used to target social assistance and anti-poverty efforts in the income distribution.[20]

In a third robustness check, I consider whether the exemption threshold is targeting sectoral, rather than employment, structure. An earlier literature has focused on agriculture as a hard-to-tax sector (Musgrave, 1981). This could suggest that the location of the threshold is set to exempt agriculture,

[18] Details on the construction of the Mexican time-series can be found in Section A.2.

[19] I collect the poverty lines that are set by the national government, rather than the international poverty lines defined by the World Bank. Details on the data-collection can be found in Section A.3.

[20] Of course, the exemption threshold redistributes through the tax system directly, as it is defined based on an individual's income. In Section 4, I provide an estimate of the impact of employee-share on the exemption threshold, in a setting where income was constant.

rather than tax employees. I find limited evidence to support this hypothesis. First, only 12% of countries in both low and middle income groups, and 5% of high-income countries, explicitly list agriculture as a source of exempt income.[21] Second, I code the sectoral composition of the workforce in all countries, and investigate the distributional profiles.[22] I show in the Appendix that in low-income and lower-middle income countries, where the agricultural-share is steeply downward-sloping in the income distribution, virtually all of the agricultural workforce is self-employed. At the same time, between upper-middle and high-income countries, where agriculture is much less quantitatively important, there continues to be an important transition from self-employment to employees, which occurs outside of the agricultural sector, and which is associated with continuous decreases in the threshold. I also show that the exemption threshold does not appear to target "easy-to-tax" sectors, such as manufacturing or the public sector. In the Appendix, I show that the income-profiles of the sectoral composition of employees does not support this hypothesis. Neither manufacturing nor the public sector profiles are upward-sloping in the income-distribution and leftward-shifting over development, in a way that could explain movements in the exemption threshold.

In a final robustness check, I use a regression setting to investigate the robustness of the employee-income slope to controlling for different characteristics. I study the impacts of sector (4 industry dummies), geography (urban dummy), and education (three levels of educational completion dummies), all measured at the individual respondent level. If controlling for one characteristic fully eliminates the employee-income slope, this could suggest that the exemption threshold is targeting this particular characteristic, rather than employment structure. At the same time, a partial reduction in the magnitude of the employee-income gradient upon inclusion of a control would be informative about the characteristics which quantitatively contribute to the slope. I find that the basic income-slope in employee-likelihood is not confounded by any one characteristic (education, sector, geography) that the government could implicitly or explicitly target. However, I find that jointly, these three characteristics can account for the bulk of difference in employee-income slopes across levels of development.

Finally, note that unobservable increases over development in the capacity to detect under-reporting of self-employment earnings would appear as a decrease in the employee share. Underlying growth in enforcement capacity thus works against my finding of gradual increases in employee share along the development path. I discuss other potential biases in Appendix Section A.6.

[21] This includes India and Pakistan. I code this in the IBFD country reports. I only count instances where agricultural income is entirely exempt. There are cases where self-employed operating in agriculture may deduct specific costs to determine their tax liability, but I do not include such instances. More details are found in Appendix Section A.4.

[22] I discuss the sector classification in Appendix Section A.4.

4 Direct estimate of impact of employment structure on the tax base and tax collection

In this section, I provide direct evidence of employment structure's impact on tax structure, using the setting of US states over time. I first provide background, data, and program details. I then discuss the identification strategy. I present graphical evidence and regression-based results, before providing results on the underlying mechanism and robustness checks.

4.1 Background and program details

Background and data

US states represent a potentially compelling setting to study development of tax systems. Each state chooses which sources of income to tax, including individual income taxes, and is responsible for the associated collection and enforcement costs. State tax administrations make investments in enforcement capacity, including in filing, withholding, and auditing processes, that are similar in nature to investments made by national administrations across countries (Snavely, 1988). As a measure of enforcement 'capacity', the cost of collecting state taxes in the average state tax administration was 4.3% of taxes collected in the early 1960s.[23] This is comparable to the average cost of collection of national tax administrations in low-income countries today, which is 3.8% (Jensen & Lagakos, 2019). Historically, the growth in state tax-to-GDP ratios has been largeley driven by an increase in personal income taxes, as seen in Panel A of Figure 6.[24] The rise of the modern income tax system in individual states matches the key tax capacity stylized fact (Besley & Persson, 2014; Kleven et al., 2016). In parallel to the rise of the modern income tax, individual states also witnessed large changes in employment structure over time. Panel B of Figure 6 shows the employee-share of total employment along the income distribution of the average state over time between 1950 and 1980. The employee-share is increasing in the income distribution, and over time the employee share increases in all income deciles. This is consistent with stylized fact #1. Panel B shows the state exemption threshold gradually moves down

[23] The cost is measured as the sum of the capital outlays and the payroll of the state's financial administration. The 1962 ediction of the *Book of the States* is the earliest year where I can construct this cost measure. More details are found in Appendix Section B.1.

[24] The change in state tax structure is discussed in detail in Wallis (2000).

the income distribution in co-movement with increases in employee share locally to its left, consistent with stylized fact #2. Finally, the employee-share above the threshold remains constant and high, consistent with stylized fact #3. Taken together, these facts increase the external validity of the US-states direct evidence for the cross-country setting.

I combine data from several sources to construct variables that are as comparable as possible to the cross-country section. I calculate employment structure in the active workforce using decennial Census data between 1940 and 2010. The 'class of worker' category in the Census data allows me to classify respondents as either self-employed or employee in a manner that is consistent over time within US states and with the cross-country classification. The advantage of the Census data is that it provides a measure of employee-share of the workforce. The disadvantage is that it is not yearly. I interpolate the numerator and denominator between Census years using a natural cubic spline (Herriot & Reinsch, 1973). As an alternative, I use the employee-share of income, which is measured continuously on a yearly basis by BEA, and find similar results. I rely on the US historical state tax calculator 1900-2007, constructed by Bakija (2009), to construct measures of the state-year income tax structure. In particular, I calculate the tax exemption-threshold for an individual single earner who claims the standard deduction. This definition of the exemption threshold matches the definition used in the cross-country setting. One important selection is that I limit the sample to all years before 1980. That is because from the 1980's onward, states began to index their exemption threshold to inflation. In the pre-1980 period, I can directly assess if changes to employment structure lead to active government decisions to change the nominal value and location of the threshold. To construct tax collection levels, I use state government finance publications between 1929 and 2010 published by the US Census. I express tax collection as a fraction of total state personal income, from the historical BEA series.[25]

Program details

To establish a direction of causality from employment structure to the income tax base, I exploit variation in the effective implementation date of the Industrial Development Bonds (IDB) program. Through the IDB program, sub-state government units (most often counties and boroughs) issued bonds to finance the acquisition or construction of facilities and equipment for lease to private firms (Cobb, 1993). Importantly, IDB issuances were revenue bonds, which are secured exclusively by the

[25]The construction of the variables is discussed in detail in Appendix Section B.1.

revenues of the project. This is in constrast to general obligation bonds, which are secured by the credit of the issuer - in this case, the local government (U.S. Department of Commerce, 1978). This distinction implies that there is not a direct relationship between the issuance under the IDB program and increased tax revenue due to a need to solidify the local government's funding capacity.

According to an early Federal definition, the IDB program was intended to attract industrial plants to rural communities characterized by under-employment among self-employed (Area Redevelopment Administration Commission (ARA), 1962), where the IDB would act as a positive level-shifter in employee-share. The dominant justification for government intervention was that inadequate local credit for industrial firms constituted the barrier to expansion in rural communities (ARA, 1968).

Financing of the IDB program was directly incompatible with the 'public purpose' provision, whereby government debt may only be issued for public purpose. Implementation therefore required the state House to vote in a legal statute which exempted IDB from the public purpose provision. In the cases where the public purpose provision was directly written in the state Constitution, the vote-in led to a constitutional amendment. But there was no legal historical precedent to such development program. The voted statute or amendment, and by extension any bond issuance, would therefore remain legally uncertain until the highest state court would litigate to uphold its constitutionality through a specific court case (Abbey, 1965; Cobb, 1993; Pinsky, 1972; Rollinson, 1976). The court case would most often be triggered by any of the administrative steps that the local government was required to fulfill, such as the set-up of local administrative units devoted to handling the IDB issuance process. The fact that IDB were issued as revenue bonds, rather than general obligation, was often the basis of the argument for not violating the 'credit for public purpose' doctrine. In the case of Wayland v. Snapp, the Arkansas Supreme Court "(...) chose to uphold the issuance of the revenue bonds by invoking the doctrine that revenue bonds do not violate a credit clause because they are retired through lease revenues of the project, not out of tax funds" (Yale Law Journal, 1961).

I code the year of vote-in and the year of upholding from both legal reviews and government publications. I find very little deviation between sources in the major cases cited as the leading case which upholds the IDB program.[26] The time-lag between the vote-in and the upholding events is substantial. The lag has mean of 6.67 years and standard deviation of 6.77. In 40 percent of cases, the time-lag exceeds 10 years. I digitize archived Moody's state financial records on issuance of IDB debt, and show below that upholding acts as a necessary condition for IDB issuance. I limit the estimation sample to

[26] In those few cases where there exists a deviation between sources, I pick the earliest case as the year of the upholding. Details on the sources and timing of vote-in and upholding events are in Appendix Section B.2.

the 28 states that have upheld the IDB program before the implementation of a major IRS reform in the early 1970s. This reform significantly widened the scope of projects approved under municipal bond projects, and shifted the nature of funding from rural industrialization towards public-goods projects in infrastructure and environmental conservation. The narrower scope of the program in the pre-1971 states arguably eases the interpretation of the main effects and limits the set of potential confounders.

4.2 Identification strategy

The legal uncertainty surrounding the IDB program suggests that the vote-in event would be a fuzzy treatment. Instead, I center the program impact around the court upholding event. I limit the sample to upholding states, since states that never implement the program are likely to be fundamentally different from implementing states in ways that are hard to control for. To account for the possibility that the decision to implement the program may be endogenous to changes in the political and economic environment, I include a dummy variable for the intermediary period between the vote-in event and the upholding event. This control captures any state-time varying observable and unobservable shocks to the political and economic environments which occur around the time of the endogenous policy-decision to to vote. This includes deciding to vote-in the program in response to shocks to employment structure, and deciding to vote-in based on changes in 'need' or 'taste' for revenue.

This discussion motivates the estimation of the following regression

$$y_{st} = \beta + \alpha \mathbf{1}\left(\text{Vote-in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \lambda \mathbf{X}_{st} + \mu_s + \gamma_t + \varepsilon_{st} \tag{1}$$

where s denotes state, t denotes time, $\mathbf{1}\left(\text{Vote-in}\right)_{st}$ is a dummy taking a value of 1 when the state has voted in the IDB program but it has not been upheld in the state court system, and $\mathbf{1}\left(\text{Upheld}\right)_{st}$ is a dummy taking on a value of 1 when the state supreme court court has upheld the legality of the constitutional amendment. Exogenous control variables in $\mathbf{X}_{st}$ include dummies for election years, and the first-stage variables used in Besley, Persson, and Sturm (2010) to instrument for political competition and tax and expenditure policies.[27] All standard errors are clustered at the state level to allow

[27] The first stage instruments measure the share of the state population which are subject to either a literacy test or a poll tax (or both), and which were abolished after the 1965 Voting Rights Act. In the Appendix (Table B.3), I show that the main estimates are robust to including additional (potentially endogenous) control variables which vary at the state-year level: a

for correlation over time within a state. The estimation in 1 assesses program impacts by comparing changes in outcomes before and after the upholding event, relative to changes before and after the vote-in event within the same state. Variation to estimate θ comes both from the within-state timing of the upholding event relative to the vote-in event, and from the cross-state variation in timing of the upholding event and in length between the vote-in and the upholding events.

To interpret the coefficient on $\mathbf{1}\left(\text{Upheld}\right)_{st}$ as the causal impact of the IDB program requires the identifying assumption that, conditional on the vote-in and the included controls, the timing of the court upholding decision is uncorrelated with the outcome variable. I provide two pieces of supporting evidence. First, in the following sub-section, I show that the regression effects are driven by sharp upon-impact changes following the upholding event. Outcomes are trending in a stable manner throughout the intermediary period between vote-in and upholding. Groups with differing length of lag between vote-in and upholding are on parallel trends prior to the upholding event, suggesting that groups with longer lag constitute a valid counterfactual for groups with shorter lag.

The second supporting evidence is that the only significant predictors of the time lag are state-specific, time-invariant historical characteristics. Table 1 reports the results from non-parametric Cox proportional hazards models. These models use state time-varying and time-invariant regressors to predict the conditional probability of the upholding event occurring, conditional on the vote in event having occurred. A civil law dummy significantly predicts a higher conditional probability of upholding. This variable is drawn from Berkowitz & Clay (2005), and codes a state with civil law origins if, by the time of American acquisition, its colonizers had a civil law legal system (as opposed to a common law system).[28] The faster time to uphold associated with civil law states is consistent with studies across US states (Berkowitz & Clay, 2005) and across countries (La Porta, Lopez-de-Silanes & Shleifer, 2008). [29] States that experienced defaults in the 19th century on debt issued for railroad-funding under a private-public scheme similar to the IDB take longer to uphold, conditional on voting in. This finding is consistent with a legal literature which argues that the 'public purpose' doctrine was more

dummy for whether the state has right to work laws, a dummy for the existence of a state corporate income tax, the firm-size coverage of state unemployment insurance, and log per capita income.

[28]Ten of the continental American states were settled by France, Mexico or Spain and had civil law legal systems by the time of the American Revolution. These ten states are: Alabama, Arizona, Arkansas, California, Florida, Louisiana, Mississippi, Missouri, New Mexico, Texas. The 38 other had a common law system or were unsettled. Note that an additional five states - Illinois, Indiana, Michigan, Ohio, Wisconsin - were also originally settled by a civil law country, but were acquired by Great Britain prior to the American Revolution.

[29]One interpretation of this result, drawing on Berkowitz & Clay (2005), is that civil law produces a Constitution with more statutory components, rather than framework provisions, and that the existence of statutory laws created more frequent demand for constitutional change among affected groups as the political and economic climates change over time. This explanation is consistent with the difference in IDB-litigation procedures observed across states: civil law origins states were more likely to vote in statutes, as opposed to Acts, which was likely to being revised more quickly.

entrenched in the states that had witnessed the adverse outcomes of IDB-type funding schemes in the past. The historical default experience is coded from Sbragia (1996). Importantly, none of the political variables or economic variables are significant predictors of the timing of the upholding decision, once these two historical state-specific variables are included. I consider a broad set of potential political factors which vary at the state-level, including: a dummy for whether state supreme court judges are appointed; a dummy for election year, a dummy for whether a neighboring state has upheld; the size of the population; a dummy for whether voting restrictions are in place; and, a dummy for whether a state has right to work laws. I also consider a broad set of state-varying economic factors, including: the labor force share of counties in a state that were considered to be eligible for IDB funding by the Federal administration; the growth rate of the labor force; the growth rate of manufacturing employment; the employee-share of employment; log of income per capita; the size of the state income tax base; the ratio of income tax to GDP; and, the ratio of total tax to GDP. As state-specific time-invariant variables, the civil law and historical default determinants of the time-lag should be absorbed in the state fixed effect in 1. Notwithstanding, I show in the Appendix (Table B.3) that the main results are robust to allowing both groups of states to be on non-parametric time-paths over the sample period.

I center the sample in every state around the year of the upholding event, in a time-window which ranges from 5 years prior to the vote-in event to 5 years following the upholding event. Consequently, θ will only pick up any 'short-run' impacts of the upholding decision, as suggested by the sharp upon-impact changes in the graphical evidence below. The main disadvantage of this sample decision is that my estimates cannot contribute to a meaningful, long-term evaluation of the overall economic impacts of IDB. The main advantage is to help increase the likelihood of capturing the isolated, direct impact of IDB on employment and tax structure. Indeed, any long-run estimate of θ is likely to be composed both of the direct impact of the court upholding decision and endogenous policy-responses to this court decision (Besley and Case, 1994). The short-run time-window increases the likelihood of excluding indirect economic impacts in θ, including sectoral spill-over and agglomeration effects, which have been found in long-run studies of place-based development programs (Kline and Moretti, 2013; 2014).

4.3 Graphical results

Before turning to the regression results, I provide graphical event-study evidence. I study outcomes in two groups: those with a time-lag from vote-in to upholding that is between 5 and 10 years, and those with a time-lag that is larger than 10 years.[30] To replicate the estimation strategy described above, I show outcomes in each group over five years prior to the vote-in, the full intermediary period, and a limited period after the upholding event. In each group, I take the raw means of outcome variables, and index it to equal 1 in the year of the upholding event. This event-study allows me to graphically illustrate the estimation strategy in 1, and to flexibly visualize the patterns of outcomes relative to both the upholding and the vote-in events directly in the primary data.

The results are displayed in Figure 8. Panel A studies the impacts on issuance of IDB debt. I code the full set of IDB issuances in every state and year, using the historical series of Moody Municipal and Government Manual, and plot the cumulative volume of issuance. Panel A displays a very large increase in issuance in the two first years immediately following the upholding event. In contrast, there is no change in trend leading up to the upholding event. This suggests that the upholding event was not fully anticipated by investors. I do not find either any significant increase in issuance following the vote-in event. This suggests that legal uncertainty following vote-in was an effective constraint on the actual implementation of the program. Finally, the trends are parallel between the long time-lag and the short time-lag groups around the year of the vote-in for the short-lag groups. This suggests that states where the time-lag is longer constitute a potentially valid counterfactual for the change in outcome of the states with shorter time-lag.

In Panel B, I consider the impacts on employment structure. I use the employee-share of earned income rather than employment.[31] I do this because it is available on a yearly continuous basis before and after both the vote in and upholding events. In the regressions, I find similar impacts of upholding on both measures of employment structure. I find a sharp upon-impact increase in the employee-share, immediately in the years following the upholding event. Both groups appear to be on a stable trend in the pre-upholding period, with no discernible change in trend around the vote-in event. Again, the longer time-lag group appears to be a valid counterfactual for the short time-lag group.

[30] Together, these two groups constitute 60 percent of the full set of IDB states. See Appendix Figure B.1 and Section B.2. for more details.

[31] This variable excluded all transfers received from Federal and state government. The denominator contains all sources of earned income: employee farming, employee non-farming, self employed farming, self-employed non-farming. The denominator excluces all sources of non-earned income, such as dividends and interest payments.

In Panel C, I study whether the upholding event is associated with changes in state tax policy. I first code every year in which I observe a change in the state tax code to the nominal value of the exemption threshold. I then construct the cumulative number of reforms over time, and plot it in the left hand graph of Panel C. This measure proxies for the likelihood of observing a reform to the exemption threshold, while controlling for differences in reform-frequency across states. With a short lag to the upholding event, I find an strong increase in tax reform decisions. This suggests that state policy-makers actively responded to the upholding event by changing the value of the threshold. The right hand side graph plots the location of the threshold in the state's income distribution. Withi a similar spaced lag to the upholding event, I observe a sharp decrease in the relative location. These two results suggest that policy-makers actively responded to the upholding event by lowering the nominal value of the exemption threshold in the income distribution and expanding the tax base.

4.4 Regression results

I now turn to a regression-based estimate of the IDB impacts, using 1. Results are displayed in Table 2. In column (1), I find that there is a 1.7 percentage point increase in the employee-share of employment in the short-run period of 5 years following the upholding event. In contrast, I find no impact of the vote-in event on employment structure. Column (3) shows that the upholding event led to a significant reduction in the value of the exemption threshold relative to per capita income. As I will show in the robustness section below, the upholding event did not lead to a change in per capita income. Consistent with the sharp graphical impacts in the previous sub-section, I interpret the decrease in this ratio as consistent with a decrease in the nominal value of the exemption threshold.

These results are consistent with a reduced-form impact of employment structure on the income tax base. I now study the distributional impacts. Specifically, I estimate 1 separately for the employee-share in every decile of the state income distribution.[32] I plot the decile-specific coefficient on the upholding event, $\widehat{\theta_j}$, against the ten deciles $j = 1, ..., 10$ in Figure 8, where I also demark the average percentile location of the exemption threshold in the pre-vote period. The figure reveals that the upholding event led to a statistically significant increase in the employee-share in the middle of the income distribution, between the third and the sixth decile.[33] Given the upward-sloping employee-

[32] I construct the employee-share in every decile in every state, using the decennial Census records. I then interpolate between years, using a natural cubic spline (Herriot & Reinsch, 1973).

[33] This is consistent with qualitative descriptions (Cobb, 1993), which highlights that the plants opened under IDB were

share profile in the average state (Figure 6), this distributional impact implies a leftward shift of the profile, consistent with patterns observed under stylized fact #1. At the same time, the IDB-induced increase in employee-share occurs strictly in deciles locally to the left of the pre-IDB location of the threshold. This is consistent with the patterns observed in stylized fact #2, suggesting that the threshold moved down the income distribution in response to (IDB-induced) increases in employee-share that occurred locally to its left. Finally, the employee-share in the tax base remains high after the expansion of the tax base, consistent with stylized fact #3. The close match-up between the IDB-induced distributional impacts and the stylized facts observed over development increase the external validity of the directly estimated evidence for the cross-country development setting of Section 3.

4.5 Mechanism

Both the stylized facts in Section 3 and the direct evidence from the previous sub-section are consistent with lowering of the exemption threshold caused by gradual increases in employee-share below the threshold. As the model in the following section highlights, there are several channels through which movements into employee-jobs can cause a lowering of the exemption threshold. In this sub-section, I provide results to suggest that employment structure impacts the decision to expand the base by changing the fiscal enforcement cost.

I leverage the interactions between the state tax administrations and the Federal Internal Revenue Service (IRS). Between the early 1950s and the mid 1970s, states would sign Exchange of Information Acts (EoI) with the IRS. The signing of the act formalized Federal involvement in the enforcement of state income taxes by creating a systematic forwarding of results from individual audits of federal tax returns to state tax authorities. Penniman (1980) documents qualitatively how states' audit tax collections increased in the period following the signing of an EoI. The signing of an EoI is therefore likely to have significantly decreased the state's cost of income tax enforcement.[34] This interpretation is consistent with the results on income tax collection in Troiano (2017).[35]

Under the assumption that passing an EoI allows state tax administrations to effectively outsource

predominantly in low-skill industries.

[34]The Federal IRS cost of collection was over four times smaller than the average IDB state administration's cost of collection in 1962 (Jensen & Lagakos, 2019). This suggests that the EoI provided the state with valuable enforcement information from an institution with significantly higher administrative capacity.

[35]As Troiano (2017) notes, it is also possible that the EoI led to an increase in the flow of audit information from state tax authorities to the Federal IRS, but there is no systematic data to suggest this is the case.

enforcement costs to the IRS, we can test whether the IDB-induced employee increase causes states to move the exemption threshold only when their administration bears the enforcement cost of this base expansion.[36] Formally, I estimate

$$y_{st} = \beta + \alpha \mathbf{1}\left(\text{Vote-in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \sigma \mathbf{1}\left(\text{Upheld}\right)_{st} \times \mathbf{1}\left(\text{EoI}\right)_{st} + \lambda \mathbf{X}_{st} + \mu_s + \gamma_t + \varepsilon_{st} \quad (2)$$

where $\mathbf{1}\left(\text{EoI}\right)_{st}$ takes a value of 1 if state s has signed an EoI with the IRS in year t. A compelling feature of the EoI is that its implementation is staggered across states over time, spanning the period of IDB implementation. This leads upholding states to be well-balanced across EoI status, with 57% of states having signed an EoI by the time the upholding event occurred.[37]

In Column 4 of Table 2, I find that the expansion of the income tax base is concentrated among states that had not signed an EoI by the time of the upholding event. In contrast, in states where the EoI preceded the upholding event, there is no reduced-form impact of upholding on the location of the exemption threshold (p-value of the F-test is 0.230). This is despite the fact that the upholding event increased the employee-share by a similar amount in both sets of states (Column 2, Table 2). This suggests that there exists an association between increases in employee-share (below the threshold) and reforms to lower the threshold only when the reform incurs a fiscal enforcement cost. In turn, this is consistent with the interpretation that employees differ from self-employed in cost of enforcement, and that a higher share of employees on a potential income segment lowers the enforcement cost of expanding the tax base to this new segment.

The final column of Table 2 studies the impact on income tax collection. I find a significant and positive impact of upholding only in the subset of non-EoI states. The p-value of the F-test in EoI states is 0.481. This result has two implications. First, the positive impact on collection is limited to the same sub-group where I also found an expansion of the base, suggesting that the expansion of the base drives the increase in collection. Second, there is no impact on collection in EoI states, even though there is a positive impact on employee-share in that sub-group. This suggests that, in this particular setting, the main impact of a change in employment structure on tax collection is through the expansion of the tax base.

The evidence in Table 2 thus provides direct evidence on a fiscal cost mechanism linking em-

[36] This heterogeneity analysis derives from a conversation with Roger Gordon, and I am grateful to him for the discussion.

[37] In results not reported, I find that neither the vote-in nor the upholding events led to a significant change in the likelihood of signing an EoI.

ployment structure to income tax base expansion and collection. Table 3 pursues a complimentary approach, by providing evidence against other mechanisms. I do this by constructing proxies for confounding determinants of the income tax base and tax collection, and then study changes to these proxies over the upholding event. I allow for heterogeneity across EoI states. The results are reported in Table 3. In Panel A, I investigate public finance outcomes. In Columns (1) to (3), I find that there was no impact of upholding on any other source of tax revenue: corporate income tax, general sales tax, property tax. This suggests the upholding decision did not coincide with a change in the state's general 'need' or 'taste' for revenue. Related, I also do not find any impact on income tax rates. The absence of non-income tax revenue impacts further alleviates concerns that issuing of IDB bonds led to an increase in revenue pressure, since the pressure would have to only be met by an increase in personal income tax and only in non-EoI states.[38] I proxy for changes in enforcement capacity in two ways. First, a decrease in the number of tax departments responsible for collecting distinct sources of tax revenue can reflect improved administrative capacity to centralize and cross-check information sources in order to enhance enforcement. Second, the introduction of withholding on wage income for remitting of state income taxes has been found to significantly improve income tax collection (Dusek and Bagchi, 2017). I do not find an impact on the court upholding decision on either of these two enforcement proxies. Interestingly, I also do not find an impact of the vote-in event, which suggests there is no substitution between revenue-enhancing tax policies when the upholding decision is delayed.

In Panel B of Table 3, I consider alternative confounding determinants of the income tax base. In Column (1), I first show that the upholding event did not lead to an economically important nor statistically significant change in income. The largest positive impact is found in the EoI states, where the upholding event is associated with a $36 increase in per capita income. Relative to a pre-treatment mean income of $7641, this small magnitude implies that changes in income cannot account for the expansion of the base and the increase in collection.[39] In the Appendix, I estimate a county-level empirical model of the local IDB impacts to explore the reasons why the program did not deliver any economically meaningful change in per capita income. I find no evidence to suggest that the IDB program led to sectoral spill-overs or agglomeration effects in the short run at the local level. This is consistent with the Federal description of the IDB program as being highly targeted and local in nature, while being narrow in scope (Advisory Commission on Intergovernmental Relations, 1962).[40]

[38] This confounding channel is also made less plausible by the fact that the personal income tax was a minor source of revenue collection, in the pre-vote period.

[39] I show in the Appendix Table B.2 that this result is robust to using alternative measures of per capita income.

[40]

In Column (2), I find no impacts of the upholding event on income inequality. The next two columns study the impact of IDB on the generosity of social assistance, and on labor laws. I find no impacts on either. The final two columns proxy for political competition and demand for redistribution. Here, the EoI states in the upholding period are found to lean marginally more Democratic (p-value of 0.098), but I found no impact on the income tax base and collection in that subset of states.

I explore the robustness of the main results in the Appendix (Section B.3 and Table B.3). I find similar results in a regression which removes $\mathbf{1}\left(\text{Vote-in}\right)_{st}$ and uses only the cross-sectional difference in timing of upholding as a source of variation. I find that results remain unchanged when allowing civil law and historical default states, the only determinants of the vote-uphold time-lag, to be on independent and fully non-parametric time paths. I show that results are robust to the inclusion of a linear trend interacted with cross-sectional difference in pre-vote per capita income. This control alleviates concerns that my main specification does not adequately capture differential convergence patterns in employment structure over time across states. Finally, I show that the results are robust to excluding the set of IDB states which were implementing the Tennessee Valley Authority program (Kline and Moretti, 2013), a concurrent Federal place-based development program.

Taken together, these results provide evidence on a direct impact of employment structure on income tax base expansion and collection. Potential threats to this interpretation remain, but the confounding state-year variation would have to cause a sharp, upon-impact change in the short-run following the upholding event; not be captured by any of the proxies in Table 3; and, only impact tax structure in the non-EoI states. The IDB-induced distributional impacts were found to match closely the patterns in the cross-country stylized facts, which increases the potential external validity of this direct evidence based on US states in the cross-country setting. Heterogeneity analysis suggests that the fiscal enforcement cost is an important channel through which employment structure impacts tax base expansions. In the following section, I derive a model which is consistent with this set of results.

5 Model

In this section, I provide a characterization of the reform decision to move the exemption threshold and derive formulas for the optimal size of the individual income tax base. The characterization of

This stands in contrast to previously US place-based redevelopment programs (including Kline and Moretti, 2013), which were broader in scope and which generated long-run income impacts through agglomeration effects.

the threshold is novel in the optimal income tax literature, which has so far focused on the schedule of marginal tax rates (Piketty and Saez, 2012; Saez and Stantcheva, 2016).[41] I study the conditions under which a small increase in the tax base is welfare improving, given an initial location of the threshold and a specified objective function. This local model-environment approximates well the tax reforms observed in reality of incremental changes to the nominal value of the exemption threshold. The agent environment is strongly simplified, because the aim of this section is to illustrate some of the main channels through which employment structure impacts the costs and benefits of local base expansions. At the end, I use the model to interpret the mechanism results in the US states and to quantitatively relate the US estimated elasticity to the cross-country variation in size of base.

5.1 Setup and empirical prediction

I consider a fixed distribution of income z across workers with pdf $h(z)$ and cdf $H(z)$. I assume exogenous employment shares of self-employed and employees at each income level, denoted respectively φ_z and $1-\varphi_z$. I do not model the development process that leads to changes in employment shares, but assume it follows the patterns documented in stylized fact #1.

If the agent reports income $z \geq K$, then she is liable to pay $\tau(z-K)$. Otherwise, she is not liable for income tax. K is the exemption threshold and τ is the marginal tax rate. I assume linear utility to abstract from income effects. I assume agents have access to an evasion technology which allows them to pay c in order to report income at K and fully evade taxes. This evasion technology generates bunching of reported income at K, in line with large set of evidence on evasion behavior. The cost is assumed to be infinite for employees: $c^E(z) = \infty$. For self-employed, the cost depends flexibly on total income z (due perhaps to a visibility effect) and on the distance between income and the threshold $z-K$ such that $c^{SE} = c(z, z-K) > 0$. The cost is assumed to be increasing and convex in z. In this setting, there will exist a marginal buncher at income $\bar{z}$ who is indifferent between bunching and full compliance : $\bar{z} - c(\bar{z}, \bar{z}-K) = K + (1-\tau)(\bar{z}-K)$. All self-employed with income $z : K \leq z \leq \bar{z}$ will under-report and bunch at the threshold. An increase in τ unambiguously leads to more evaders: $\frac{\partial \bar{z}}{\partial \tau} > 0$. An increase in the threshold will lead to less evaders if the marginal gain from compliance is

[41] In the indirect tax design literature, Yitzhaki (1979) and Wilson (1989) derive optimal tax formulae for the number of taxable commodities, while Keen & Mintz (2004) focus on the optimal threshold for a value-added tax.

larger than the marginal gain from under-reporting after the threshold decrease, that is

$$\frac{\partial \bar{z}}{\partial K} < 0 \text{ if } \tau > c_K\left(\bar{z}, \bar{z} - K\right) \tag{3}$$

I will assume the condition in 3 holds. The revenue base reflects evading self-employed between K and $\bar{z}$:

$$R = \int_{z \geq \bar{z}} \tau\left(z - K\right) \varphi_z dH\left(z\right) + \int_{z \geq K} \tau\left(z - K\right)\left(1 - \varphi_z\right) dH\left(z\right)$$

Consider a reform which locally decreases the threshold: $dK < 0$. This reform will have two effects on revenue: a mechanical gain and a behavioral loss. The mechanical gain, dM, reflects the increase in revenue collected on the inframarginal agents, assuming no behavioral responses

$$\begin{aligned} dM &= -dK\tau\left[\int_{z \geq \bar{z}} \tau\left(z - K\right) \varphi_z dH\left(z\right) + \int_{z \geq K} \tau\left(z - K\right)\left(1 - \varphi_z\right) dH\left(z\right)\right] \\ &\geq 0 \text{ if } dK < 0 \end{aligned} \tag{4}$$

The behavioral loss, dB, reflects loss in revenue due to behavioral responses of the marginal agents

$$\begin{aligned} dB &= -\frac{\partial \bar{z}}{\partial K} dK\tau\left(\bar{z} - K\right)\varphi_K \\ &\leq 0 \text{ if } dK < 0 \end{aligned} \tag{5}$$

where I have used the local approximation that $\varphi_K \approx \varphi_{\bar{z}}$, which is plausible if the last buncher is not located too far above the threshold. At the revenue maximizing optimum, K^{Rev}, it must be that $dB + dM = 0$. This yields the characterization for the location of the threshold

$$\frac{K^{Rev}}{\bar{z}} = \frac{1}{\left[1 + \left[\frac{\text{Mech gain}}{\text{Beh loss}}\right] \cdot \left[\varepsilon_{\bar{z},K}\varphi_K\right]^{-1}\right]} \tag{6}$$

where Mech gain $= \int_{z \geq \bar{z}} \tau\left(z - K\right) \varphi_z dH\left(z\right) + \int_{z \geq K} \tau\left(z - K\right)\left(1 - \varphi_z\right) dH\left(z\right)$, Beh loss $= h\left(\bar{z}\right)\bar{z}$, and where $\varepsilon_{\bar{z},K}$ denotes the elasticity of the marginal buncher with respect to the threshold. By changing the mass of agents who respond to the local reform, the model predicts the main empirical result

Empirical prediction: An increase in employee share locally around the threshold leads to an expansion of the income tax base through lowering of the threshold

$$\frac{\partial K^{Rev}}{\partial(1-\varphi_K)} < 0 \tag{7}$$

Extension: administrative costs

A policy literature in developing and developed countries (summarized in IMF, 2015) discusses differences in administrative costs between reconstructing information trails for self-employment earnings and aggregation of employee information trails by employers. I model the administrative cost of taxing an income segment z as an increasing function of the self-employed share on the income segment, $c(z) = c(\varphi_z)$. This unit cost is constant for all segments. Revenue net of administrative costs equals

$$\begin{aligned} R = &\int_{z\geq\bar{z}} \tau(z-K)\,\varphi_z dH(z) + \int_{z\geq K} \tau(z-K)(1-\varphi_z)\,dH(z) \\ &- \int_{z\geq K} c(\varphi_z)\,dH(z) \end{aligned}$$

The local threshold decrease $dK < 0$ will lead to an additional administrative marginal cost $dC = dK \cdot c(\varphi_K) < 0$ if $dK < 0$. The revenue maximizing threshold now equals

$$\frac{K^{AdminRev}}{\bar{z}} = \frac{1}{\left[1 + \left[\frac{\text{Mech gain}_{-dC(\varphi_K)}}{\text{Beh loss}}\right] \cdot \left[\varepsilon_{\bar{z},K}\varphi_K\right]^{-1}\right]} \tag{8}$$

where the threshold is now predicted to increase due both to behavioral distortions and administrative costs that increase as the self-employed share goes up.

Discussion: objective functions

The stylized facts and main US states reduced-form results results are consistent with an objective function of revenue maximization over the exemption threshold. An extension to the objective function which is also consistent with this evidence is to include a social preference for a 'fair tax base'. Discussed especially in a setting of low enforcement capacity countries, fairness relates to the idea that the tax base should not discriminate against particular groups in terms of compliance (Brautigam,

Fjeldstad & Moore, 2008). On the income tax base, such fairness would imply that a group's share in effective contribution to tax revenue should be equal to its share in statutory contribution. This can be formalized by modeling a mis-representation index, given by the ratio of employee-share of income on the statutory income tax base to the ratio of employee-share of income on the compliant income tax base. Society faces social loss with parameter μ from any deviation of this index from a situation of perfect representation (with index value 1)

$$\text{Horizontal inequity} = \mu \left(1 - \left| \frac{\text{Employee-share on statutory income tax base}}{\text{Employee-share on compliant income tax base}} \right| \right) \tag{9}$$

So long as self-employed evade more than employees, the inequity cost associated with a lowering of the income tax threshold, dE, will always be smaller when the employee-share at the local threshold is larger. This horizontal equity channel delivers a non-trivial prediction for movements in exemption threshold driven by gradual increases in employee share (stylized fact #2) and for constant employee share above the threshold (stylized facts #2 & #3) in the simplest possible setting of costless full evasion by self-employed (unlike the behavioral distortions channel and the administrative cost channel).[42] Future work is required on the existence and importance of horizontal equity concerns for tax policies.

5.2 Relating theory to US states direct evidence and cross-development stylized patterns

The objective of the previous sub-section has been to show that the basic empirical prediction 7 is consistent with a range of objective functions and mechanisms through which employment structure impact the decision to expand the income tax base. I now discuss how the evidence from the US states is consistent with one simplified version of 8, in which the only fiscal cost of a threshold reform is the administrative cost, which depends positively on the self-employed share locally around the threshold. I model the existence of an EoI as a cost-function which is independent of the employee-share: $c^{EoI}(z) = c(z)$. This could be because the EoI provides the state tax authority with the enforcement results of all relevant income sources, including findings on evaded income of the self-employed versus employees, such that the state tax authority did not need to duplicate the enforcement effort across employment-types. In non-EoI states, the tax authority has to incur the cost of uncovering evasion dif-

[42] In a setting where agents differ in skill level and in cost of avoidance, Kopczuk (2001) derives conditions under which high marginal tax rates can exacerbate horizontal inequity.

ferences between self-employed and employees: $c^{Non-EoI}(z) = c(\varphi_z)$. This set-up predicts

$$\frac{\partial K^{Non-EoI}}{\partial(1-\varphi_K)} < 0 \text{ and } \frac{\partial K^{EoI}}{\partial(1-\varphi_K)} = 0$$

which is consistent with the negative impact of employee-share on the threshold empirically found in non-EoI states and the absence of an impact in EoI states (Table 3).

The theory suggests that an appropriate elasticity to summarize the impact of changes in employment structure on the tax base is $\varepsilon_{[1-F(K)],[1-\varphi_z|z<K]}$. This elasticity indicates the percent increase in the size of the base, $1-F(K)$, following a 1 percent increase in the employee-share beneath the threshold, $1-\varphi_z|z<K$. I derive a value of 3.89 for this elasticity, using the identified evidence in the US states.[43] I then investigate what share of the variation in size of the base across countries (Panel B, Figure 4) can be accounted for by combining the observed cross-country variation in employee-share (below the threshold) and the US states-based value of the elasticity. I find that the elasticity-based predicted change in base can account for 29.9% of the difference in size of the tax base between a typical low-income country, India, and the US. The estimated elasticity can account for 68.3% of the variation between a typical lower-middle income country, Peru, and the US.[44]

6 Conclusion

This paper has provided evidence and supporting theory to show that transitions into employee jobs over development explains the rise of the modern tax system. I empirically identify a new channel, in which movements from self-employment into employee jobs occurring gradually further down the income distribution cause broadening of the tax base through lowering the exemption threshold. In a novel research design in taxation and development, I combine micro-descriptive with quasi-experimental evidence. The descriptive evidence is derived from a micro data-base covering 100 countries at all levels of development and the US over 140 years (1870-2010), and which allows me

[43] Using the evidence from Figure 8, the upholding event led to a 3.6 percent change in the employee-share beneath K. Using the regression results from column 3 in Table 2, the upholding event led to a 14 percent increase in the size of the base. This implies an elasticity value of $14/3.6 = 3.89$.

[44] Using Figure 1, between India and the US, there is a 347% increase in the employee-share beneath the 98th percentile (the location of the Indian exemption threshold). Using the elasticity estimate, the implied increase in the size of the tax base is 1349%, while the actual increase in size of the base between India and the US is 4500% (moving from the 2nd to the 92nd percentile of the gross income distribution). Using Figure 1, between Peru and the US, there is a 117 percent increase in the employee-share beneath the (Peruvian) threshold, which implies a predicted increase in the size of the base by 455%. The actual variation in size of the base between Peru and the US is 666%.

to measure employment-type, sources of income and tax structure in a consistent manner across space and time. The identified evidence is based on exploiting variation in the effective implementation date of a previously un-studied state-led development program implemented across US states in the 1950-1960s. To bridge the descriptive and identified evidence, I show that the identified changes in employment and tax structure replicate the stylized facts underlying the descriptive evidence, and that the mechanism uncovered in the US states is plausible in the cross-development setting.

The evidence suggests the importance of studying jointly the drivers of development and their impacts on taxation. A robust finding has been the close match between less developed countries and currently advanced countries at similar levels of development. This suggests that a small income tax base in a developing country reflects the same factors which lead to a low employee share. Future work is needed to reconcile current models of structural change with the finding that transitions into employee-jobs occur gradually further down the country's income distribution.

The existence of a constrained income tax base violates the assumptions underlying the Atkinson-Stiglitz (1976) theorem and has implications for the optimal mix in developing countries. The inability to redistribute flexibly through the income tax system suggests it may be optimal to provide redistribution through indirect taxes. This provides a normative redistributive role for indirect tax systems in developing countries (Bachas, Gadenne & Jensen, 2019), when its prominence has previously been justified on the positive grounds of its favorable revenue performance (Keen & Lockwood, 2009).

My research design has highlighted the usefulness of building micro evidence to answer macro questions. Micro evidence on enforcement gains derived from sales connections between firms (Pomeranz, 2015) can be combined with macro-development evidence on growth in interconnectedness between firms (Carvalho & Tahbaz-Salehi, 2018), to explain sales tax structure over development.

References

[1] Abbey, A. [1965]. "Municipal Industrial Development Bonds", *Vanderbilt Law Review,* 19.

[2] Acemoglu, D., J. Robinson [2000]. "Why did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective", *Quarterly Journal of Economics*, 115(4).

[3] Akcigit, U., J. Grigsby, T. Nicholas, S. Stantcheva [2018]. "Taxation and Innovation in the 20th Century", *NBER Working Paper Series*, No. 24982.

[4] Area Redevelopment Administration [1963]. *Records of the Office of Planning and Research. Several volumes*. Washington D.C: Department of Commerce and National Archives.

[5] Area Redevelopment Administration [1961]. *Statistical Profiles for Rural Redevelopment Areas. All states* . Washington D.C: Department of Commerce and National Archives.

[6] Atkinson, A., J. Stiglitz [1976]. "The Design of Tax Structure: Direct versus Indirect Taxation", *Journal of Public Economics*, 6.

[7] Bachas, P., M. Soto [2018]. "Not(ch) Your Average Tax System: Corporate Taxation Under Weak Enforcement", *World Bank Policy Research Working Paper*, No.8524.

[8] Bachas, P., L. Gadenne, A. Jensen [2019]. "Informal Consumption and Indirect Tax Design: Evidence from Consumption Diaries around the World", *Working Paper*, Harvard Kennedy School.

[9] Baicker, K., J. Clemens, M. Singhal [2012]. "The Rise of the States: U.S. Fiscal Decentralization in the Postwar Period", *Journal of Public Economics*, 96.

[10] Bakija, J. [2009]. "Documentation for a Comprehensive Historical U.S. Federal and State Income Tax Calculator Program", *Working Paper*, Williams College.

[11] Banerjee, A., E. Duflo [2007]. "The Economic Lives of the Poor", *Journal of Economic Perspectives*, 21(1).

[12] Banerjee, A., A. Newman [1993]. "Occupational Choice and the Process of Development", *Journal of Political Economy*, 101(2).

[13] Benzarti, Y. [2018]. "How Taxing is Tax Filing? Using Revealed Preferences to Estimate Compliance Costs", *NBER Working Paper Series*, No. 23903.

[14] Best, M. [2014]. "The Role of Firms in Workers' Earnings Response to Taxes: Evidence from Pakistan", *Working Paper*, Stanford Institute for Economic Policy Research.

[15] Best, M., A. Brockmeyer, H. Kleven, J. Spinnewijn, M. Waseem [2015]. "Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan", Journal of Political Economy, 123(6).

[16] Besley, T., A. Case [1994]. "Unnatural Experiments? Estimating the Incidence of Endogenous Policies", *NBER Working Paper Series*, No. 4856.

[17] Besley, T., T. Persson [2011]. *Pillars of Prosperity: The Political Economics of Development Clusters*. Princeton University Press.

[18] Besley, T., T. Persson [2014]. "Why Do Developing Countries Tax So Little?", *Journal of Economic Perspectives*, 28(4).

[19] Besley, T., T. Persson, D. Sturm [2010]. "Political Competition, Policy and Growth: Theory and Evidence from the US", *Review of Economic Studies*, 77(4).

[20] Berkowitz, D., K. Clay [2005]. "American Civil Law Origins: Implications for State Constitutions", *American Law and Economic Review*, 7(1).

[21] Bicks, A., N. Fuchs-Schundeln, D. Lagakos [2018]. "How do Average Hours Worked Vary with Development? Cross-Country Evidence and Implications", *American Economic Review*, 108(1).

[22] Bird, R., O. Oldman [1964]. *Readings on Taxation in Developing Countries*. Johns Hopkins University Press.

[23] Brautigam, D., O. Fjeldstad, M.Moore [2008]. *Taxation and State-Building in Developing Countries: Capacity and Consent*. Cambridge University Press

[24] Brockmeyer, A., M. Hernandez [2018]. "Taxation, Information, and Withholding: Evidence from Costa Rica", *Working Paper*, World Bank Group.

[25] Cage, J., L. Gadenne [2018]. "Tax Revenues and the Fiscal Cost of Trade Liberalization, 1792-2006", *Explorations in Economic History*, 70.

[26] Carillo, P., D. Pomeranz, M. Singhal [2017]. "Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement", *American Economic Journal: Applied Economics*, 9(2).

[27] Carvalho, V., A. Tahbaz-Salehi [2018]. "Production Networks: A Primer", *Annual Review of Economics,* forthcoming.

[28] Cobb, J. [1993]. *The Selling of the South: the Southern Crusade for Industrial Development 1936-1990*. University of Illinois Press.

[29] Council of State Governments [years between 1939 and 2009]. *Book of the States, all Parts*. Washington DC: The Council of State Governments.

[30] Dincecco, M., U. Troiano [2015]. "Broadening the State: Policy Responses to the Introduction of the Income Tax", *NBER Working Paper* No. 21373.

[31] Dusek, L., S. Bagchi [2017]. "Third-Party Reporting, Tax Collections, and the Size of Government: Evidence from Withholding", *Working Paper*.

[32] Feng, Y., D. Lagakos, J. Rauch [2018]. "Unemployment and Development", *Working Paper*, UC San Diego.

[33] Gallman, R. [2000]. "Economic Growth and Structural Change in the Long Nineteenth Century", in S. Engermand & R. Gallman (eds.), *The Cambridge Economic History of the United States: Volume II: The Long Nineteenth Century*. Cambridge University Press.

[34] Gerard, F., G. Gonzaga [2018]. "Informal Labor and the Efficiency Cost of Social Programs: Evidence from the Brazilian Unemployment Insurance Program", *NBER Working Paper* No. 22608.

[35] Gollin, D. [2008]. "Nobody's Business But My Own: Self-employment and Small Enterprise in Economic Development", *Journal of Monetary Economics*, 55(2).

[36] Gollin, D., D. Lagakos, M. Waugh [2013]. "The Agricultural Productivity Gap", *Quarterly Journal of Economics*, 129(2).

[37] Gordon, R., W. Li [2009]. "Tax Structure in Developing Countries: Many Puzzles and a Possible Explanation", *Journal of Public Economics*, 93(7-8).

[38] Haines, M. [2006]. *Cost of Living of Industrial Workers in the United States and in Europe*. ICPSR study 7711. Ann Arbor, Michigan: Inter-university Consortium for Political and Social Research.

[39] Herriot, J. C. Reinsch [1973]. "Algorithm 472: Procedures for Natural Spline Interpolatation (E1)", *Communications of the ACM*, 16(12).

[40] Hinrichs, H. [1966]. *A General Theory of Tax Structure Change During Economic Development*. The Law School of Harvard University.

[41] ICPSR based on US Department of Labor [2009]. *Study of Consumer Purchases in the United States, 1935-1936*. ICPSR study 8908. Ann Arbor, Michigan: Inter-university Consortium for Political and Social Research.

[42] ICPSR based on US Department of Labor [1986]. *Cost of Living in the United States, 1917-1919*. ICPSR study 8299. Ann Arbor, Michigan: Inter-university Consortium for Political and Social Research.

[43] International Labour Organization [2002]. "Report VI: Decent Work and the Informal Economy", in *International Labour Conference 90th session*. Geneva: International Labour Office.

[44] International Labour Organizaion [2009]. "International Definitions and Prospects of Underemployment Statistics", *Working Paper Series*, ILO Bureau of Statistics.

[45] International Monetary Fund [2015]. "Current Challenges in Revenue Mobilization - Improving Tax Compliance", *IMF Policy Paper*.

[46] Jensen, A., D. Lagakos [2019]. "Government Productivity in Tax Collection", *Working Paper*, Harvard Kennedy School.

[47] Keen, M. [2012]. "Taxation and Development - Again" IMF Working Paper 12/220.

[48] Keen, M., B. Lockwood [2009]. "The Value-Added Tax: Its Causes and Consequences", *Journal of Development Economics*, 92.

[49] Keen, M., J. Mintz [2004]. "The Optimal Threshold for a Value-Added Tax", *Journal of Public Economics*, 88(3).

[50] Kleven, H., M. Knudsen, C. Kreiner, S. Pedersen, E. Saez [2011]. "Unwilling or Unable to Cheat? Evidence from a Randomized Tax Audit Experiment in Denmark", *Econometrica*, 79(3).

[51] Kleven, H., C. Kreiner, E. Saez [2016]. "Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries", *Economica*, 83.

[52] Kline, P. E. Moretti [2014]. "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs", *Annual Review of Economics*, 6.

[53] Kline, P., E. Moretti [2013]. "Local Economic Development, Agglomeration Economies and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority", *Quarterly Journal of Economics*, 129(1).

[54] Kopczuk, Wojciech [2011]. "Redistribution when Avoidance Behaviour is Heterogeneous", *Journal of Public Economics*, 81(1).

[55] La Porta, R., A. Shleifer [2014]. "Informality and Development", *Journal of Economic Perspectives*, 28(3).

[56] La Porta., R., F. Lopez-de-Silanes, A. Shleifer [2008]. "The Economic Consequences of Legal Origins", *Journal of Economic Literature*, 46(2).

[57] Lindert, P., J. Williamson [2016]. *Unequal Gains: A New History of American Growth and Inequality*. Princeton University Press.

[58] Moody's Investor Service [1974]. *Municipal and Government Manual of Issuers*. New York: Moody's Investor Service Inc.

[59] Musgrave, R. [1966]. *Fiscal Systems*. New Haven: Yale University Press.

[60] Musgrave, R. [1981]. "Income Taxation of the Hard-to-Tax Groups", in R. Musgrave (editor), *Fiscal Reform in Bolivia*, Harvard University Press.

[61] Naritomi, J. [2018]. "Consumers as Tax Auditors", *Working Paper*, London School of Economics.

[62] Olken, B., M. Singhal [2011]. "Informal Taxation", *American Economic Journal: Applied Economics*, 3(4).

[63] Penniman, C. [1980]. *State Income Taxation*. Johns Hopkins University Press.

[64] Pinsky, D. [1972]. "State Constitutional Limitations on Public Industrial Financing: an Historical and Economic Approach", *University of Pennsylvania Law Review*, 111(3).

[65] Pomeranz, D. [2015]. "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax", *American Economic Review*, 105(8).

[66] Rollinson, M. [1976]. *Small Issue Industrial Development Bonds*. Capital Publishing Corporation.

[67] Saez. E., T. Piketty [2012]. "Optimal Labor Income Taxation", *NBER Working Paper Series*, No. 18521.

[68] Saez, E., S. Stantcheva. [2016]. "Generalized Social Welfare Weights for Optimal Tax Theory", *American Economic Review*, 106(1).

[69] Sbragia, A. [1996]. *Debt Wish: Entrepreneurial Cities, U.S. Federalism, and Economic Development*. University of Pittsburgh Press.

[70] Serrato, J.C., O. Zidar [2016]. "Who Benefits from State Corporate Tax Cuts? A Local Labor Market Approach with Heterogeneous Firms", *American Economic Review*, 106(9).

[71] Serrato, J.C., O. Zidar [2018]. "The Structure of State Corporate Taxation and its Impacts on State Tax Revenues and Economic Activity", *Journal of Public Economics*, 167.

[72] Snavely, K. [1988]. "Innovations in State Tax Administrations", *Public Administration Review*, 48(5).

[73] Tazhitdinova, A. [2018]. "Reducing Evasion through Self-Reporting: Evidence from Charitable Contributions", *Journal of Public Economics*, 165.

[74] Troiano, U. [2017]. "Intergovernmental Cooperation and Tax Enforcement", *NBER Working Paper Series*, No. 24153.

[75] US Advisory Commission on Intergovernmental Relations [1963]. *Industrial Development Bond Financing: A Commission Report*. Washington D.C., US: Advisory Commission on Intergovernmental Relations.

[76] US Internal Revene Service [years between 1945 and 2000]. *Statistics of Income: Individual Income Tax Returns*. Washington DC: Department of the Treasury.

[77] Wallis, J. [2000]. American Government Finance in the Long-Run: 1790 to 1990", *Journal of Economic Perspectives*, 61-82.

[78] Wilson, J. [1989]. "On the Optimal Tax Base for Commodity Taxation", *American Economic Review*, 79(5).

[79] Yale L.J. [1961]. "The Public Purpose of Municipal Financing for Industrial Development", *Yale Law Journal*, 70(5).

[80] Yitzhaki, S. [1979]. "A Note on Optimal Taxation and Administrative Costs", *American Economic Review*, 475-480.

FIGURE 1: EMPLOYEE SHARE OVER INCOME DISTRIBUTION AND DECREASE IN INCOME TAX EXEMPTION THRESHOLD

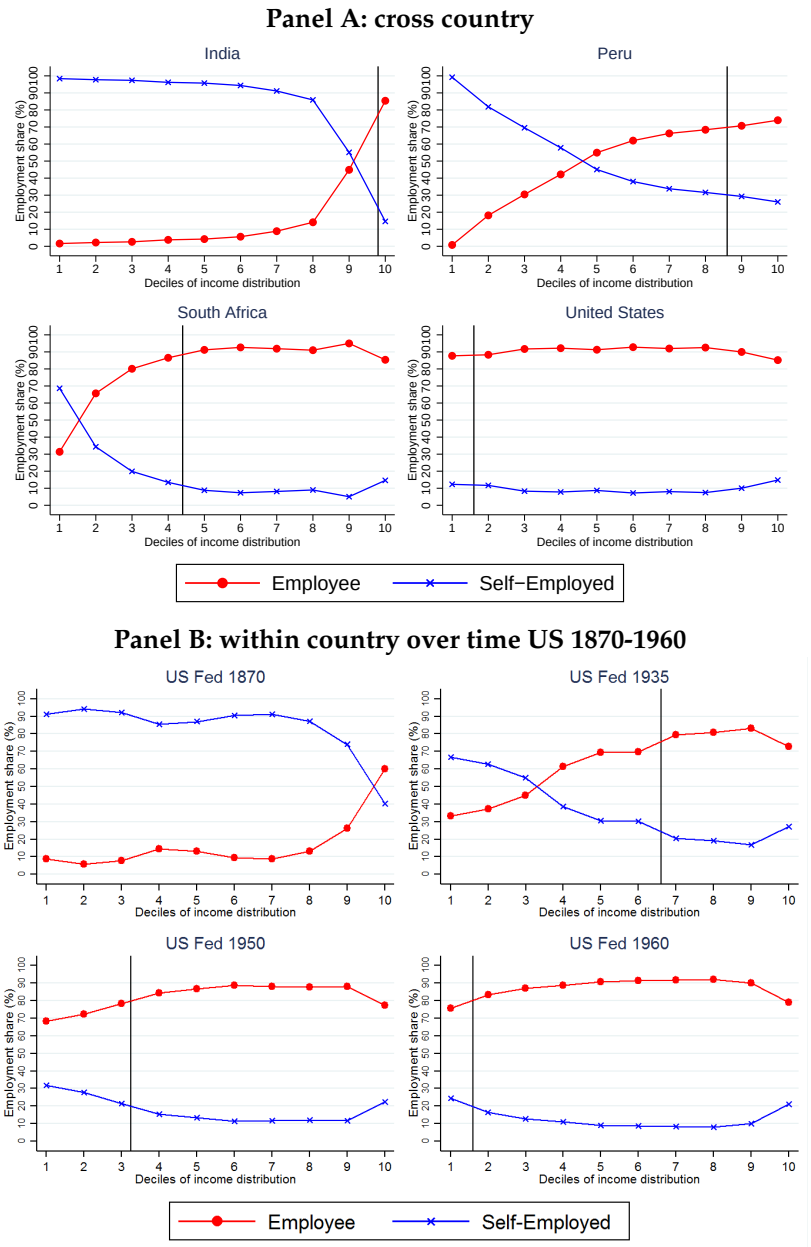

Notes: These graphs plot the employment-shares of employees and self-employed over deciles of the income-distribution, for different countries (Panel A) and within the US over time (Panel B). The share of each work-type is defined as the share of total employment in the decile of the income-distribution. Employees are defined as individuals working in jobs whose activity generate information trails that can be leveraged for income tax enforcement (Section 3.2). In each graph, the black solid denotes the location of the personal income tax (PIT) exemption threshold, which is the value of gross income above which a tax-filer becomes liable to pay income taxes (Section 3.2). Each graph is constructed from nationally representative micro surveys. Source: Section 3.2 and Appendix Section A.1.

FIGURE 2: EMPLOYEE SHARE: REPRESENTATIVE DEVELOPMENT PROFILES

Profile for average country at $1065 pc [LHS] and $2226 pc [RHS]

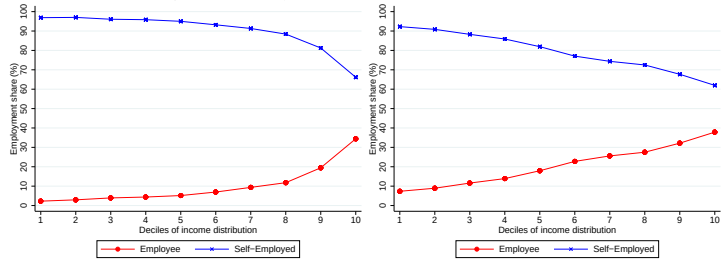

Profile for average country at $3239 pc [LHS] and $5796 pc [RHS]

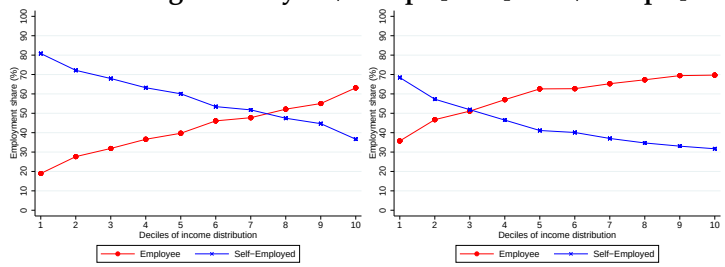

Profile for average country at $8826 pc [LHS] and $11257 pc [RHS]

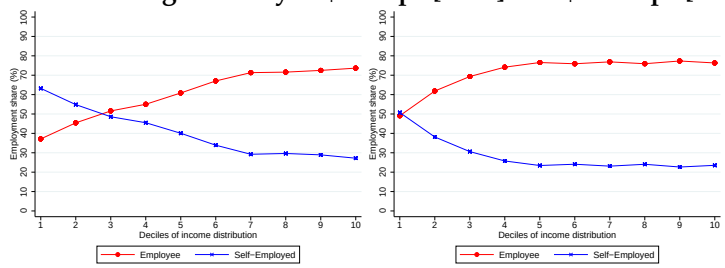

Profile for average country at $17141 pc [LHS] and $27960 pc [RHS]

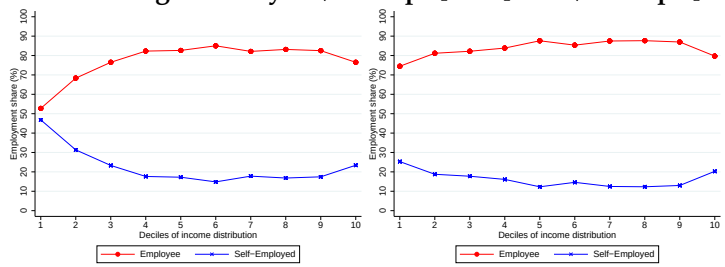

Profile for average country at $38224 pc [LHS] and $53878 pc [RHS]

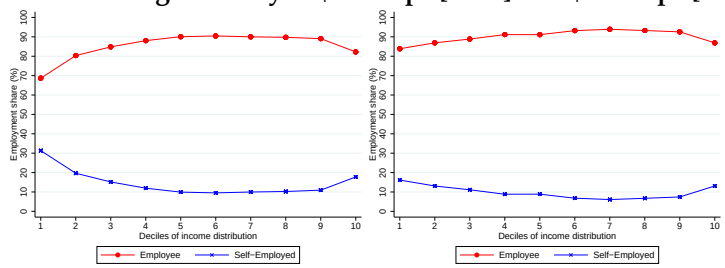

Notes: These figures plot the employment-shares of employees and self-employed over deciles of the income-distribution, for representative countries at different levels of per capita income. The share of each work-type is defined as the share of total employment in the decile of the income-distribution. Employees are defined as individuals working in jobs whose activity generates information trails that can be leveraged for income tax enforcement (Section 3.2). To construct this graph, I partition the cross-country sample into ten groups of ten countries, based on their level of per capita income. Within each group, I calculate the unweighted average of the employee-share and the self-employed share profiles across the 10 countries. I plot this average profile for every group, and indicate the average per capita income of the group. I use expenditure-side real GDP at chained PPPs in 2011 US$ from the same year as the country-survey year. Source: Section 3.2 and Appendix Section A.1.

FIGURE 3: EMPLOYEE SHARE: WITHIN US AND PAIRED SYNTHETIC CROSS-COUNTRY PROFILES

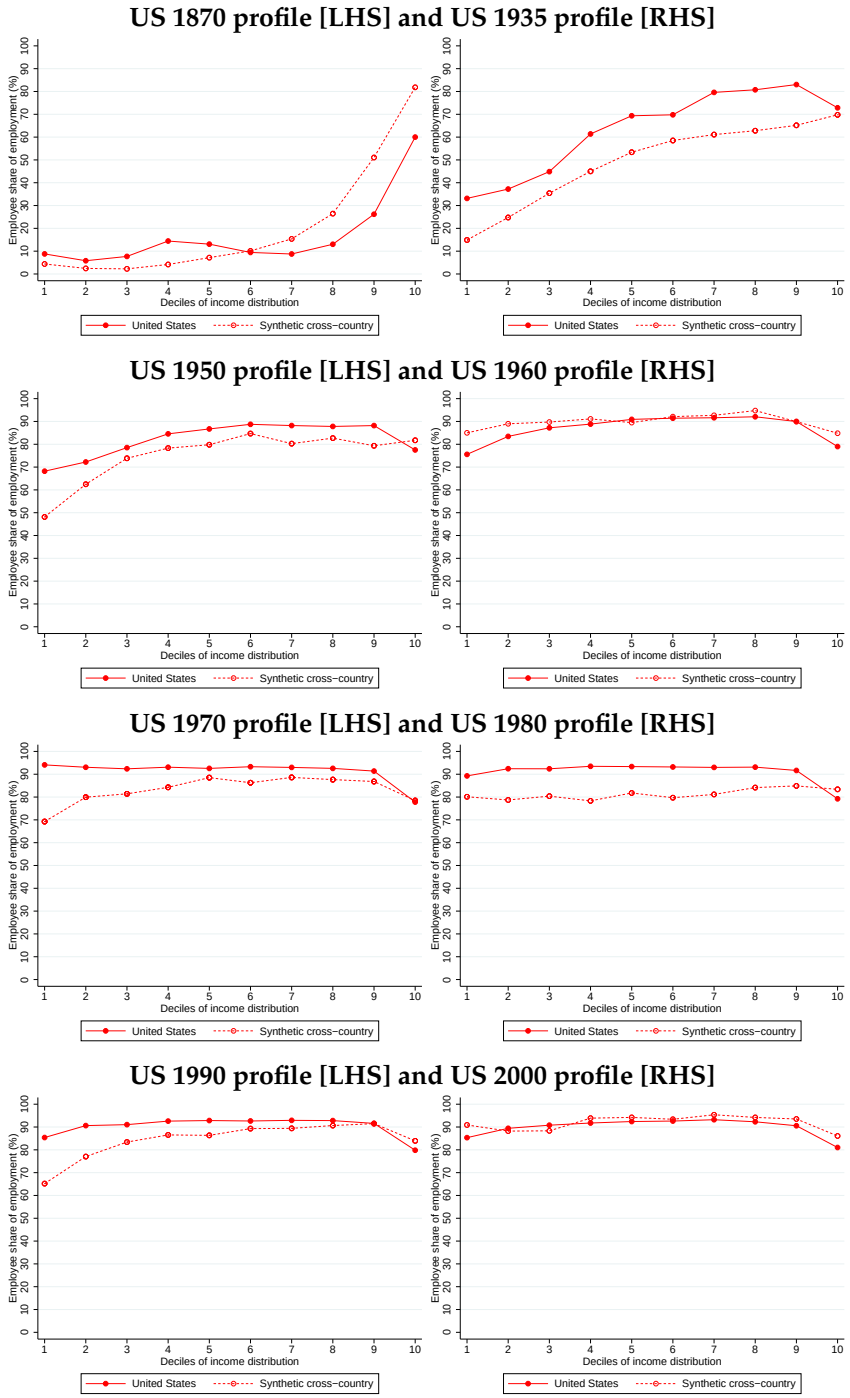

Notes: These panels plot employee-shares over the income distribution in the US between 1870 and 2000, in comparison to employee-profiles based on current countries at similar levels of per capita income. In every graph, the solid red line indicates the employee-share of economically active employment across deciles of individual gross income in the US, in the indicated year. In every graph, the hollow and dashed line indicates the employee-share of economically active employment across deciles of individual gross income in a synthetic country based on the cross-country sample. This synthetic country is constructed as the average of the countries in the contemporaneous cross-country sample which have real per capita income that lies within 10% of the per capita income of the US in the historical year. As an example, the synthetic country for the US in 1870 includes contemporaneous India and Pakistan, while the synthetic country for the US in 1935 includes Indonesia and Peru. I use the expenditure-side real GDP at chained PPPs in 2011 US$ from Penn World Tables to construct the synthetic country profiles. Source: Section 3.2-3.3 and Appendix Section A.1.

FIGURE 4: EMPLOYEE SHARE, SIZE AND EMPLOYMENT-COMPOSITION OF INCOME TAX BASE

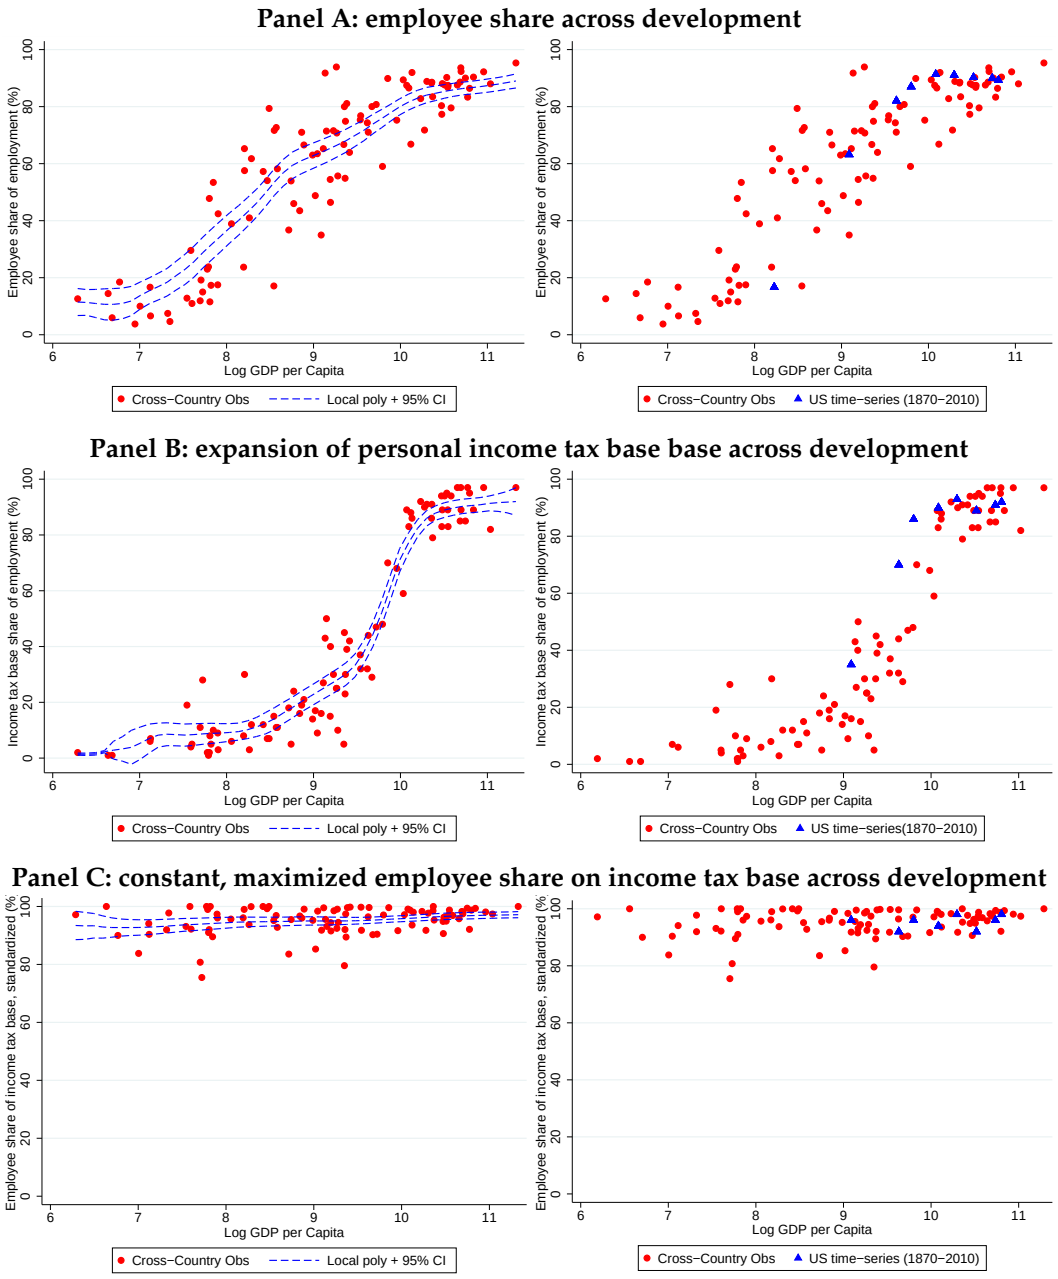

Notes: These figures document changes in employment and tax structure across development, using the 100 countries in the cross-country micro-database. In each panel, the left-hand scatter-plot is based on the full cross-country sample, while the right-hand scatter-plot overlays the US historical series onto the cross-country plot. Income per capita in the cross-country sample and in the US historical sample is measured with expenditure-side real GDP at chained PPPs in 2011 US$ from Penn World Tables. In the left-hand plot, the dashed lines indicate a local polynomial fit with 95% confidence bands. Panel A plots the employee-share of employment against per capita income. Employee-jobs are defined as those which activity generate information trails that can be used for income tax enforcement . Employment is the total number of individuals in the economically active workforce, following the definition of the U.N. System of National Accounts. Panel B plots the size of the income tax base against per capita income. The size is defined as the sum of percentiles in the gross individual income distribution that lie above the income tax exemption threshold, and represents the share of the economically active workforce that is liable to pay income taxes. See Section 3.2 for more details on the measurement of employee-jobs and the exemption threshold. Panel C plots the employee-share composition of the tax base. This is constructed as the employee-share of employment whose gross income lies above the exemption threshold. This share is expressed relative to the saturated share, which is calculated as the highest possible employee-share that can be obtained by locating the threshold anywhere in the gross income distribution. The ratio of the actual to the saturated share will therefore be equal to 1 if the actual location of the threshold maximizes the employee-share above the threshold in every country's income distribution. Source: Section 3.2 and Appendix Section A.1.

FIGURE 5: INCOME TAX BASE, RATE AND COLLECTION ACROSS DEVELOPMENT

Panel A: PIT base, rate and collection across development

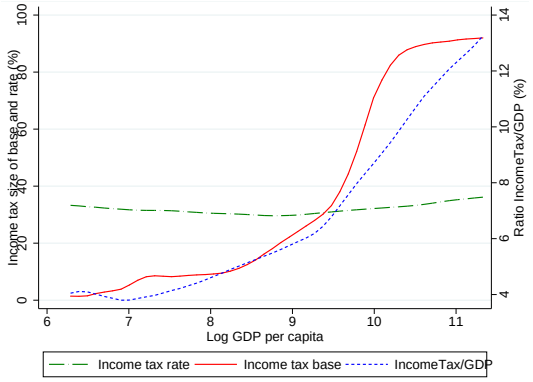

Panel B: conditional association between PIT collection and base [LHS], rate [RHS]

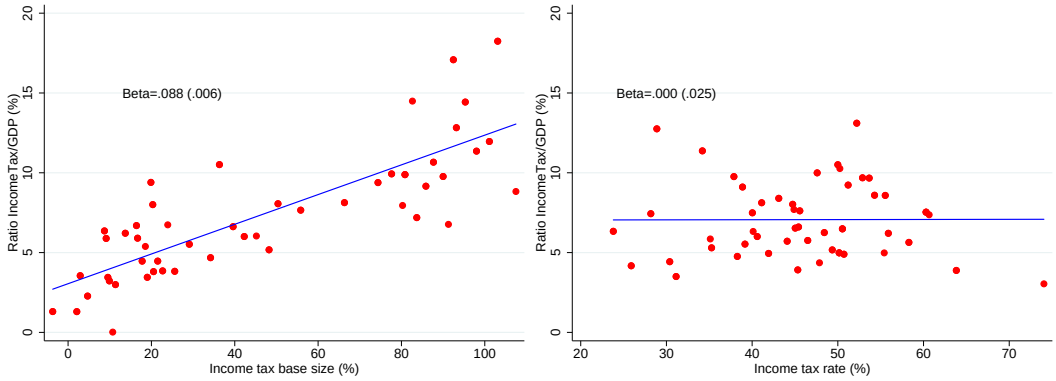

Panel C: collection across development, conditional on PIT base [LHS], rate [RHS]

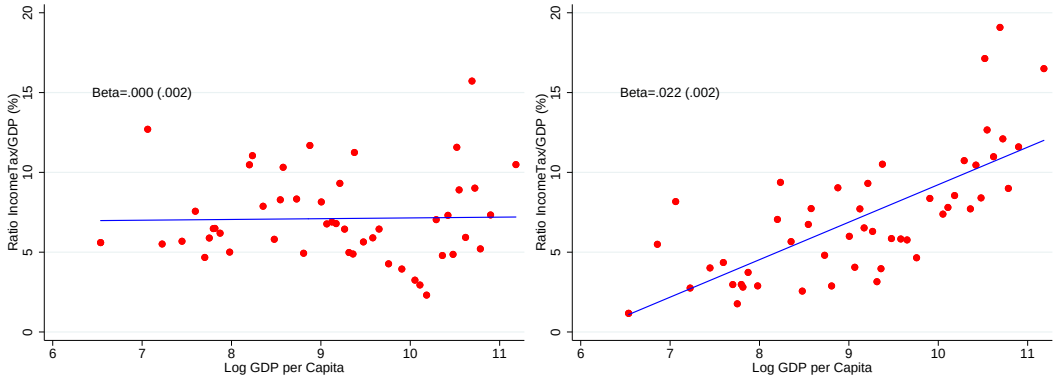

Notes: All graphs use the 100 countries contained in the cross-country micro data-base. Panel A plots the evolution of PIT collection, base, and rate across development. Each line corresponds to a local polynomial fit between log GDP per capita: income tax to GDP (red solid line); size of income tax base (blue short dash line); and, top marginal income tax rate (MTR, green long dash line). The size of the income tax base in a country is defined as the sum of percentiles in the gross individual income distribution that lie above the income tax exemption threshold, and represents the share of the economically active workforce that is liable to pay income taxes. Panel B plots the conditional association between income tax to GDP and income tax base (MTR), controlling for MTR (income tax base). Per example, to construct the right-hand side figure, I first regress both [income tax/GDP] on income tax base on MTR, and calculate residuals. I then group observations into fifty equal-sized (2 percentile-point) bins based on the tax base residuals, and scatter the means of [income tax/GDP] and tax base residuals within each bin, adding back the sample mean of each variable to ease interpretation. Panel C plots the association between residual income tax and development, after regressing income tax on PIT base (left-hand graph) or MTR (right-hand graph). In Panels B and C, the solid line shows the best linear fit, estimated on the underlying cross-country data. Source: Section 3.2. and Appendix Section A.1.

FIGURE 6: RISE OF INCOME TAX, TAX BASE EXPANSION AND EMPLOYEE SHARE: US STATES

Panel A: State income tax share of total taxes: all states 1939-2010

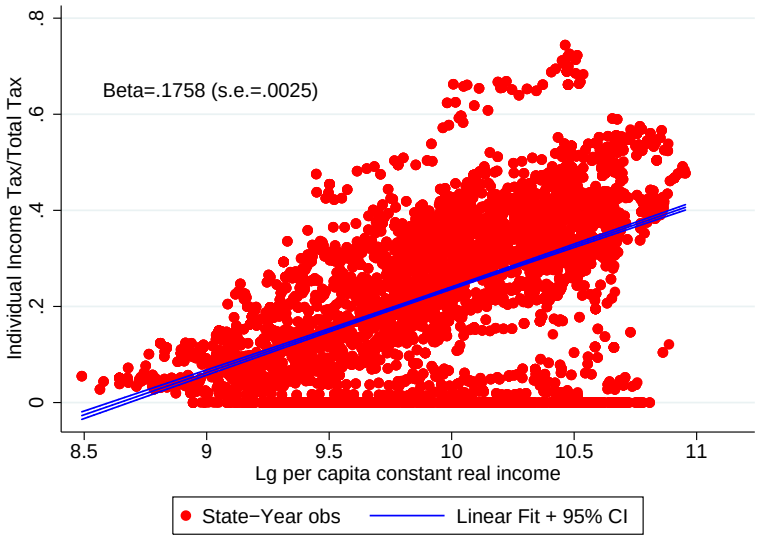

Panel B: employee share and state exemption threshold: average state 1950-1980

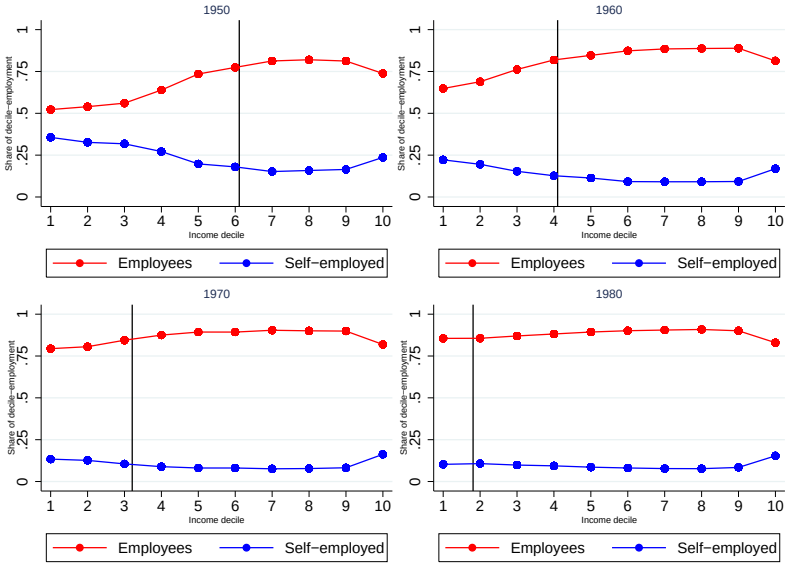

Notes: Panel A plots the state income tax share of total state taxes, using all state-year between 1939 and 2010, against log per capita income. State income tax is the sum of corporate income and personal income tax revenues. The tax revenue data is from the State Government Finance series, published by the Census. The real per capita income is constructed as the per capita income in a state-year from the historical BEA series, deflated by the historical CPI. The solid lines denote the linear fit with a 95% confidence interval from the regression on the full underlying state-year observations. Panel B plots the employment-shares of employees and self-employed over deciles of the gross income income distribution, for the average state in the US between 1950 and 1980. An employee-job is defined as a job which activity generates an information trail that is relevant for income tax enforcement. Employees are expressed as a share of the total economically active workforce. In every graph, the black solid line represents the location of the State income tax exemption threshold in the gross income distribution. This threshold represents the value of gross income above which single filer becomes liable to pay State income taxes, and is constructed using the state tax-calculator from Bakija (2009). For more details on the construction of the exemption threshold and employee-job, please see Section 4.1. Every graph in Panel B represents the average US continental State in the indicated year. Source: Section 4.1 and Appendix B.1.

FIGURE 7: IMPACT OF UPHOLDING EVENT ON OUTCOMES

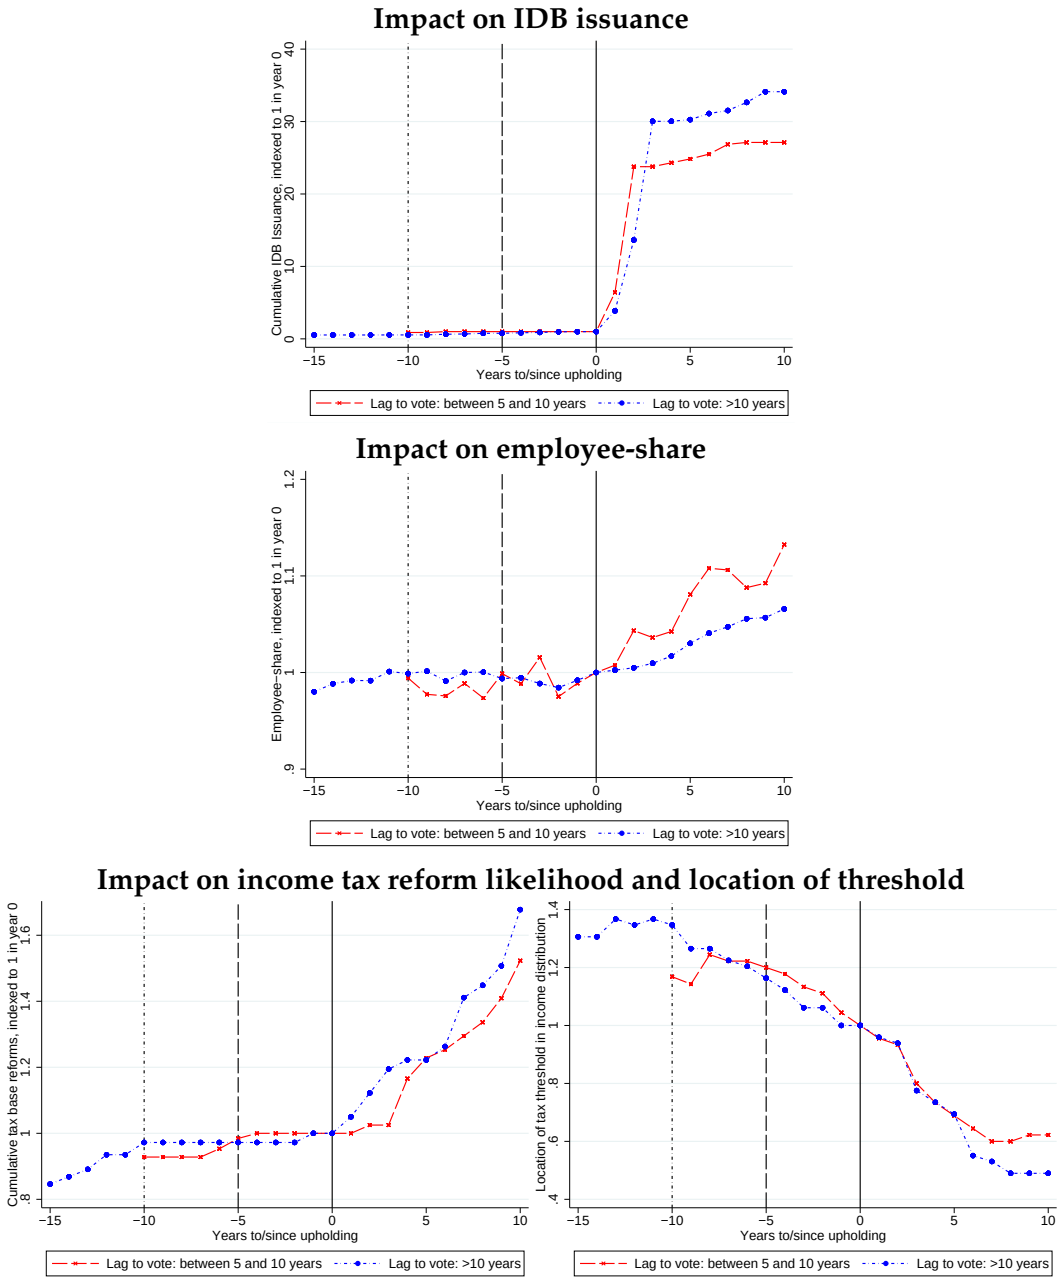

Notes: These graphs provide graphical evidence on the impact of the Industrial Development Bonds (IDB) program. Every graph shows two series: the long-dash, square line indicates the group of States for which the lag between the vote and upholding event lasted between 5 and 10 years; the short-dash, circle line represents that group of States which the lag between vote and upholding lasted more than 10 years. The vote event is the year that the State legislature decided to vote on an IDB statute or provision, while the upholding event is the year that the voted IDB statute or provision was upheld as constitutional by the State supreme court. For more details on these events and the causes of the lag between them, please see Section 4.1. In every series, the variable of interest is indexed to 1 in the year of the upholding event, which is denoted with a solid vertical line. The distance between the vertical and the long-dashed (short-dashed) line demarks the intermediary time-period after the vote event but before the upholding event in the States for which the lag from vote to upholding was between 5 and 10 years (longer than 10 years). In Panel A, the outcome is the cumulative number of IDB issuances +1 (to avoid dividing by 0 in years with no cumulative issuances). In Panel B, the outcome is the employee-share of gross income. In the left-hand graph of Panel C, the outcome is the cumulative number of legislative reforms to the nominal value of the State income tax exemption threshold; in the right-hand graph, the outcome is the decile-location of the income tax exemption threshold in the State's income distribution. Source: Section 4.1-4.2 and Appendix Section B.2.

FIGURE 8: DISTRIBUTIONAL IMPACTS OF UPHOLDING EVENT ON EMPLOYEE SHARE

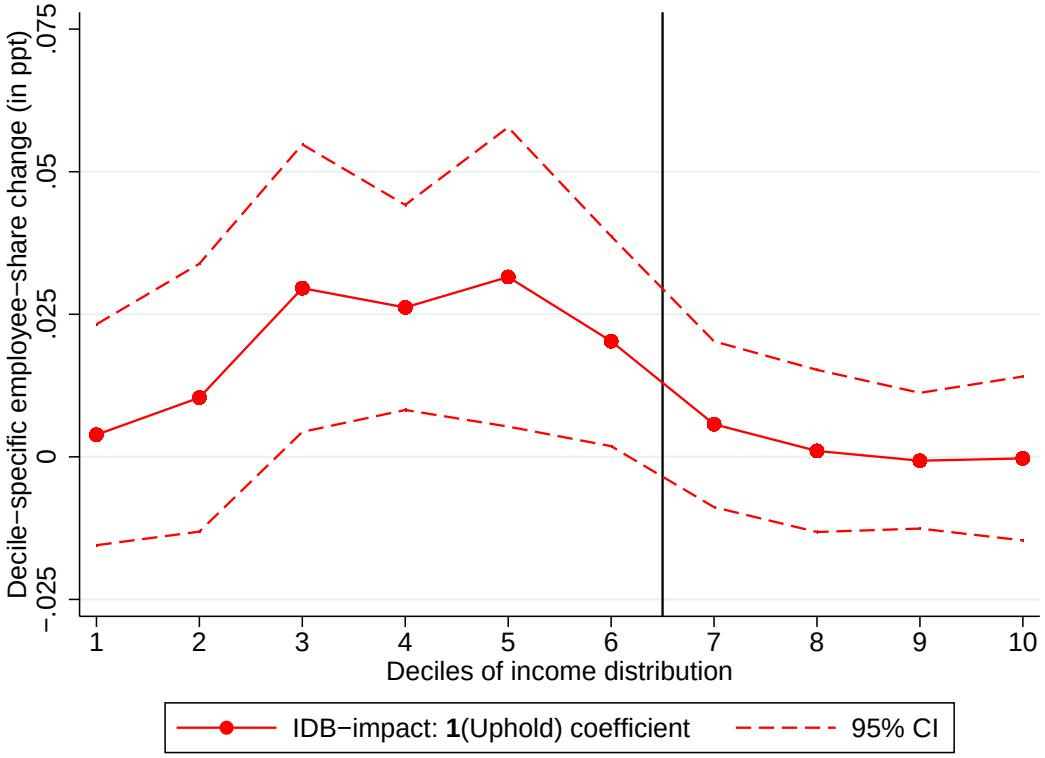

Notes: This graph displays the impacts of the IDB upholding event on the distribution of employee-share across deciles of the State income distribution. The graph plots coefficients $\widehat{\theta_j}$ on the $\mathbf{1}(\text{Uphold})$ indicator in a regression on employee-share in decile $j = 1, ..., 10$, using specification 1. Each circle denotes the decile-j point estimate $\widehat{\theta_j}$. The dashed lines denote the 95% confidence interval of the point-estimate, where standard errors are clustered at the state level. The black solid line denotes the average location of income tax exemption threshold in the IDB States, calculated in the year prior to the vote event. Source: Section 4.4 and Appendix Section B.1.

TABLE 1: DETERMINANTS OF CHANGE TO IDB LITIGATION-STATUS

	LHS=1(Upheld)							
Panel A: Politics								
1(Civil law)	5.435	12.781	12.117	9.789	12.558	10.291	12.956	12.320
	(3.201)***	(5.212)***	(4.592)***	(4.815)***	(5.489)***	(3.884)***	(5.031)***	(4.921)***
1(Hist rail default)		.337	.567	.378	.340	.671	.336	.428
		(.105)***	(.334)	(.120)***	(.116)***	(.449)	(.108)***	(.215)*
Covariate	-	-	SupremeC appointed?	Election dummies	Voting restrictions	Population total	Neighbor uphold?	RightWork laws?
Signif 10 pct?	-	-	N	N	N	N	N	N
Panel B: Economics								
1(Civil law)	11.020	10.266	11.029	12.803	13.102	10.547	12.776	11.461
	(6.981)***	(3.905)***	(4.182)***	(5.134)***	(6.031)***	(4.370)***	(5.221)***	(5.988)***
1(Hist rails)	.358	.677	.571	.332	.294	.461	.338	.358
	(.092)***	(.455)	(.350)	(.148)**	(.078)***	(.208)*	(.121)***	(.115)***
Covariate	Redev counties	LabForce growth	Manuf growth	Employee share	Income per cap	PIT base	PIT take	TotT take
Signif 10 pct?	N	N	N	N	N	N	N	N

Notes: This table reports the results of estimating non-parametric Cox proportional hazard models, where hazard rates are reported. The unit of observation is a state-year. A State enters the sample in the year where the State legislature votes in a statute or provision authorizing the IDB program. The State drops from the sample once the highest instance of the State court system upholds the constitutionality of the voted statute or provision in a leading court case. Each column in each panel reports the coefficients on a dummy for civil law origins, a dummy for historical rail defaults, as well as a description of the included covariate and whether the included covariate is significant at the 10 percent level. The test for significance relates to the null that the coefficient is equal to one. Panel A includes political covariates that could determine a change in litigation status: whether State supreme court members are appointed or elected; whether it is a State election year; whether voter restrictions are in place; the total population size; whether a geographically adjacent State has upheld the IDB statute; and, whether right-to-work laws are in place. Panel B includes economic potential determinants: the share of counties in a State that are classified as "redevelopment areas" and eligible for local development financing according to Federal criteria; the yearly labor force growth; the yearly manufacturing workforce growth; the employee-share of employment; the income per capita; the size of the State income tax base; the ratio of State income taxes to GDP; the ratio of State total taxes to GDP. $*, **, ***$ denote significance at the 10 percent, 5 percent, 1 percent level. Standard errors robust to clustering at the state level. Sources: Section 4.1-4.2 and Appendix Section B.2.

TABLE 2: EFFECTS OF IDB PROGRAM ON EMPLOYMENT AND INCOME TAX OUTCOMES

	(E-share)		(K/y)		Log(PIT/GDP)
	(1)	(2)	(3)	(4)	(5)
1(Vote)	.003	.003	-.357	-.343	.027
	(.005)	(.235)	(.232)	(.232)	(.039)
1(Uphold)	.017	.015	-.639	-.794	.176
	(.005)***	(.006)***	(.278)**	(.334)**	(.083)**
1(Uphold)x**1**(EoI)		.003		.356	-.244
		(.007)		(.197)*	(.137)*
F-test: **1**(Uphold) + **1**(Uphold)x**1**(EoI)		10.17		1.50	0.51
(p-value)		(.003)		(.230)	(.481)
Mean outcome variable	0.771	0.771	7.084	7.084	.972
State FE	x	x	x	x	x
Year FE	x	x	x	x	x
State-year controls	x	x	x	x	x
States	28	28	28	28	28
State-year Obs	466	466	466	466	466

Notes: This table reports results from estimating the following regression

$$y_{st} = \beta + \alpha \mathbf{1}\left(\text{Vote-in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \beta \mathbf{1}\left(\text{Upheld}\right)_{st} \times \mathbf{1}\left(\text{EoI}\right)_{st} + \lambda \mathbf{X}_{st} + \mu_s + \gamma_t + \varepsilon_{st}$$

where s denotes state and t denotes time. $\mathbf{1}\left(\text{Vote-in}\right)_{st}$ indicates whether a vote has occurred in the state House to allow issuance of IDB but the IDB has not yet been upheld, $\mathbf{1}\left(\text{Upheld}\right)_{st}$ indicates whether the State court system has upheld the constitutionality of the voted IDB statute or provision. The vote-in and upholding events are mutually exclusive. $\mathbf{1}\left(\text{EoI}\right)_{st}$ is an indicator variable taking a value of 1 when a State has passed an exchange of information agreement with the Federal Internal Revenue Service. In Columns 1-2, the outcome variable is the employee-share of the economically active workforce. In Columns 3-4, the outcome variable is the ratio of the State income tax exemption threshold, K, to State per capita income, y. K represents the value of gross income above which an individual filer becomes liable to pay State income taxes. In Column 5, the outcome variable is (log) ratio of State individual income taxes to GDP. In the middle of the table is reported the mean of the outcome variable in the period before the vote event. The state-year controls, $\mathbf{X}_{st}$, are indicator variables for election year, and indicator variables for the existence of voting restrictions in the form of poll tax and literacy tests. These are the first stage instruments used by Besley et al. (2010) to study political competition and policy-making in US states. $*, **, ***$ denote significance at the 10%, 5%, 1% level. Robust standard errors clustered at the state level in parentheses. Source: Section 4.2 and Appendix Section B.1.

TABLE 3: EFFECTS OF IDB PROGRAM ON CONFOUNDING OUTCOMES

Panel A

	Log(CorpIncTax/GDP)	Log(SalesTax/GDP)	Log(PropertyTax/GDP)	**1**(PIT Withholding)	PIT MTR	Tax Administration
	(1)	(2)	(3)	(4)	(5)	(6)
1(Vote)	-.030	.067	-.038	-.000	.123	.006
	(.050)	(.056)	(.026)	(.001)	(.104)	(.043)
1(Uphold)	-.058	-.047	-.030	-.001	.034	-.025
	(.136)	(.051)	(.049)	(.002)	(.152)	(.080)
1(Uphold)x**1**(EoI)	.003	.114	-.056	-.000	.006	-.020
	(.101)	(.098)	(.070)	(.001)	(.142)	(.076)
Mean outcome variable	.991	2.521	.860	.014	.129	2.492
State FE	x	x	x	x	x	x
Year FE	x	x	x	x	x	x
State-year controls	x	x	x	x	x	x
States	28	28	28	28	28	28
State-year Obs	466	466	466	466	466	466

Panel B

	Income per Capita	Top 1 percent income share	Max Unemp Benefits	**1**(Right to Work Laws)	Political Competition	Democratic Vote Share
	(1)	(2)	(3)	(4)	(5)	(6)
1(Vote)	27.401	.162	-17.910	.008	-.001	-.022
	(26.799)	(.114)	(9.226)*	(.009)	(.009)	(.019)
1(Uphold)	-8.165	.040	-10.448	.050	-.004	-.016
	(35.421)	(.223)	(18.530)	(.046)	(.014)	(.017)
1(Uphold)x**1**(EoI)	35.958	.067	-2.422	-.037	.004	-.022
	(36.605)	(.202)	(18.085)	(.035)	(.009)	(.013)*
Mean outcome variable	7641.297	13.433	307.366	.234	-.112	.539
State FE	x	x	x	x	x	x
Year FE	x	x	x	x	x	x
State-year controls	x	x	x	x	x	x
States	28	28	28	28	28	28
State-year Obs	466	466	466	466	466	466

Notes: This table reports results from estimating the following regression

$$y_{st} = \beta + \alpha \mathbf{1}\left(\text{Vote-in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \beta \mathbf{1}\left(\text{Upheld}\right)_{st} \times \mathbf{1}\left(\text{EoI}\right)_{st} + \lambda \mathbf{X}_{st} + \mu_s + \gamma_t + \varepsilon_{st}$$

where s denotes state and t denotes time. $\mathbf{1}\left(\text{Vote-in}\right)_{st}$ indicates whether a vote has occurred in the state House to allow issuance of IDB but the IDB has not yet been upheld, $\mathbf{1}\left(\text{Upheld}\right)_{st}$ indicates whether the State court system has upheld the constitutionality of the voted IDB statute or provision. The vote-in and upholding events are mutually exclusive. $\mathbf{1}\left(\text{EoI}\right)_{st}$ is an indicator variable taking a value of 1 when a State has passed an exchange of information agreement with the Federal Internal Revenue Service. In Panel A, the outcome variables are: ratio of corporate income tax to GDP; ratio of sales taxes to GDP; ratio of property taxes to GDP; an indicator for whether a State has implemented withholding on State income taxes; the State top marginal income tax rate; the number of State tax departments. In Panel B, the outcome variables are: income per capita; the top 1% share of gross income; the maximum State weekly unemployment benefits; an indicator for whether right-to-work laws are in place; the measure of political competition from Besley et al. (2010); and, the Democratic vote-share across all State-specific elections. The state-year controls, $\mathbf{X}_{st}$, are indicator variables for election year, and indicator variables for the existence of voting restrictions in the form of poll tax and literacy tests. $*, **, ***$ denote significance at the 10%, 5%, 1% level. Robust standard errors clustered at the state level in parentheses. Source: 4.2 and Appendix Section B.1.

Online Appendix for 'Employment Structure and the Rise of the Modern Tax System'

Anders Jensen *

January 2019

*Harvard Kennedy School and NBER. Contact: anders_jensen@hks.harvard.edu.

A Cross-development material

A.1 Data sources and construction of variables

In this sub-section, I provide additional details on the novel micro data-base used in Section 3. I first outline the underlying data and construction of variables in the cross-section of countries. I then outline the data and variables construction used in the historical US time-series.

A.1.1 Cross-country: sources and methodology

The cross-country database contains micro-data collected from 100 countries around the world, to document changes in employment structure transformation in as many incremental stages over development as possible. I chose to focus on countries with at least 1 million citizens. The selection of a survey in a particular country had to satisfy three criteria. First, it must be nationally representative. Second, it must survey respondents in all forms of work arrangement as opposed to, per example, only salaried workers. Third, it must contain continuous information on all sources of income, instead of, say, only wage earnings.

Given these criteria, the preferred type is a living conditions survey. This type of survey will often dominate a labor force survey, for three reasons. First, the living conditions survey usually contains information on a broader range of income sources which, especially in the context of less-developed countries, can be quite important in order to construct the lower deciles of the country's income distribution. Second, it is not always clear what the underlying sample design is for the labor force survey, and it could potentially omit individuals which in the context of this study should be included in the survey, such as casual wage day laborers and household family workers; on the other hand, the scope of a living conditions survey is usually to assess the conditions of a nationally representative sample of individuals, which should include all the alternative work type patterns. Third, the sample size of a living condition survey is typically larger than that for a labor force survey, which does not have to imply better quality of data, but usually is due to sampling design which attempts to survey all geographical areas in the country. Basic health and demographics surveys are discarded, because they do not contain information on work arrangements and income.

The data collection effort resulted in 100 surveys, which are detailed in Table A.1, displaying for each country: the year of the survey; the per capita income group; the survey type; the coverage; the sample size; and, the original source. The income group corresponds to the World Bank classification

of the country in the year of the survey. The micro data-base covers all levels of development: 20% of surveys from low-income countries; 28% from lower-middle income countries; 21% from upper-middle income countries; and, 31% from high-income countries. 93 out of the 100 data-sets are living condition surveys, 5 are labor force surveys, and the remaining 2 are censuses. In low and lower-middle countries, I obtain almost all surveys directly from the national statistics office, or the relevant government agency. In these countries, the average sample size is substantially larger than the corresponding Living Standards and Measurement Survey (LSMS) from the same country.

The construction of the employee variable is based on questions similar to the 'class of worker' question in the US Census. All cross-country surveys were chosen to ensure the highest possible international comparability. Two features in particular serve that purpose, and are common across all surveys. The first feature is the high level of detail in the categories of the 'worker-class' question. In all surveys, I can therefore distinguish between employees and employers. This removes the possibility that employers of large firms are counted as employees, in which case the comparison of employee versus self-employed would partially be confounded by a firm size comparison. In addition, I can systematically distinguish between employees and both family and non-family workers in household enterprises. I can also systematically distinguish between employees that work for a salary versus employees that work for in-kind payments. Finally, and related to the previous point, I can distinguish between casual daily wage laborers and 'regular' employees in the countries where seasonal work is arguably most prevalent. It is true, however, that I cannot systematically distinguish casual wage laborers, and non-regular wage earners more generally, from contract-based regular employees. Taken together, this discussion implies that, with the exception of daily wage laborers, I can construct employee and self-employment categories in a consistent and internationally comparable manner across all countries. The second advantageous feature of all surveys is that my definition of employee versus self-employed is systematically based on an 'objective' worker-class question. In contrast, certain surveys allow respondents to choose 'informal sector' in response to the worker-type question. As discussed in the main text, my employee classification is closely related to the ILO concept of formality. Nonetheless, the specific definition of formality embedded in surveys is likely to vary across countries in ways that are hard to measure, and relying on such responses would reduce the transparency of comparisons across countries. As such, I discard all surveys where I cannot construct the employee classification based on an detailed and objective 'class of worker' question.

I focus on calculating gross income from all sources in order to be conceptually consistent with the

broadest possible income-definition in the tax code. This leads me to calculate four sources of income: wage income, self-employment income, capital income, and miscellaneous income (such as lottery receipts). Most importantly, I ensure that I can calculate both employee and self-employment income with precision. In this context, the most significant challenge is to calculate self-employment income in agriculture in less-developed countries. Agricultural earned revenue includes the value of crops sold to others. I do not attempt to create a monetary value of in-kind sales, as offering and receipt of in-kind goods and services is not subject to tax. Agricultural capital revenue includes the sale of livestock, income from rental of equipment, and share-cropping income. From this revenue I attempt to subtract costs, which include expenditure on inputs, wages paid out to workers, and new investments. In a limited number of countries, I do not observe any agricultural revenue for respondents that are self-employed in agriculture. These are most often contributing family workers on farms where the full output is consumed by the family. In this limited number of cases, I construct the income as the market value of the own-consumed output, as estimated by the respondent. In all surveys, I exclude two sources: social transfers, and in-kind goods and services. I exclude social transfers because it falls outside the concept of taxable income. The monetary value of in-kind goods and services are sometimes included in taxable income, often on a presumptive basis. However, apart from the mentioned case above, I exclude this source of income because I cannot measure it consistently across all surveys. Non-monetary income is often more important for less wealthy individuals, and is more prevalent in less developed countries. In the surveys where there exists systematic data on the monetary value of non-monetary income, I can confirm that the inclusion of these sources of income does not change the distributional employee-profile. That is because these sources of non-monetary income are too small in magnitude to overturn the decile-ranking of individual income.

In 7 countries, I cannot calculate gross individual income with precision. These countries, also reported in Table A.2, are: Democratic Republic of Congo, Liberia, Ethiopia, Malawi, Mali, Burkina Faso, Cambodia. In the case of DRC, Liberia, Ethiopia, and Malawi, I do not comprehensively observe either agriculture sales or costs, so I cannot calculate agricultural self-employment income. In Mali, Burkina Faso, and Cambodia, I do not comprehensively observe costs of non-agriculture own-account workers, so I cannot calculate non-agricultural self-employment income. In these 7 cases, which are among the poorest in the micro-database, I instead calculate total individual expenditure, and use it as a proxy for total income. There exists a set of low-income countries in which I have both good income and expenditure data. In results not reported, I can confirm that the employee-share profiles are very

similar when using either income or expenditure to calculate the x-axis distribution. As mapping expenditure into income is difficult, I do not attempt to locate the income tax exemption threshold in these 7 countries.

While I define the employee-status based on the respondent's primary job activity, I attempt to calculate income from all activities reported during the reference period. The main issue that arises in this context is the allocation of income which is reported at the household, rather than individual, level. For sources of earned income that are not at the individual level, I assign equal portions of this income to each economically active member of the household that reports having undertaken this activity during the reference period. Per example, the value of sold crops will be distributed equally among all household members that report having contributed to the family farm, either as a first or secondary activity. For sources of non-earned income reported at the household level, I assign an equal portion to each economically active member, such as in the case of property rental income.

Whenever a country's tax code is based on annual amounts and the reference period in the country's survey module is not, I construct the annual income distribution. I multiply the regular amount by the number of periods in the year – e.g. if wage income was reported monthly, I multiply it by the number of months that the wage income is reported to have been received during the past year. In the case where no periodicity exists, I assume that the flow was occurring during the whole year with the same pattern as during the reference period.

In every country survey, I limit the sample to the economically active population, following the definition of employment from the U.N. System of National Accounts. This definition is also used in Bicks, Fuchs-Schundeln, & Lagakos (2018), and in Feng, Lagakos, & Rauch (2018), which study respectively how hours worked and unemployment vary with development. I code employment-type based on the primary job in the reference period. The primary job is often explicitly defined as the job in which the respondent spent most hours during the reference period. The reference period in the Luxembourg Income Study (LIS) is annual, while it is predominantly monthly in the remaining surveys. The extent to which the periodicity and the focus on the primary job introduce biases in the representativeness of my employment-categories is discussed in Section A.6.

In addition to income and employment-categories, the micro-database also contains variables on education, sector, and geographical location. The geographical location measures whether a respondent lives in an urban area or not. I do not attempt to harmonize this variable, and use the urban definition in the surveys, which may therefore vary from country to country. I use variables to indi-

cate three levels of education completion: not completed primary; completed primary but not high school; completed high school. I chose to not distinguish further levels of education, in order to maximize the number of surveys where I could create consistent measures. Finally, I code the sector of the primary job. The aim was to create a set of sectoral categories which are consistent with the ISIC classification. I create four sectoral categories: agriculture; manufacturing ; services; and, public administration. I define these four categories in relation to the divisions of the ISIC 4.4 classification, where: agriculture contains Section A; manufacturing and construction contains Sections B to F; services contains Sections G to M, and S to U; and, public administration and education contains Sections N to R. As such, the manufacturing sector also contains mining and construction; the services sector also contains wholesale and retail trade, transportation, IT, finance, and activities of household enterprises; and, public administration also contains education, social work, and entertainment. Most of the industry codes in the surveys do not contain a pre-existing ISIC classification. To the best extent possible, I therefore first map the survey-categories to ISIC divisions, and then to my 4 sectoral categories. I do not include the sectoral variable if the survey has data only on a subset of the categories - per example, if a survey records that a job is not in the agricultural sector, but does not specify which non-agricultural sector it belongs to. The availability of the geography, education, and sector variables across surveys is described in Table A.2. These variables are used in the regression analysis in Section A.5.

A.1.2 Historical US time-series: sources and methodology

The historical federal profiles in the US between 1950 and 2010 were constructed using the decennial Census samples, extracted from the IPUMS USA database. I exclude all respondents that are not active in the labor force during the reference period. I calculate the individual income distribution, based on the measure of gross income at the individual level. To construct the income distribution, I use the measure of total, pre-tax, personal income. Farm and non-farm business income, as well as wage income, are consistently recorded in every Census sample. I use the detailed 'class of worker' question, which allows me to assign unpaid family workers to the self-employed category. Consequently, the self-employed category includes employers, own account workers, self employed that are not incorporated, and self-employed that are incorporated. Given the resemblance with the categories contained in the cross-country surveys, there is strong comparability between these US historical pro-

files and the cross-country profiles constructed in Figure 3. I apply individual weights to estimate the employee-share of every decile of the income distribution in every decade.

Before 1950, the decennial Census does not report total personal income at the individual level. The 1940 1 percent sample does contain wage and salary income, but no business income nor farm income, which are required to construct a personal gross income distribution. Instead, I use the 1935-36 Study of Consumer Purchases. The scope of the study was to "ascertain for the first time in a single national survey the earning and spending habits of inhabitants of large and small cities, villages, and farms" (ICPSR Study 8908, 2009). The survey was the result of a joint effort by the Bureau of Labor Statistics and the Bureau of Home Economics of the Department of Agriculture, and is meant to have been the sampling-methodology predecessor for the income-component in Census. The survey contains both a labor force component, where respondents gave information on income and housing, and for a subset of the total sample, a living conditions component where respondents gave additional information on expenditure. The primary sampling units were chosen to represent "the demographic, regional, and economic characteristics of the United States" (ICPSR, 2009). From these areas, a randomly selected group of approximately 700,000 families were screened in a first wave. From this first wave, 300,000 families were chosen to supply basic income and housing info, and a subset of 61,000 families were selected to provide additional expenditure information. It is important to understand the selection criteria into the different waves. The ICPSR accompanying documentation explains that in order to be selected out of the first wave, the requirements were: "families include at least two members, with husband and wife married for at least one year, and with no more than the equivalent of ten boarders for the survey year (...) farm families had to live in a setting that met the Census definition of a farm; the family itself must operate the farm (or in the southeast, be a sharecropper) and have conducted farming activities for at least one year" (ICPSR Codebook, 2009). Families were admitted to the first wave "without restriction in terms of occupation, income, employment status, or whether they were drawing or had drawn relief during the year." Selection into the second-wave where the survey included expenditure components, was based on the following criteria: "non-farm families must have had at least one wage earner in a clerical, professional, or business occupation. A minimum income for the survey year of $500 was required in the largest cities and $250 in the smaller cities and rural areas (...) Families that had received relief were excluded from this third wave." These criteria produce a highly selected sample for the second-wave respondents, and hence I base the analysis on the sample of first-wave respondents.

The ICPRS data-sample that I use for the 1935 Federal profile is based on a random sub-sample of approximately 5,000 families who only completed the first-wave 'labor force' component of the survey.[1] The ICPSR sub-sample was created in the following way: "a sampling fraction of 1 schedule for entry for every 83 schedules counted was chosen" from the urban sample, creating 3200 schedules from the larger urban areas and 1800 schedules from the more rural areas"; the ICPSR sample consists of schedules "spread across both the rural and urban portions of the original investigation." The employee classification is based on 'status of employment' question, which is identical to the (non-detailed) 'class of worker' question used in all US Censuses from 1950 onward. I code as an employee any individual respondent who reports being a "salaried worker/wage earner." I code as self-employed any respondent who reports being "self-employed", and any respondent who does not specify a type of work but declares to be working, is above age 20 and who has substantial work-related income. I exclude all respondents that are employed on work-relief projects in their primary job. As such, the sample closely resembles the economically active workforce definition used in the cross-country sample. Total gross income only exists at the household level. Rather than try to assign income at the individual level within the household, I focus on the work-type of the head of household. I then rank individuals based on the reported total income, and estimate the employee-share in each income decile.

The 1935-36 survey marked a clear shift in focus of the surveys conducted by the Bureau of Labor Statistics. Indeed, the surveys carried out prior to the 1930s focused on measuring family income and expenditure patterns of the U.S. employed workers and their families. Consequently, the available surveys, including the "Cost of living in the United States, 1917-1919" (ICPRS 7711, 1986) and the "Cost of living of industrial workers in the United States and Europe, 1888-1890" (Haines, 2006) contain data from families of wage earners or salaried workers in industrial locales scattered throughout the U.S. In order to construct a historical profile before the 1930s, I use data from Lindert & Williamson (2016), which studies incomes in the U.S. between 1650 and 1870.

Unlike previous work which approaches the measurement of income during this historical period from the production-side or the expenditure-side, Lindert & Williamson build estimates of income based on personal income records, assembling nominal earnings from free labor and property income. The approach to estimating income in Lindert & Williamson derives from combining informa-

[1] The ICPSR data available from the 1935-36 survey has also been used in Collins & Wanamaker (2014), Costa (2001), Margo (1993).

tion about income and labor force participation counts across occupation-space-time. This amounts to building 'social tables' across occupations within a given space-time frame, and the approach is conceptually similar to social accounting matrices that were used in development economics in the 1970s and 1980s. The authors provide a significant effort to capture all occupation categories in a given space-time. They draw on data from local tax assessments and occupational directories for 'registered' occupations, and local censuses for 'unregistered occupations'. These same data sources usually provide counts of the total number of individuals across the different occupations. The authors combine previous work with new estimates from local sources to derive personal earned income across occupation-space-time. In some instances, the occupation-space-time income reported was not at the annual level, and the authors bring the estimates to such level by making assumptions on the full-time number of hours spent (the assumptions are discussed in Lindert & Williamson, 2016). The authors also collect data on property income by assuming rates of return on wealth estimates that vary across occupation-space-time, and combine this with earned income to derive measures of total income.

I construct a historical 1870 profile based on the data kindly provided by Peter Lindert. This cross-section builds upon the 1870 1 percent US Census sample delivered to the authors by IPUMS USA, which included sampling weights at the individual-level. The 1 percent sample contains space-occupation counts, which are then merged with the authors' estimate of total income at the same level. I extend their analysis and classify all available occupation categories as either self-employed or employee. I use the detailed description of each occupation category to code employment-type. Per example, all occupations where a reference is made to 'manager' are coded as employee cells. The enumerator instructions for the sample design are particularly useful for my exercise in that they highlight very clearly the need to distinguish between self-employed and employee status: "Do not call a man a 'shoemaker', 'bootmaker', unless he makes the entire boot or shoe in a small shop. If he works in a boot and shoe factory, say so (...) Cooks, waiters, etc., in hotels and restaurants will be reported separately from domestic servants." The occupation category only exists for the head of household. The measure of total income includes own labor earnings in agriculture and non-agriculture, farm and non-farm operating income, and property income. This is a comprehensive measure of gross income before taxes and transfers which is not identical to, but closely resembles, the measure used in the more recent Federal US and cross-country samples. I apply the sampling weights initially provided by IPUMS USA. I estimate the employee-share in every decile of the individual gross income

distribution, for the population that is active in the labor force.

In all the profiles, I locate the Federal income tax exemption threshold in the income distribution. Note that there was no Federal income tax in 1870. In all profiles from 1935 onward, I use the historical IRS series which provide details on the nominal value of the standard deduction of a single filer.

Table A.1: Cross-Country Data Sources

Country	Year	Per Capita Income Group	Survey type	Coverage	Sample Size	Original source
Albania	2009	Upper Middle	Labor Force	National	18,997	National Institute of Statistics
Argentina	2009	Upper Middle	Living Conditions	Urban	47,862	National Institute of Statistics and Census
Australia	2014	High	Living Conditions	National	16,801	Luxembourg Income Study
Austria	2013	High	Living Conditions	National	5,102	Luxembourg Income Study (LIS)
Azerbaijan	1995	Low	Living Conditions	National	8,901	Living Standards Measurement Study (LSMS)
Bangladesh	2010	Low	Living Conditions	National	19,664	Bangladesh Bureau of Statistics
Belgium	2000	High	Living Conditions	National	2823	Luxembourg Income Study (LIS)
Belize	1999	Lower Middle	Labor Force	National	15,167	Central Statistical Office
Bolivia	2007	Lower Middle	Living Conditions	National	16,130	National Institute of Statistics
Brazil	2009	Upper Middle	Living Conditions	National	191, 810	National Institute of Geographics and Statistics
Bulgaria	2007	Upper Middle	Living Conditions	National	6,941	National Institute of Statistics
Burkina Faso	2014	Low	Living Conditions	National	32,023	National Institute of Statistics and Demographics
Cambodia	2009	Low	Living Conditions	National	31,959	Ministry of Planning
Cameroon	2007	Lower Middle	Living Conditions	National	51,836	National Institute of Statistics
Canada	2013	High	Living Conditions	National	27,344	Luxembourg Income Study (LIS)
Chile	2009	Upper Middle	Living Conditions	National	90,610	Social Observatory, University Alberto Hurado
China	2013	Upper Middle	Living Conditions	National	14,782	Luxembourg Income Study (LIS)
Colombia	2009	Upper Middle	Living Conditions	National	170,220	National Directory of Statistics
Costa Rica	2009	Upper Middle	Living Conditions	National	19,594	National Institute of Statistics and Census
Czech Republic	2013	High	Living Conditions	National	7,653	Luxembourg Income Study (LIS)
Cote d'Ivoire	2008	Lower Middle	Living Conditions	National	59,699	National Institute of Statistics
Dem. Rep. of the Congo	2004	Low	Living Conditions	National	72,685	National Institute of Statistics
Denmark	2013	High	Living Conditions	National	88,696	Luxembourg Income Study (LIS)
Dominican Republic	2009	Upper Middle	Living Conditions	National	30,430	National Statistics Office
Ecuador	2009	Lower Middle	Living Conditions	National	78,865	National Institute of Staistics and Censuses
Egypt	2010	Lower Middle	Living Conditions	National	34,069	Economic Research Forum (ERF)
El Salvador	2014	Lower Middle	Living Conditions	National	20,361	Center for Labor and Social Studies (CEDLAS)
Estonia	2013	High	Living Conditions	National	6,576	Luxembourg Income Study (LIS)
Ethiopia	2010	Low	Living Conditions	National	18,864	Living Standards Measurement Study (LSMS)
Finland	2013	High	Living Conditions	National	11,112	Luxembourg Income Study (LIS)
France	2010	High	Living Conditions	National	14,440	Luxembourg Income Study (LIS)

Table A.1: Cross-Country Data Sources (continued)

Country	Year	Per Capita Income Group	Survey type	Coverage	Sample Size	Original source
Georgia	2010	Lower Middle	Living Conditions	National	4,811	Luxembourg Income Study (LIS)
Germany	2014	High	Living Conditions	National	14,915	Luxembourg Income Study (LIS)
Ghana	2010	Low	Living Conditions	National	62,042	Ghana Statistical Service
Greece	2013	High	Living Conditions	National	6,115	Luxembourg Income Study (LIS)
Guatemala	2014	Lower Middle	Living Conditions	National	22,118	Luxembourg Income Study (LIS)
Honduras	2009	Lower Middle	Living Conditions	National	98,028	National Institute of Statistics
Hungary	2014	High	Living Conditions	National	2,718	Luxembourg Income Study (LIS)
Iceland	2010	High	Living Conditions	National	4,133	Luxembourg Income Study (LIS)
India	2004	Low	Living Conditions	National	59,487	Luxembourg Income Study (LIS)
Indonesia	2011	Lower Middle	Living Conditions	National	111,824	Statistics Indonesia
Iraq	2011	Lower Middle	Living Conditions	National	176,042	Economic Research Forum (ERF)
Ireland	2010	High	Living Conditions	National	3,508	Luxembourg Income Study (LIS)
Israel	2014	High	Living Conditions	National	11,770	Luxembourg Income Study (LIS)
Italy	2014	High	Living Conditions	National	6,258	Luxembourg Income Study (LIS)
Jamaica	2002	Lower Middle	Living Conditions	National	18,943	Living Standards Measurement Study (LSMS)
Japan	2008	High	Living Conditions	National	7,840	Luxembourg Income Study (LIS)
Jordan	2010	Upper Middle	Living Conditions	National	15,472	Economic Research Forum (ERF)
Kenya	2005	Low	Living Conditions	National	62,175	National Bureau of Statistics
Kosovo	2000	Lower Middle	Living Conditions	National	14,167	Living Standards Measurement Survey (LSMS)
Liberia	2014	Low	Living Conditions	National	18,089	Institute for Statistics
Lithuania	2008	Upper Middle	Living Conditions	National	15,837	National Statistics Office
Luxembourg	2013	High	Living Conditions	National	4,373	Luxembourg Income Study (LIS)
Malawi	2011	Low	Living Conditions	National	56,218	National Statistical Office
Mali	2014	Low	Living Conditions	National	37,175	Living Standards Measurement Study
Mexico	2011	Upper Middle	Living Conditions	National	17,682	National Institute of Statistics and Geography
Mongolia	2003	Low	Labor Force	National	49,948	National Statistical Office
Morocco	2009	Lower Middle	Living Conditions	National	10,769	Ministry of Economy and General Affairs
Mozambique	2014	Low	Living Conditions	National	9,128	National Institute of Statistics
Namibia	2009	Upper Middle	Living Conditions	National	44,614	National Planning Commission
Netherlands	2013	High	Living Conditions	National	23,935	Luxembourg Income Study (LIS)
Nicaragua	2014	Lower Middle	Living Conditions	National	9,250	Center for Labor and Social Studies (CEDLAS)
Niger	2011	Low	Living Conditions	National	3,859	Living Standards Measurement Survey (LSMS)
Nigeria	2011	Lower Middle	Living Conditions	National	23,289	National Bureau of Statistics
Norway	2013	High	Living Conditions	National	23,993	Luxembourg Income Study (LIS)

Table A.1: Cross-Country Data Sources (end)

Country	Year	Per Capita Income Group	Survey type	Coverage	Sample Size	Original Source
Pakistan	2001	Lower Middle	Living Conditions	National	75,519	Federal Bureau of Statistics
Palestine	2011	Lower Middle	Living Conditions	National	25,947	Economic Research Forum (ERF)
Panama	2010	Upper Middle	Population and Housing Census	National	314,118	IPUMS-International
Papua New Guinea	1996	Lower Middle	Living Conditions	National	8,660	Living Standards Measurement Survey
Paraguay	2009	Lower Middle	Living Conditions	National	18,419	National Statistics Office
Peru	2009	Upper Middle	Living Conditions	National	95,199	National Institute of Statistics
Poland	2013	High	Living Conditions	National	39,993	Luxembourg Income Study (LIS)
Puerto Rico	2005	High	Population and Housing Census	National	35,416	IPUMS-International
Romania	1997	Lower Middle	Living Conditions	National	35,995	Luxembourg Income Study (LIS)
Russia	2013	High	Living Conditions	National	6,079	Luxembourg Income Study (LIS)
Rwanda	2000	Low	Living Conditions	National	32,679	National Institute of Statistics
Serbia	2007	Upper Middle	Living Conditions	National	17,375	Living Standards Measurement Survey (LSMS)
Sierra Leone	2003	Low	Living Conditions	National	23,022	National Office of Statistics
Slovakia	2009	High	Living Conditions	National	4,704	National Statistical Office
South Africa	2012	Upper Middle	Living Conditions	National	7,105	Luxembourg Income Study (LIS)
South Korea	2006	High	Living Conditions	National	13,178	Luxembourg Income Study (LIS)
Spain	2013	High	Living Conditions	National	10,728	Luxembourg Income Study (LIS)
Sri Lanka	2008	Lower Middle	Labor Force	National	66,381	Department of Census and Statistics
Sudan	2009	Lower Middle	Living Conditions	National	48,845	Economic Research Forum (ERF)
Sweden	2005	High	Living Conditions	National	11,607	Luxembourg Income Study (LIS)
Switzerland	2013	High	Living Conditions	National	7,961	Luxembourg Income Study (LIS)
Taiwan	2013	High	Living Conditions	National	23,474	Luxembourg Income Study (LIS)
Tajikistan	2007	Low	Living Conditions	National	1,503	State Statistical Agency
Timor Leste	2007	Lower Middle	Living Conditions	National	9,094	National Statistics Directorate
Tunisia	2009	Upper Middle	Living Conditions	National	50,371	Economic Research Forum (ERF)
Turkey	2011	Upper Middle	Labor Force	National	37,121	National Statistical Institute
Tanzania	2010	Low	Living Conditions	National	20,559	National Bureau of Statistics
Uganda	2011	Low	Living Conditions	National	13,618	National Bureau of Statistics
Ukraine	2010	Lower Middle	Living Conditions	National	10,428	State Statistics Service
United Kingdom	2013	High	Living Conditions	National	20,002	Luxembourg Income Study (LIS)
United States	2013	High	Living Conditions	National	63,859	Luxembourg Income Study (LIS)
Uruguay	2009	Upper Middle	Living Conditions	National	132,559	National Institute of Statistics
Venezuela	2006	Upper Middle	Living Conditions	National	166,506	National Institute of Statistics
Zambia	2014	Lower Middle	Living Conditions	National	11,921	Central Statistical Office

Notes: for details on this table, please see Section A.1.

Table A.2: Cross-Country Data Variable Availability

Country	Year	Income	Sector	Education	Location
Albania	2009	x	x	x	
Argentina	2009	x	x	x	x
Australia	2014	x	x	x	
Austria	2013	x		x	x
Azerbaijan	1995	x	x		x
Bangladesh	2010	x	x	x	x
Belgium	2000	x	x	x	x
Belize	1999	x	x	x	x
Bolivia	2007	x	x	x	x
Brazil	2009	x	x	x	x
Bulgaria	2007	x		x	x
Burkina Faso	2014		x	x	x
Cambodia	2009			x	x
Cameroon	2007	x	x	x	x
Canada	2013	x		x	x
Chile	2009	x	x	x	x
China	2013	x	x	x	x
Colombia	2009	x	x	x	x
Costa Rica	2009	x	x	x	x
Czech Republic	2013	x	x	x	x
Cote d'Ivoire	2008	x	x	x	x
Dem. Rep. of the Congo	2004			x	x
Denmark	2013	x	x	x	x
Dominican Republic	2009	x	x	x	x
Ecuador	2009	x	x	x	x
Egypt	2010	x	x	x	x
El Salvador	2014	x	x		
Estonia	2013	x	x	x	
Ethiopia	2010			x	x
Finland	2013	x	x	x	x
France	2010	x	x	x	x

Table A.2: Cross-Country Data Variable Availability (continued)

Country	Year	Income	Sector	Education	Location
Georgia	2010	x	x	x	x
Germany	2014	x	x	x	x
Ghana	2010	x		x	x
Greece	2013	x	x	x	x
Guatemala	2014	x	x	x	x
Honduras	2009	x	x	x	x
Hungary	2014	x	x	x	x
Iceland	2010	x	x	x	x
India	2004	x	x	x	x
Indonesia	2011	x	x	x	x
Iraq	2011	x	x	x	x
Ireland	2010	x	x	x	x
Israel	2014	x	x	x	x
Italy	2014	x	x	x	x
Jamaica	2002	x	x	x	x
Japan	2008	x	x	x	x
Jordan	2010	x	x	x	x
Kenya	2005	x	x	x	x
Kosovo	2000	x	x	x	x
Liberia	2014			x	x
Lithuania	2008	x		x	x
Luxembourg	2013	x	x	x	x
Malawi	2011		x	x	x
Mali	2014		x	x	x
Mexico	2011	x	x	x	x
Mongolia	2003	x	x	x	x
Morocco	2009	x	x	x	x
Mozambique	2014	x	x	x	x
Namibia	2009	x	x	x	x
Netherlands	2013	x		x	x
Nicaragua	2014	x	x		
Niger	2011	x	x	x	x
Nigeria	2011	x	x	x	x
Norway	2013	x			x

Table A.2: Cross-Country Data Variable Availability (end)

Country	Year	Income	Sector	Education	Location
Pakistan	2001	x	x	x	x
Palestine	2011	x	x	x	x
Panama	2010	x	x		x
Papua New Guinea	1996	x	x	x	
Paraguay	2009	x	x	x	x
Peru	2009	x	x	x	x
Poland	2013	x	x	x	x
Puerto Rico	2005	x		x	x
Romania	1997	x	x	x	x
Russia	2013	x	x	x	x
Rwanda	2000	x	x	x	x
Serbia	2007	x	x	x	x
Sierra Leone	2003	x	x	x	x
Slovakia	2009	x		x	x
South Africa	2012	x	x	x	x
South Korea	2006	x			x
Spain	2013	x	x	x	x
Sri Lanka	2008	x		x	x
Sudan	2009	x	x	x	x
Sweden	2005	x	x	x	
Switzerland	2013	x	x	x	
Taiwan	2013	x	x	x	
Tajikistan	2007	x	x		x
Timor Leste	2007	x	x	x	x
Tunisia	2009	x	x	x	x
Turkey	2011	x		x	x
Tanzania	2010	x	x	x	x
Uganda	2011	x	x	x	x
Ukraine	2010	x	x		
United Kingdom	2013	x	x	x	
United States	2013	x	x	x	
Uruguay	2009	x	x	x	x
Venezuela	2006	x	x	x	x
Zambia	2014	x	x	x	x

Notes: for details on this table, please see Section A.1.

A.2 Additional historical profile: Mexico 1960-2010

As a robustness check to the stylized facts, I show that they also hold over the long-run in a currently developing country, Mexico. I focus on Mexico because it has variables of income and employee-jobs that are consistently defined over a long period of time, namely 1960-2010. The data is extracted from IPUMS International. The disadvantage is that only earned income is measured consistently over this period - as opposed to total income, which further includes capital income and 'other' income. I use answers to the 'class of worker' question. The only inconsistency over time in this question is that the 2010 sample groups household assistants together with salaried workers, whereas in previous samples, these categories are separated. As such, I am over-estimating the true employee-share in the 2010 profile. Importantly, day laborers are separated from salaried workers, and I can assign the former to the self-employment category in all years. There also exists a category for unpaid family workers, which I assign to the self-employment category. I construct the sample of respondents that are economically active, and use survey weights to construct individual earned income distributions in 1960, 1990, and 2010. For the years 1990 and 2010, I code the value of the exemption threshold from OECD's Personal Taxes database. For 1960, I use the historical archives of the Mexican Tax Authority.[2]

The results from this exercise are displayed in Figure A.1. I uncover the same stylized facts that were found both in the cross-country sample and in the historical US series: the employee-share profile is upward-sloping and gradually moves leftward in the gross income distribution; the exemption threshold gradually moves down the distribution and expands the size of the income tax base; and, the employee-composition on the tax base is constantly maximized.

[2] Available at: http://www.dof.gob.mx/index.php.

Figure A.1: Mexico 1960-2010

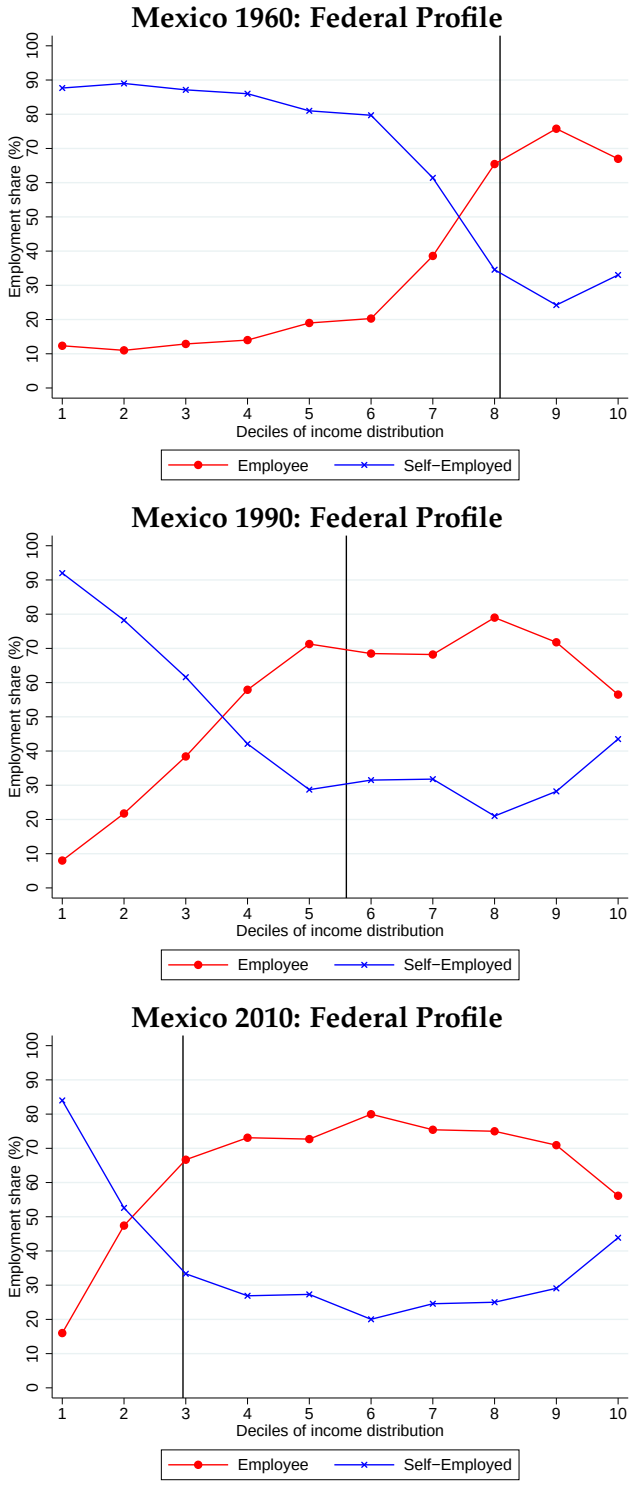

Notes: The circle-line (small cross-line) indicate the employee-share of the economically active workforce in a decile of the Mexican personal gross income distribution. An employee-job is defined as a job whose activity generates an information trail that can be leveraged for income tax enforcement purposes. For more details on this variable, please see Section 3.2. In every profile, the vertical solid line denotes the location of the Federal individual income tax exemption threshold. This threshold is the nominal value of gross (pre-tax) income above which a single filer becomes liable to pay income tax. Each historical profile is built from the Census micro-data from IPUMS International. The values of the exemption thresholds are from the OECD's Personal Taxes database, and the official archives of the Mexican revenue service. Source: Section 3.4 and Appendix Section A.2.

A.3 Redistributive targeting

In this robustness check, I provide evidence to suggest that the exemption threshold is not set to target social assistance or anti-poverty in the income distribution. Governments define thresholds of income that are used as inputs in formulas to provide social assistance and anti-poverty relief. I use the national poverty line and the minimum wage values as proxies for the 'social redistribution' threshold. I first show that only very rarely is the income tax threshold explicitly defined to be either equal to, or a multiplicative of, this social redistribution threshold. In 5% of countries in the cross-sectional sample, the tax code defines the exemption threshold to be a multiple of this redistribution threshold. These countries are: Mozambique, Bolivia, Paraguay, Turkey, and Slovakia. As an example, in Mozambique the exemption threshold is equal to 36 times the minimum wage, while in Paraguay it is equal to 120 times the minimum wage. I use the country-specific IBFD tax summaries to document this pattern. There exists a much more frequent explicit relation between redistributive thresholds and social security contributions. Indeed, several countries use (a multiplicative of) the minimum wage to define an exemption threshold for employee contributions.

Even if there exists no explicit relation defined in the tax code, governments may nonetheless implicitly maintain an association between the tax threshold and the social assistance threshold. To investigate this, I collect data on the value of the national poverty line and the minimum wage in all countries in the cross-sectional sample. I try to collect the data in as close a year as possible to the survey and tax exemption threshold year. I use harmonized data from ILO on the statutory nominal gross monthly minimum wage. Data is missing in 8 countries: Austria, Denmark, Finland, France, Kosovo, Sweden, Switzerland, Palestine. There does not exist a similar harmonized database on the value of the national poverty line for my sample. The World Bank collects cross-country data on the share of the population that falls below both international and national poverty lines, but such data does not directly disclose the value of the national lines used. I was able to collect relevant data in 88 of the 100 countries in my sample. The missing countries are: Albania, Austria, Hungary, Kosovo, Panama, Papua New Guinea, Romania, Serbia, Slovakia, Timor Leste, Ukraine, and Venezuela. Importantly, I collect the poverty line that is set by the national government, rather than the value of the international poverty line in local currency. Some governments do incorporate international criteria to determine poverty lines. Per example, some low-income countries base their poverty calculations on the minimum nutritional intake concept used by the World Bank to define international poverty; and, some European countries adopt the EU-wide definition of poverty as 60% of median income. The

important point is that the poverty lines I collect are based on an active decision made by the government, similarly to the definition of the tax exemption threshold. In some countries, the government defines several poverty lines, per example on a regional basis or on an urban-rural basis. I always pick the poverty line in each country with the highest value. Since poverty lines in developing countries are most often below the tax exemption threshold, this decreases the likelihood to observe that the two thresholds are far away from each other in value.

The results are displayed in Figure A.2. The three panels separate countries into development groups: low and lower middle income; higher middle income; and, high income. I construct the ratio of the income tax exemption threshold to the minimum wage, and of the exemption threshold to the poverty line. In the left-hand graphs, the bars represent country-specific ratios using the minimum wage, while the right-hand graphs display the ratio using the poverty line. Finally, within each graph, I sort the countries by GDP per capita. I take the log of the ratio, as this allows me to display all country-ratios on the same graph. Therefore, a bar-value below 0 means that the exemption threshold is located below the minimum wage/poverty line in the specific country. There is no obvious, confounding trend which emerges from Figure A.2. Within all development groups, countries with similar per capita income, and hence similar size of tax base (Figure 4), display very large variation in the relative value of the tax threshold to the redistribution threshold (note the log-scale of the y-axis). This holds even for countries at similar levels of development within the same region: the ratio for the minimum wage (poverty) is 0.48 log points (1.90 log points) in Burkina Faso, while it is 3.07 log points (6.09 log points) in Uganda; it is 2.09 log points (2.79 log points) in Bolivia, and 0.51 log points (2.65 log points) in Honduras. The highest-income countries often locate both the poverty and the minimum wage thresholds above the tax exemption threshold. But apart from this feature, there is not any systematic relationship between the relative location of tax and redistribution thresholds, and per capita income. Taken together, these findings suggest that the tax exemption threshold is not set to target social assistance in the income distribution.

Figure A.2: Redistributive Targeting

Panel A: Low and Lower-Middle Income Countries

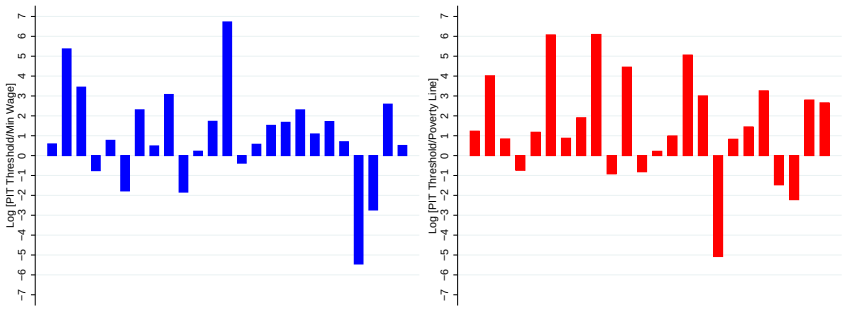

Panel A: Upper-Middle Income Countries

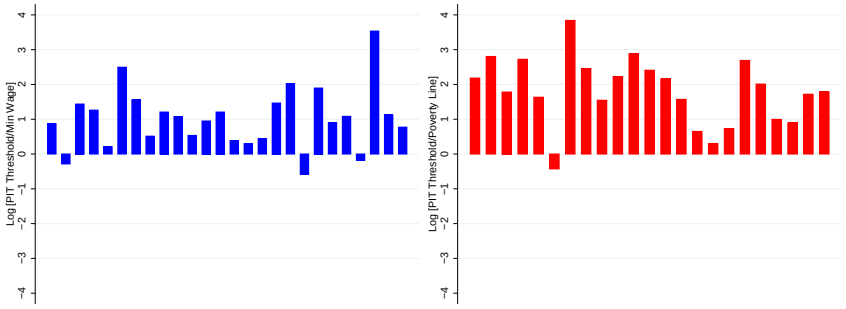

Panel C: High-Income Countries

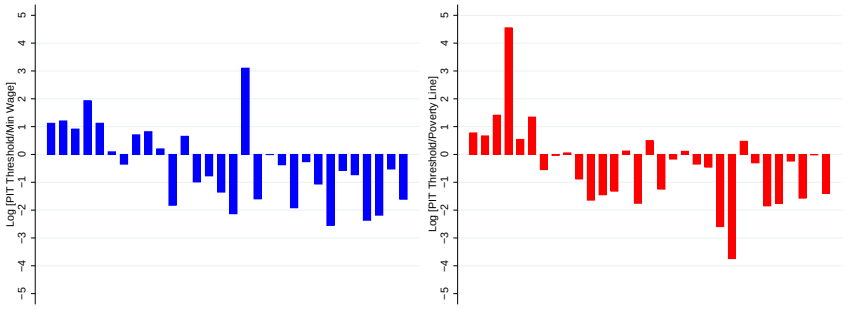

Notes: In every graph, a bar represents a country-observation from the cross-country micro-database. The three panels demark countries according to their per capita income group: low and lower-middle; upper-middle; high income. Within each graph, countries are ranked in ascending order of per capita income. Within each group, the left-hand graph shows the log of the ratio of the income tax exemption threshold to the minimum wage; the right-hand graph shows the log of the ratio of the income tax exemption threshold to the poverty line. All thresholds are expressed in annual and local currency. Source: Section 3.4 and Appendix Section A.3.

A.4 Sectoral distributional profiles

In this robustness check, I consider whether the location of the exemption threshold is targeting sectoral structure, rather than employment structure. I study whether the threshold appears to be set such as to avoid a 'hard to tax' sector, agriculture, or whether the threshold is set to capture the 'easy to tax' sectors of manufacturing and public administration (Musgrave, 1981). In order to investigate this confounding hypothesis, I first consider whether the tax exemption explicitly targets any sector. In particular, I use the IBFD country-reports in all countries in the cross-section, and report whenever income from agriculture is fully exempt from individual income taxation. I do not take into account instances where tax codes allow self-employed to deduct costs specifically related to agricultural work - per example, from the purchase of a tractor for farming. This is because my measure of the exemption threshold in all countries is the standard deduction, which is granted regardless of taxpayer behavior, and not the itemized deduction, which requires the taxpayer to itemize deductions. I chose the former measure because it can be constructed in a comparable way across space and time, as discussed in Section 3.2. I find that agricultural income is fully exempt only in 11% of low-income countries; 12% of middle-income countries; and, 5% of high-income countries. This list includes Mali, Morocco, and Sierra Leone in Africa; and, India and Pakistan in Asia.

As an alternative approach, I consider whether changes in sectoral distributional profiles over development could account for the movement in the exemption threshold. I create four sectoral categories in all the surveys in the cross-section: agriculture; manufacturing and construction; trade and services; and, public administration. I define these four categories in relation to the divisions of the ISIC 4.4 classification. The construction of the sector variable is described in detail in Section A.1.1.

Using these harmonized sector variables, I first study the distributional profiles of agricultural employment. I construct these profiles in the same way as the employment profiles in the main text (Section 3.2). The results are displayed in Figure A.3. At lowest levels of development, agriculture is prevalent everywhere except for the top of the income distribution. And, in the transition from low-income to middle-income group, the downward-sloping agriculture-profile gradually shifts leftward in the distribution. This pattern is similar to stylized fact #1, such that stylized fact #2 would be consistent with a setting where the exemption threshold targets the non-agricultural sector which increases gradually further down the income distribution. However, Figure A.3 also reveals that in these same income groups, virtually all agricultural work is concentrated among self-employed with no information trails. On the other hand, Figure A.3 reveals that in the transition from middle-income to the

high-income group, the agricultural profile has become very small in magnitude and almost entirely flat in the distribution. During this same transition, there continues to be an important transition between self-employment and employee-jobs outside of the agricultural sector, which is associated with further decreases in the location of the threshold. These facts suggest that movement out of agriculture could account for the expansion of the tax base, but only in a limited range of the development path, where it is fully confounded by movements out of self-employment. In contrast, movements out of self-employment can account for the expansion of the base over the full development path, including over a range of development where it cannot be confounded by movements out of agriculture.

I now consider whether the movement of the exemption threshold is consistent with targeting of 'easy to tax' sectors. I focus on manufacturing and public administration. Since work in these sectors is strongly correlated with having an employee-job, I study the sectoral profiles conditional on employee-job. Results are displayed in Figure A.4. The distributional profile of easily taxable sectors would have to be upward-sloping in the income distribution, and move leftward as the country develops, in order to be a confounding factor. This is not borne out in the observed profiles. The public administration profile is upward-sloping at some development levels, but the magnitude of the slope is quantitatively small, and there is no consistent left-ward shift over development. The public sector share at the top of the income distribution is most likely driven by central administration workers, while the share towards the lower end of the distribution is probably made up in part by field-workers in health and education. While located at very different parts of the income distribution, these jobs share the common feature of being easy to tax - in the sense that the government, as the direct employer, perfectly observes the salaries. The manufacturing distributional profile is largely flat in the income distribution. The level-shift upward and then downward of the manufacturing profile is consistent with the inverse-U shaped aggregate importance of manufacturing over development that other work has documented. Taken together, these facts do not suggest that the stylized facts #1-#4 are confounded by sectoral transitions over the development path.

Figure A.3: Distributional Profiles of 'Hard-to-Tax' Sectors

Profile for average country at $1065 pc [LHS] and $2226 pc [RHS]

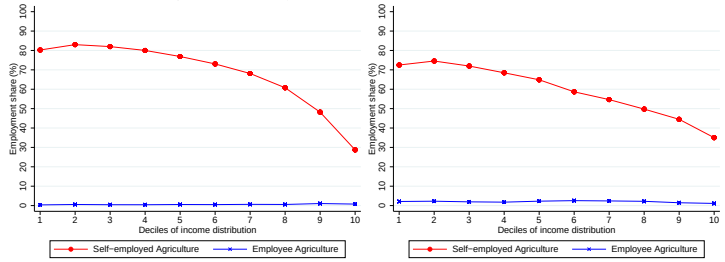

Profile for average country at $3239 pc [LHS] and $5796 pc [RHS]

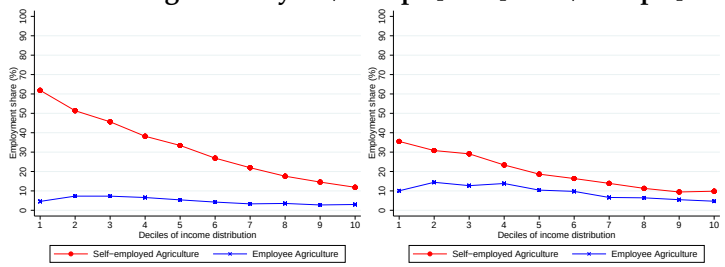

Profile for average country at $8826 pc [LHS] and $11257 pc [RHS]

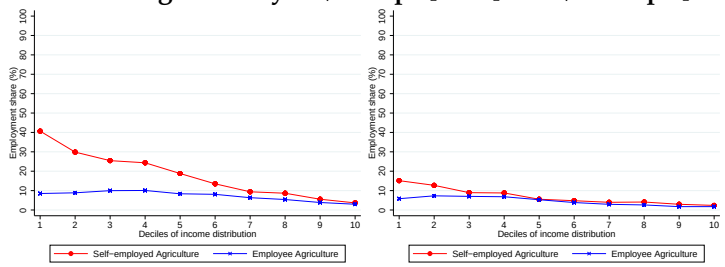

Profile for average country at $17141 pc [LHS] and $27960 pc [RHS]

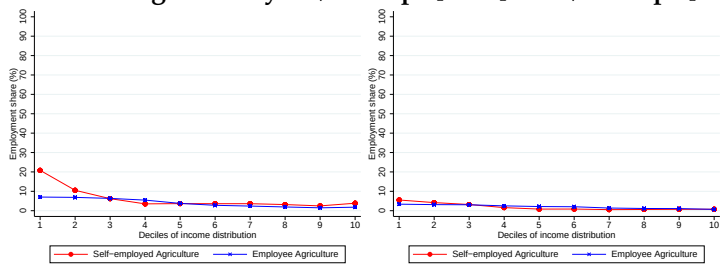

Profile for average country at $38224 pc [LHS] and $53878 pc [RHS]

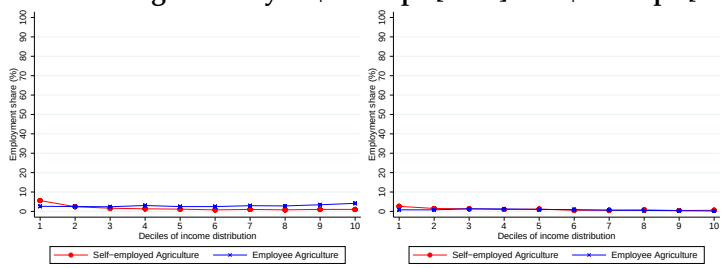

Notes: These figures plot the employment shares of self-employed agricultural workers and of employee agricultural workers, over deciles of the income distribution, for representative countries at different levels of per capita income. Employees (self-employed) are defined as individuals working in jobs which generate (no) information trails for the purposes of income tax enforcement. The share of each group is defined as the share of the total economically active workforce in the decile of the income distribution. To construct this graph, I partition the cross-country sample into ten groups of equal size, based on their level of per capita income. Note that I am limited to the group of countries where there exists sectoral data (see Table A.2). Within each group, I calculate the unweighted average employment-share of agricultural self-employed and agricultural employee. I plot this average profile for every group, and indicate the average per capita income of the group. I use expenditure-side real GDP at chained PPPs in 2011 US$ from the same year as the country-survey year. Source: Section 3.4 and Appendix Section A.4.

Figure A.4: Distributional Profiles of 'Easy-to-Tax' Sectors

Profile for average country at $1065 pc [LHS] and $2226 pc [RHS]

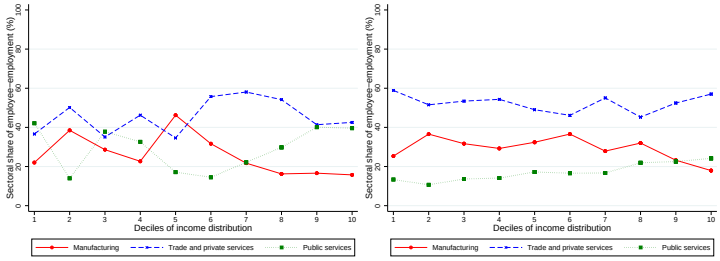

Profile for average country at $3239 pc [LHS] and $5796 pc [RHS]

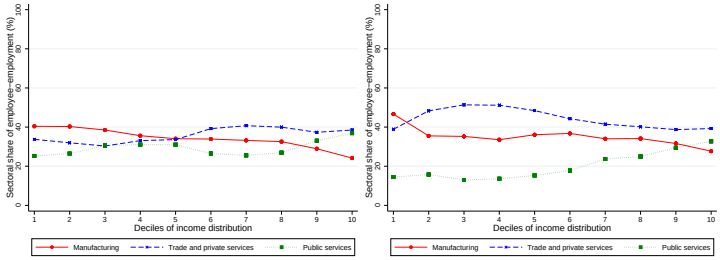

Profile for average country at $8826 pc [LHS] and $11257 pc [RHS]

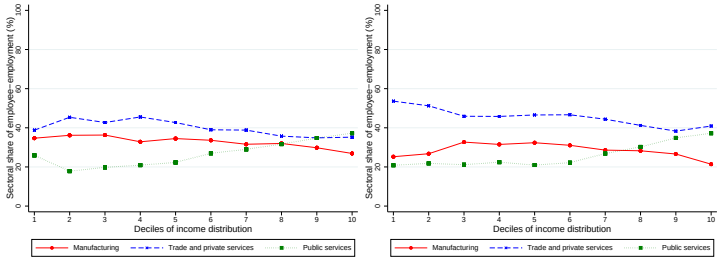

Profile for average country at $17141 pc [LHS] and $27960 pc [RHS]

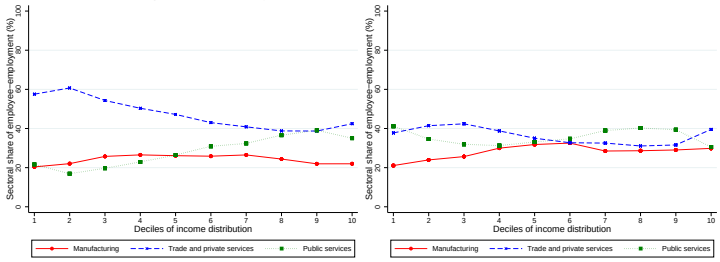

Profile for average country at $38224 pc [LHS] and $53878 pc [RHS]

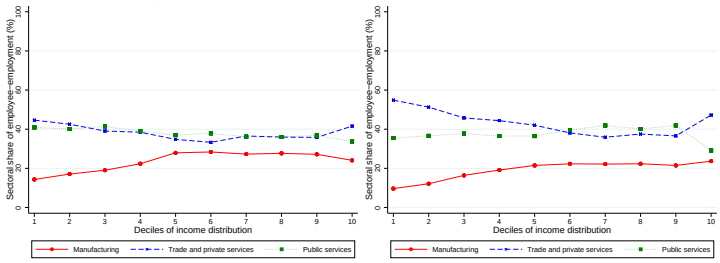

Notes: These figures plot the sectoral shares of employees over deciles of the income distribution, for representative countries at different levels of per capita income. Sectors are defined accordig to the ISIC classification (Section A.4). The share of each sector is defined as the share of the total employee workforce in the decile of the income distribution. To construct this graph, I partition the cross-country sample into ten groups of equal size, based on their level of per capita income. Note that I am limited to the group of countries where there exists sectoral data (see Table A.2). Within each group, I calculate the unweighted average sectoral shares by income decile. I plot this average profile for every group, and indicate the average per capita income of the group. I use expenditure-side real GDP at chained PPPs in 2011 US$ from the same year as the country-survey year. Source: Section 3.4 and Appendix Section A.4.

A.5 Robustness of employee-income gradient in regression setting

In this subsection, I investigate the employee-income gradient in a regression setting. This serves two purposes. First, it provides a complementary method to the distributional profiles approach, to study the robustness of the employee-income gradient. Second, it provides a more formal setting to study which characteristics partially contribute to the steepness of the observed slope. I focus on three characteristics: sector, location, and education. These are individual characteristics that the government could, albeit imperfectly, seek to target for redistributive purposes. As such, if controlling for one such characteristic eliminates the employee-income gradient, this could suggest that the threshold in fact targets this confounding characteristic. At the same time, these are observable characteristics which vary over development, including from the sectoral movement from agriculture to manufacture to services; the rural-urban migration; and, the rise in higher education. As such, the partial reduction in magnitude due to controlling for a particular characteristic would be informative of the importance of this characteristic in quantitatively explaining the change in employee-income gradient over development.

I use the four sectoral categories described in the Section A.1.1. I further create a dummy variable equal to 1 if a respondent lives in an urban area. I do not attempt to harmonize this variable, and use the urban definition directly in the surveys. Finally, I use education variables to code four dummies, indicating if a respondent has: not completed primary; completed primary but not high school; completed high school. I chose to not distinguish further levels of education, in order to maximize the number of surveys where i could create consistent measures. The availability of these different variables is described in Table A.2.

To visualize the impact of controlling for a characteristic on the employee-income gradient, I employ the methodology used in Bachas, Gadenne & Jensen (2019). In particular, in every country c, I estimate the following regression

$$\mathbf{1}(\text{Employee})_i = \alpha + \theta \boldsymbol{X}_i + \beta log(\text{income})_i + \varepsilon_i$$

where income_i is the individual gross income of individual i used to construct the income distribution (Section 3.2), $\mathbf{1}(\text{Employee})_i$ is a dummy equal to 1 if an individual is an employee (Section 3.2), and $\boldsymbol{X}_i$ contains the control indicator variables (sector, education, urban). I obtain a country-specific slope-coefficient β^c from estimating this regression separately in every country. In every graph, I

plot these coefficients $\beta^c_{with\,control}$ together with coefficients from estimating the regression without controls, $\beta^c_{no\,control}$, against log per capita income. The two coefficients for a particular country are denoted by the beginning ($\beta^c_{no\,control}$) and end ($\beta^c_{with\,control}$) of a vertical arrow. This regression is a linear probability model, which has the advantage that the slope-coefficient is directly interpretable. The disadvantage is that the slope-coefficient is not informative in settings where the relationship between employee and log(income) is strongly non-linear. This is the case in less (most) developed countries, where the likelihood of being an employee is very small (large) apart from the very top (top and bottom) of the income distribution (Figure 3). As an alternative, I can estimate the employee-share differential between the top and bottom deciles. This yields very similar qualitative results (not reported).

The results are displayed in Figure A.5. The top two panels control for geography (left graph) and education (right graph). The impact of geography is limited, but the inclusion of education significantly reduces the income-employee gradient especially in middle-income countries. The bottom left graph controls for sectors. This leads to the strongest reduction in magnitude, both in low-income and middle-income countries. It does not, however, fully eliminate the slope in most countries, and the potential confounding movement out of agriculture has been addressed in Appendix Section A.4. The bottom right graph includes all the control variables. This leads to a further reduction in slopes in most countries, compared to the sector control specification. This suggests that within sectors, location and, perhaps more likely, education, continues to be associated with higher income and employee-job status. Interestingly, the full set of controls almost fully eliminates the variation in the magnitude of employee-income gradient across development. This suggests that the joint movement over development of these three characteristics could drive the distributional employment patterns in stylized fact #1.

Figure A.5: Employee-Income Gradients Across Countries Without and With Controls

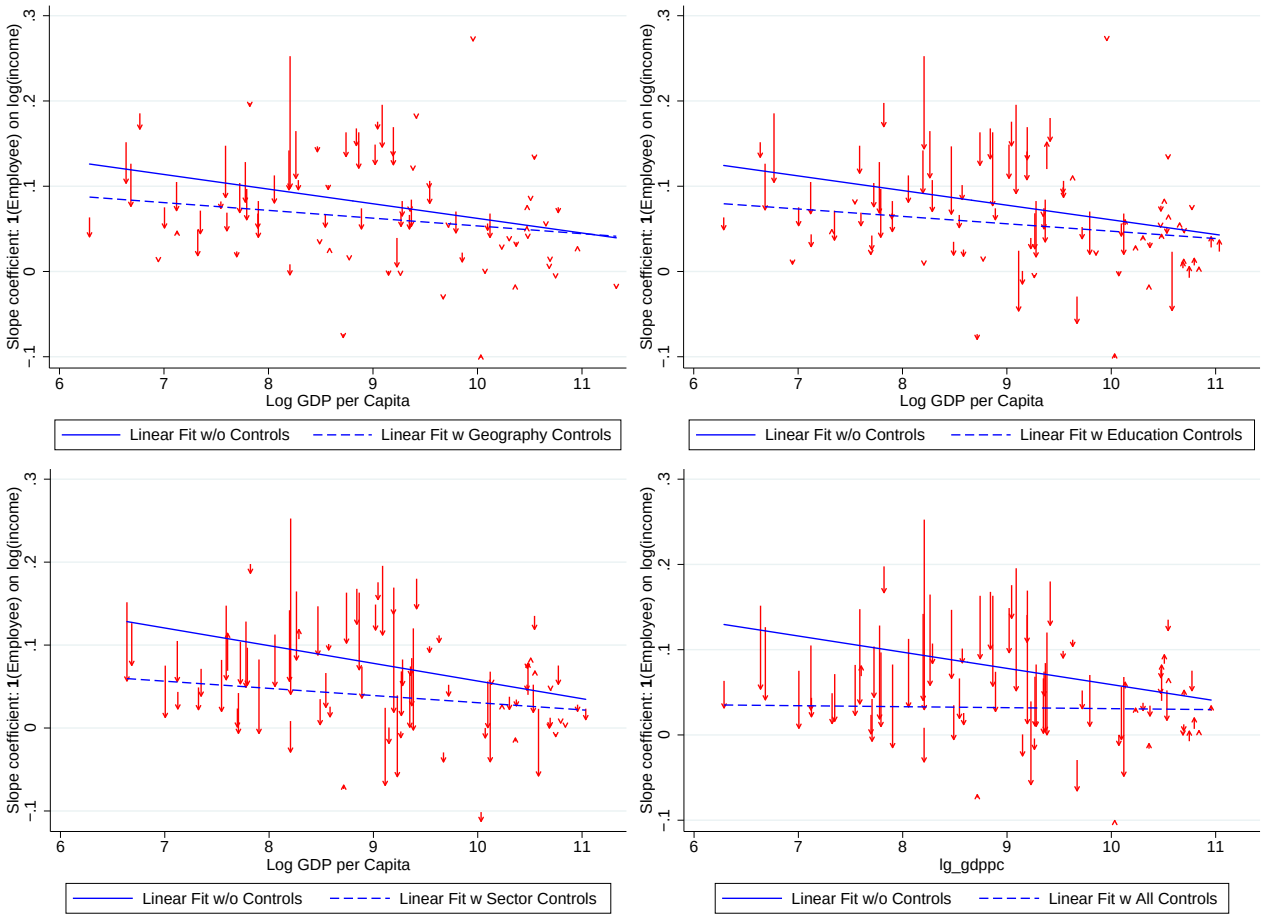

Notes: Each dot in every scatter-plot represents a country-specific slope coefficient based on estimating the regression in Section A.5. Each of the four graphs show slope-coefficients when including controls for: geography (North-West quadrant); education (NE); sectors (SW); geography, education, and sector (SE). In each graph, the start-point of an arrow represents the country-specific slope-coefficient without the control, and the end-point of an arrow represents the slope-coefficient after including the control. All slope-coefficients are plotted against log GDP per capita, measured using expenditure-side real GDP at chained PPPs in 2011 US$. In every graph, the solid (dashed) line represents the linear OLS fit of the slope-coefficients without control (with control). For more details on the construction of the different control variables, please see Section A.1.1. These graphs are constructed using the full cross-country survey sample. Source: Section 3.4 and Appendix Section A.5.

A.6 Potential biases resulting from methodology

In this subsection, I discuss the potential biases that can arise from the survey methodology and the measurement and construction of variables. I code employment type based on the primary job in which the respondent spent the most hours during the reference period. Many individuals have many jobs at the same time (Banerjee & Duflo, 2007). But this will this will affect the representativeness of my estimates only to the extent that these jobs fall in different categories in my classification. An individual who contributes on the family farm while being an own-account worker within the same reference period would be classified as 'self-employed' in both jobs. In surveys where the reference period is not yearly, there may be bias in the measure of employment structure if the employment type in the reference period is not representative of the entire year. This is potentially important in developing countries, where there is strong seasonality in job type. This introduces bias to the extent that the jobs at different periods of the year fall in different employment structure categories, which I argue is unlikely in a developing country context. Indeed, individuals that are casual wage laborers during the harvest season are unlikely to be regular full-time employees in the non-harvest season. Rather, they are likely to be own-account workers or contributing family workers. In this case, the individual would be classified as self-employed during all periods of the year, despite the different jobs held at different periods of the year.

A second source of bias comes from the fact that I cannot systematically separate casual wage work from contract-based wage-work. I can always distinguish between working for someone for pay versus for in-kind payment, and I exclude the latter from the employee category. As such, casual wage-laborers that receive in-kind payment are systematically classified as self-employed. This leaves the group of casual workers that are not paid in-kind as the group that I potentially mis-classify as employee, whenever the survey answers do not provide sufficient precision about the nature of the employee-work. Since the transition over development involves a movement out of casual wage labor into contract-based wage labor, this mis-classification will lead me to under-state the true growth in employee-share along the development path.

Another potential source of bias arises from the possibility that self-employed misreport their true amount of income. This is unlikely to introduce a major bias, for three reasons. First, unlike on tax returns, self-employed do not directly have any incentive to mis-report their income to surveyors. Second, the model in Section 5 does predict under-reporting of income among self-employed locally around the exemption threshold. But while the standard bunching model predicts a steep-

ened employee-share locally around the exemption threshold, it also predicts a decrease in employee-share further to the left of the threshold. This is not borne out in the data: instead, I observe a gradual increase in the employee-share over the full distribution. More generally, under-reporting of income by the self-employed would imply that the true self-employed distributional profile lies to the right of the observed one. If development is associated with increases in the ability to detect under-reporting among self-employed, this would generate gradual leftward shifts of the employee-share profile. Under-reporting of income by self-employed could also be due to by non-evasion motives. Woodruff et al. (2009) show that recall error, which is more present when the reference period is not annual, lead self-employed to under-estimate their income. If development is associated with a decrease in recall error, either due to changes in survey methodology or to an increase in accounting tools and book-keeping, this would similarly imply a rightward shift of the employee-profile at increasing levels of development. Both evasion detection capacity and measurement precision, which plausibly grow with development, therefore lead me to under-estimate the true progressive rightward shifts in the employee-profile due to structural transition out of self-employment.

Finally, bias could be introduced from the construction of the income tax base. I construct the tax base as the share of the individual income distribution that lies above the single-filer standard deduction (or allowance). As explained in the main text (Section 3.2), this choice is made to construct the tax base in the most transparent way without making any behavioral assumptions and in a way that can meaningfully be compared across countries. Notwithstanding, there exists features of tax systems which allow taxpayers to further reduce their tax liability, including deductible expenses. If a significant number of filers makes use of such additional features, this introduces a wedge between the size of base measured in this paper, and the size of the 'effective' base. The extent of existence of these features varies significantly across countries. Per example, there is a growing policy debate in the US on the large number of taxpayers that do not pay any Federal income tax. There exists no consistent evidence across countries at different levels of development on the extent to which the effective tax base is reduced through credits and deductions. Even if taxpayers in all countries in my sample made use of these deductions, it is likely that the size-wedge between my measured base and the effective base is larger in more developed countries. This is simply because the potential wedge in less-developed countries is bounded above by the small size of my measured base. In this case, I am overstating the variation in size of base across levels of development (Panel B, Figure 4). Perhaps more importantly, a size-wedge that increases with development means that I am understating the strength

of the association between size of tax base and income tax collection (Panel B, Figure 5). This point is also supported by the observation that the variance in residual tax collection, controlling for the statutory size of tax base, is larger in more developed countries (Panel C, Figure 5). This discussion suggests that bias introduced by the wedge between my measure of the base and the effective size of tax base only strengthens the main finding of the tax base being a first-order determinant of tax collection across development.

B US states material

B.1 Data sources and construction of variables

In this sub-section, I describe the construction of variables used in the US states analysis (Section 4).

B.1.1 Employment and earnings

I construct the aggregate employment-share variables using decennial Census data at the state level between 1930 and 2010. The data is extracted from IPUMS USA. In each decennial data-extract, I exclude from the sample any individual that is not economically active during the reference period and for whom the general class of worker variable is 0 ("N/A"). I also exclude, when possible, any individual who reports total personal income either equal to 9999999 ("N/A") or strictly negative. In the IPUMS USA data, total personal income corresponds to the respondent's total pre-tax personal income or losses from all sources for the previous year. I code as self-employed (employee) a respondent who responds 'self-employed' ('works for wages') in the class of worker category. This classification in IPUMS USA is consistent with the classification used in the cross-development sample, in the sense that I code the employment-type based on the primary job of the respondent in which they spent the most time during the reference day or week. Within each decennial extract, I apply person-weights to estimate, for each state, the representative total number of respondents, the total number of employee respondents, and the total number of self-employed respondents. I then calculate the employee-share as the ratio of total number of employee respondents to the total number of employee and self-employed respondents. I interpolate the numerator and denominator between Census years using a natural cubic spline (Herriot & Reinsch, 1973).

I construct the employment shares by income decile of the income distribution of each state, in 1935 and in every decade between 1950 and 2010. The 1950-2010 data is extracted from the IPUMS USA database. The definitions of type of work and industry are the same as those used to construct the state-year aggregate employment shares. I rank all respondents within a given state according to the reported total personal income. The personal income reported measures each respondent's total pre-tax personal income. Importantly, throughout the sample period, this measure is largely comparable: it includes in all samples, wage, farm and business components. I then apply person-weights and partition each state's income distribution into ten deciles (ten bins of equal sample size).

Within each decile, I estimate the conditional proportions of employees and self-employed to construct the employee-shares by income decile. In years before 1950, the decennial US Census does not provide reported income and occupation-category at the level of the individual. I use the 1935-36 Study of Consumer Purchases in the United States, which had the scope to 'ascertain for the first time in a single national survey the earning and spending habits of inhabitants of large and small cities, villages, and farm'. I access this data under the ICPSR data archive reference #08908. I discuss the 1935 data-sample and construction of variables in more detail in Section A.1.2. I construct the deciles of the state-specific income distribution and estimate the employment shares specific to each decile-state. I use these data to construct the profile of employment-share and self-employment share over deciles of each state's income distribution, for all continental states, between 1935 and 2010. I again interpolate both the numerator and denominator between data-years.

The earning structure is constructed for all states and all years between 1929 and 2001 by combining the two historical series, namely SA5H and SA5 'Personal Income by Major Components and Earnings by Industry' published by the US Bureau of Economic Analysis. The denominator for earnings-structure is line-item 45 'Net earnings by place of residence', which equals total earnings less contributions for government social insurance plus 'adjustment for residence'. The employee-share uses in the numerator line-item 90 'private non-farm earnings', while the self-employed share of income uses line-item 70 'proprietors' income'. The line-item 45 is also used as the denominator y to construct the ratio of the PIT-threshold K to average earnings, K/y. Importantly, this measure y of personal income excludes transfers from all levels of government, similarly to the gross income variable used in the cross-development sample.

B.1.2 Tax revenue

The tax-revenue sources by state and year are based on the historical series on state government finances published by the US Census Bureau. The State Government Finances series publishes series on yearly tax-revenue collected over the fiscal year of each state. I proxy for tax-take by constructing the ratio of tax-revenue collected to total personal income in the state, where the denominator is based on the BEA historical series of state personal income. This tax-take ratio differs from a more standard construction of the variable, used in the cross-development sample, where the denominator use a measure of aggregate output. Unfortunately, continuous GDP data at the state-year level in the US is

only available from 1963 onward. Instead, I follow previous papers studying growth in the US states (e.g. Barro and Sala-i-Martin, 1992; Besley et al., 2010) and use state personal income as a measure of state output. In the State Government Finances, T40 is the line-code corresponding to personal income tax; T41 corresponds to corporate net income tax; and, T09 corresponds to general sales tax.

B.1.3 Personal income tax structure: thresholds, rates, and reforms

To construct measures of the state PIT-base and state PIT-rate structure, I use data from the Bakija (2009) historical U.S. Federal and state income tax calculator program. I thank Jon Bakija for kindly providing me access to the calculator. The calculator models federal and state personal income taxes based on legal text, covering the period from 1900 to 2007 for state income tax laws. I construct the income tax threshold K for an individual earner who files under the status of being single and who claims the standard deduction. This filng behavior is directly comparable to the filing behavior chosen to calculate the exemption threshold in the cross-development sample. As such, the measures of thresholds and income tax base are comparable between the US states time-series and the development cross-country series. The choice of a single earner, as opposed to household earnings, is also consistent with the income distribution which is calculated based on ranking of total personal earned income. Finally, an appealing feature of the standard deduction is that, unlike the itemized deduction, the filer does not deduct state personal income tax from her federal income tax liability. This provides additional incentives for the filer to under-report state income taxes, and makes the filing-choice more similar to the under-reporting model derived in Section 5. Evidence from IRS statistics suggest that standard deduction filers are systematically more prevalent at lower levels of gross income (the Statistics of Income series on individual income tax returns regularly documents on this: see e.g. IRS, 1982). I construct the ratio K/y where y is the state-year per capita personal income, extracted from the historical US BEA series.

I use the same state tax calculator to construct measures of the tax-rate structure. The calculator provides data on the number of brackets for the specific filing-type, and the marginal tax rate which applies to each bracket. Some states have multi-bracketed structure with progressive marginal tax rates, other states apply a single-rate flat income tax over all taxable income. I use the marginal tax rate that applies on the first bracket in Table 3.

The measure for income tax reforms is coded in the following way. States began in the 1980s to

automatically adjust the nominal values of the exemption threshold (and rate-brackets) for inflation. Prior to this period, no state provided inflation adjustments. Prior to the 1980s, the dollar value of the calculated threshold K would therefore remain constant unless a legislative reform occurred which changes the value of the exemption threshold. I therefore code a year of reform as a year, before 1980, during which the nominal value of the threshold changed. I then construct the state-specific cumulative series of exemption reforms over time. I use this measure of reform likelihood in Panel A of Figure 7. In a graph that pools several States, the cumulative distribution measure has the advantage of controlling for cross-state heterogeneity in the frequency of threshold reforms.

B.1.4 Covariates

The poll tax and literacy test dummies are taken from Besley et al. (2010). They provide state-time varying measures of the share of the state population subject to either a literacy test or a poll tax. Prior to the 1965 Voting Rights Act, such measures were in place in predominantly Southern states. The 1965 VRA gave the Attorney General the authority to appoint federal examiners to oversee voter registration in states using literacy or qualification tests, and the power to seek legal action against poll taxes as a prerequisite for voting in state elections. Besley et al. use variation in these dummies to instrument for political competition, which they find to have a significant impact on the share of non-farm income and tax revenues. I also use the election year dummies from Besley et al.

I construct proxies for the state-year policy environment. These different proxies are meant to capture variation in state-policies which may have affected location decisions of private firms. The choice of proxies is based on historical readings which provide qualitative evidence that these policies contributed to the workforce transition into manufacturing and services jobs, especially in Southern and Midwestern states (Cobb, 1993; Newman, 1984). First, a dummy for the existence of a corporate income tax is constructed, which takes value 1 in all years in a state where there exists such a tax-base. The date of creation of stat corporate income tax is taken from Table 4.1 of Newman (1984). The dummy for the existence of right-to-work laws is extracted from Besley et al. (2010). Right-to-work laws make it illegal to demand that employees join a union, or to automatically deduct union fees from wages. The continuous measure of state unemployment insurance firm-size coverage is taken from the historical publication series 'Significant Provisions of UI State Laws" published by the US Department of Labor. I download all publications between 1937 and 1979. In each state-year, I code

the firm-size coverage, that is the lower-bound on firm-size above which an employee in a given firm is entitled to receive state UI benefits. This measure is defined consistently over the entire series. Federal-time varying regulation provided an upper-bound on the allowed firm-size, but states were free to legislate in order to define a firm-size below the Federally mandated size. Some states chose to lower the firm-size coverage earlier on, ahead of Federal regulations, while some states followed the Federal upper-bound throughout time. After 1979, Federal regulations extended coverage to all firms with one employee or more, and I code the state-time coverage as equal to 1 from 1979 onward. I also wanted to code the employer UI-contribution, expressed as a percentage of wages, but this measure is not consistently reported throughout.

B.1.5 Additional outcome variables

I construct a proxy for tax administrative reforms based on the historical series of the Book of the State, published annually from 1993 until today by the Council of State Governments. I collect data at the state-year level on the number of agencies administering major taxes: property, income, sales, gasoline, motor vehicle, tobacco, death, liquor. I code the total number of state tax agencies in operation in every state-year. This variable is available from 1939 to 2009. This variable is intended to proxy for investments in enforcement capacity, through consolidation of the number of tax agencies, and is used in the robustness checks (Table 3). I also collected state-year data from the same source on the annual salaries of the chief state administrative official in different departments: revenue-collection and taxation; treasury; attorney general. I then constructed the ratio of the annual salary in revenue-taxation relative to the salary in the Treasury and to the salary as Attorney General. These ratios were meant to proxy for investment in enforcement capacity through funding higher wages to tax administrators (relative to other state administrators). In results not reported, I do not find an impact of the upholding event on this measure of relative pay. These variables represent, to my knowledge, the first long-run time-series on proxies for tax administrative capacity of individual states in the US. As an additional proxy for enforcement capacity, I code the year when each state adopted withholding of state personal income taxes by employers. There exists both micro-evidence from Denmark (Kleven et al., 2011) and state-level evidence from US states (Dusek & Bagchi, 2017) on withholding's positive impacts on income tax collection. I use the historical IRS 'Annual Report' series to code the years of adoption. This variable is used in robustness checks (Table 3).

I use data from Besley et al. (2010) to build proxies for political outcome-variables. I use their measure of party-neutral political competition, which is defined as (minus) the absolute value of the deviation of the democratic vote-share from 50 percent, where the vote-share is the average vote-share over all state-wide races. Further, I use the Democratic vote-share averaged across all state-wide elections, and the Democratic seat-share in the state House. These measures are used as outcome variables in robustness checks (Table 3).

In the robustness checks (Table 3), I also study the impact of upholding on the generosity of the state's unemployment benefits. In particular, I use the measure of state maximum unemployment benefits. This variable is taken from the 'Correlates of State Policy' database (Jordan & Grossman, 2017).

Finally, in the robustness checks (Table 3), I study the impact of upholding on level of income, and income inequality. I use the 'net earnings' measure of income from BEA, and the top 1 percent income share from Frank et al. (2015).

B.1.6 Exchange of information agreements

In the main heterogeneity analysis (Table 2), I study whether the impact of upholding on tax structure and collection differs according to whether a state has an exchange of information agreement in place by the time of the court upholding decision. I code the year of implementation of the agreement from the historical IRS series 'Annual Report.' The signature of the exchange of information acts has been found to increase income tax revenue (Troiano, 2017). Troiano's source for the year of implementation is Penniman (1980). There are only minor differences in the year of implementation between the annual IRS publication series and Penniman (1980), and my results are robust to using this alternative measure of implementation dates.

B.1.7 Cost of collection

I construct the measure of cost of collection used in Section 4.1 from the Book of the States. The earliest year where the required data exists is 1962. The cost-components of collecting state taxes are capital outlays, operating costs, and payroll. In 1962, these measures exist for the state's financial admin-

istration, which includes the revenue administration and the procurement administration. As such, my cost measure constitutes an upper bound, since I cannot separate the administrative costs of the revenue division from the procurement division. I divide this total cost by the total gross tax revenue collected within the same financial year. This measure of cost of collection is similar in construction to Jensen & Lagakos (2019), which studies variation in cost of national tax administrations across levels of development. Interestingly, I find that the cost of collection in the average US state in 1962 is slightly higher than the average low-income country's tax administration from Jensen & Lagakos (2019).

B.2 Program details: Industrial Development Bonds

In this sub-section, I provide additional information on the Industrial Development Bonds program (IDB), and the legal uncertainty which generates variation at the state-level in the effective implementation date.

The IDB was a place-based local development program that sought to attract industrial facilities to predominantly rural areas characterized by 'surplus labor' concentrated among self-employed farmers (Advisory Commission on Intergovernmental Relations, 1962). The first state to implement an IDB program was Mississippi, when it launched 'Balancing Agriculture with Industry' in 1936. Practice of IDB did not, however, become a multi-State practice until the mid-1950s, when several other States decided to implement similar programs.

The official justification for government intervention was that these rural areas were "deficient in credit facilities" (ACIR, 1962), and capital for local firms was not readily available from conventional credit sources. Through the IDB, the local government therefore sought to relieve a local credit constraint. In the IDB program, sub-state government units (counties, boroughs, and cities) issue bonds to finance the acquisition or construction of facilities and equipment for lease to private firms. Importantly, IDB issuances were revenue bonds, which are secured exclusively by the revenues of the project. This is in contrast to general obligation bonds, which are secured by the credit of the issuer - in this case, the local government. This distinction implies that there is not a direct relationship between the issuance under the IDB program and increased tax revenue due to a need to solidify the local government's funding capacity.

The interest received from IDB securities was exempt from Federal income taxes. This meant that IDB securities commanded more favorable terms in the financial markets in relation to corporate securities with comparable risk. The Federal exemption is thought to have been one of the main reasons behind the growth of the IDB market. The growth of IDB issuances in the late 1960's implied an amount of forgone Federal government tax revenue which became intolerable. This triggered legislation in the early 1970s to remove the IDB exemption for Federal tax purposes and and to significantly limit the per issuance volume of IDB. These reforms also significantly widened the scope of projects that could be approved under municipal bond projects, with a shift away from rural industrialization towards public-goods projects in infrastructure and environmental conservation.

For identification purposes, I exploit the institutional features of implementation. In particular, the particular methods under IDB were unprecedented in the context of postwar state financing. The use

of public credit for an otherwise private purpose was considered to be in direct violation of the public purpose doctrine, which prohibits such usage. Constitutions of many States explicitly contained such public purpose statutes. The implementation of IDB therefore required, in a first instance, a legislative vote of constitutional or statutory provision that authorizes industrial development financing.

The lack of historical precedent, however, meant that the voted provision required judicial testing in order to be effectively implemented. Indeed, investors were reluctant to hold IDB securities in the period where the legality of the voted state provision had not been confirmed in the state's judicial system (Cobb, 1993). Judicial testing was most often delivered by a court case brought before the State's supreme court. This court case could be triggered in several ways. Most often, the issuance of an IDB required a significant amount of pre-issuance preparation, including a detailed description of the local workforce needs and a justification for why a particular candidate private firm would satisfy those needs. These preparations were often done by a local government agency, created specifically for this purpose. The case would then be brought against the legality of this local development agency. More generally, any legal step required to issue IDB could be targeted in a court case. In several States, including Tennessee, the IDB statute featured the requirement of a vote of approval by the relevant electors as a special municiapl election. The court case could also directly involve the issuance of an IDB bond itself. But as the graphical evidence in Panel A of Figure 7 shows, this was only very rarely the case.

In several instances, the fact that IDB were issued as revenue, rather than general obligation, bonds, was the basis of the argument for not violating the 'credit for public purpose' doctrine. In the case of Wayland v. Snapp, the Arkansas Supreme Court "(...) chose to uphold the issuance of the revenue bonds by invoking the doctrine that revenue bonds do not violate a credit clause because they are retired through lease revenues of the project, not out of tax funds" (Yale Law Journal, 1961).

I collect information on the dates of the legislative vote and the upholding from several sources. Importantly, I collect information from both administrative sources and legal reviews: Abbey (1965), ACIR (1963), Pinsky (1972), and Economic Development Administration (1978). The date for the vote is the year of appearance of the constitutional statute or provision authorizing local development financing. The date for the upholding event is the publication year of the leading case that upholds the constitutionality of the statute or provision. There is only little conflict in the reported dates of the vote and the upholding between the administrative and legal sources. In the case of upholding, there are sometimes several leading cases, when the first case upholds the constitutionality of the

statute allowing cities to issue development bonds, and the second (later) case extends these powers to counties. I always choose the earliest date across sources for both the vote and the upholding events.

Table B.1 provides information for each IDB state in my time-period of study. The table reports the year of vote and the year of upholding that I use in the main analysis. In Figure B.1, I plot the cumulative distribution of the time-lag between the vote-year and the upholding-year. The average lag is 6.67 years, with a standard deviation of 6.77. In just under 40 percent of States, the time-lag exceeds 10 years. In the main analysis, my estimation is helped by the existence of a significant lag between the vote and the upholding events within state; the variance in lag across States; and, the differential timing of court upholding decisions across States.

Figure B.1: Time-Lag Between Vote and Upholding Events Across IDB States

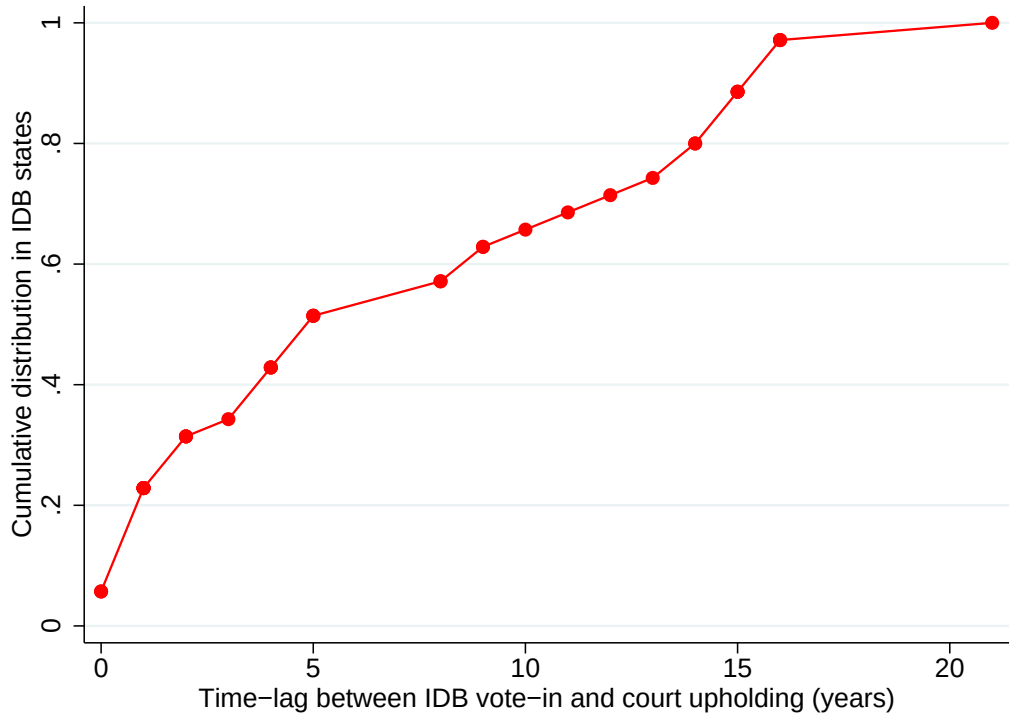

Notes: This graph displays the empirical cumulative distribution function of the time-lag in years between the vote event and the upholding event within each State that has upheld IDB by 1980. The year of the vote is the year where the State legislature voted in a statute or provision authorizing IDB. The year of the uphold event is the year where the State supreme court upheld the legality of the voted IDB statute or provision through a leading court case. The time-lag is defined as the difference in years between these two events. Source: Section 4.1 and Appendix Section B.2.

Table B.1: Industrial Development Bonds Program Legal Timing

State	Year vote	Year uphold	Leading court case
Alabama	1949	1950	Newberry v. City of Andalusia, 257 Ala. 49, 57 So. 2d 629
Arizona	1963	1973	Industrial Development Authority of Pinal County v. Nelson, 109 Ariz. 368, 509 P. 2d 705
Arkansas	1958	1960	Wayland v. Snapp, 232 Ark. 57, 334 S.W. 2d 633
Colorado	1955	1970	Allardice v. Adams County, 173 Colo. 133, 476 P. 2d 982
Delaware	1961	1962	In re Opinion of the Justices, 177 A. 2d 205
Georgia	1957	1970	In re Opinion on Sub. H. B. 24
Illinois	1951	1972	People ex rel. City of Salem v. McMackin, 53 Ill. 2d 347
Iowa	1963	1964	Green v. City of Mount Pleasant, 131 N.W. 2d 5
Kansas	1961	1962	State ex rel. Ferguson v. Pittsburgh, 364 P. 2d 71
Kentucky	1946	1950	Faulconer v. City of Danville, 313 Ky. 468, 232 S.W. 2d 80
Louisiana	1952	1954	Miller v. Washington Parish, 75 Southern So. 2d 394
Maine	1958	1966	Northeast Shoe Company v. Industrial and Recreational Finance Approval Board, 233 A. 2d 423
Maryland	1960	1974	Wilson v. Board of County Commissioners of Allegheny County, 273 Md. 30, 327 A. 2d 488
Michigan	1963	1966	City of Gaylord v. Beckett, 144 N.W. 2d 460
Minnesota	1961	1970	City of Pipestone v. Madsen, 178 N.W. 2d 594
Mississippi	1936	1944	Albritton v. City of Winona, 178 So. 799
Missouri	1960	1975	Atkinson v. Planned Industrial Expansion Authority of St. Louis, 517 S.W. 2d 36
Montana	1965	1970	Fickles v. Missoula County, 470 P. 2d 287
Nebraska	1960	1962	State ex rel. Meyer v. County of Lancaster, 113 N.W. 2d 63
Nevada	1959	1973	State ex rel. Brennan v. Bowman, 512 P. 2d 1321
New Hampshire	1955	1971	Opinion of the Justices, 278 A. 2d 357
New Mexico	1955	1956	Village of Deming v. Hosdreg Co., 62 N.M. 18, 303 P. 2d 920
North Dakota	1955	1964	Gripentrog v. City of Wahpeton, 126 N.W. 2d 230
Ohio	1955	1966	State v. Greater Portsmouth Growth Corporation, 218 N.E. 2d 446
Oklahoma	1960	1961	Application of The Oklahoma Industrial Financial Authority, 360 P. 2d 720

Table B.1: Industrial Development Bonds Program Legal Timing (end)

State	Year vote	Year uphold	Leading court case
Pennsylvania	1956	1968	Basehore v. Hampden IDA and Walker v. Butler County IDA, 248 A. 2d 212
Rhode Island	1958	1974	In re Advisory to Governor, 324 A. 2d 641
South Carolina	1962	1967	Elliott v. McNair, 156 S.E. 2d 421
South Dakota	1964	1968	Clem v. City of Yankton, 160 N.W. 2d 125
Tennessee	1951	1952	Holly v. Elizabethon, 241 S.W. 2d 1001
Utah	1953	1968	Allen v. Toole County, 445 P. 2d 994
Virginia	1962	1967	Industrial Development Authority of the City of Chesapeake v. Suthers, 208 Va. 51 155 S.E. 2d 326
West Virginia	1963	1964	State ex rel. Marion County v. Demus, 135 S.E. 2d 352
Wisconsin	1957	1973	Hammermill Paper Co. v. LaPlante, 205 N.W. 2d 784
Wyoming	1963	1967	Uhls v. State, 429 P. 2d 74

Notes: This table provides details on the legal timing of the IDB program in all States. The year of the vote is the year where the State legislature voted in a statute or provision authorizing IDB. The year of the uphold event is the year where the State supreme court upheld the legality of the voted IDB statute or provision through a leading court case. This leading court case is indicated in the final column of the table. The years of the vote and upholding event are drawn from administrative and legal reviews. Source: Section 4.1 and Appendix Section B.2.

B.3 Robustness of main regression results

In this sub-section, I provide robustness checks for the main results in the US states regression setting. In Table B.2, I maintain the same specification as in the main text, but consider alternative measures of the main outcome variables. In the main text, I studied the impact of the upholding event on employment structure using the employee-share of the active workforce. The disadvantage of this variable is that it is interpolated between Census years. In Column 1, as an alternative I use the employee-share of income. This variable is drawn from the SA5H BEA series and is continuous throughout the sample period. The variable is constructed as the ratio of total wages and salaries to total resident income.[3] Column 1 indicates that the upholding event led to a large and significant increase in the employee-share of income, in line with the main finding of an increase in the employee-share in employment.

The final three columns of Table B.2 investigate the robustness of the absence of a per capita income effect. In column 2, I use the BEA 'net earnings by place of residence', which equals total earnings less contributions for government social insurance plus 'adjustment for residence'. In Column 3, I use the Census-based measure of total personal income. This income measure is interpolated between Census years. In Column 4, I use the IRS-based measure of income, adjusted gross income (AGI), drawn from the top-income share series in Frank et al. (2015). I find an insignificant impact of both the upholding event and the vote-in event across these three alternative measures of income. Both the BEA and Census measures suggest an insignificant positive impact, while the IRS measure suggests an insignificant, but negative impact. The absence of an impact on income at the state-level using various measures is consistent with the regressions in Section B.4 which also fail to detect a per capita income impact, but at the local county-level. Note that the absence of an impact on income in this context is not inconsistent with other place-based program evaluations which have found positive development impacts. Indeed, the findings in those studies, including Kline & Moretti (2014), are based on long-run estimates, while my estimates only capture the short-run program impacts.

In Table B.3, I consider the robustness of the impact on employment-structure to alternative specifications. Column 1 replicates the result from the main specification. In Column 2, I remove the vote-in dummy from the main specification. The counterfactual is now entirely built from states that uphold IDB at a later date. This has only a minor impact on the estimated coefficient, which changes from 1.7 percentage points to 1.5 percentage points. In Column 3, I remove the main covariates from the main specification, which are the first stage instruments used in Besley et al. (2010), and election year

[3] It was also used in the graphical evidence in Figure 7.

dummies. This has no effect on the estimated coefficient. In Column 4, I augment the main specification with additional controls. These additional controls are plausible determinants of employment-structure, but were not included in the main specification because of their potential endogeneity. They are: log per capita income; an indicator for the existence of right to work laws; an indicator for the existence of a corporate income tax; and, a firm-size measure of the state's unemployment insurance. The sources and construction of these variables is described in Section B.1. The inclusion of these controls has no impact on the main estimate.

In Column 5, I allow for the determinants of the time-lag between the vote-in and upholding event to have an independent impact. In particular, Table 1 showed that the time-lag was shorter in states with civil law origins, and longer in states that had witnessed defaults for a historical public-private funding initiative. While these are state-specific but time-invariant characteristics, they may nonetheless be correlated with state-time varying determinants of employment structure. This would confound the impact of the upholding event. I therefore allow civil law states and historical default states to be on fully non-parametric time-paths throughout the sample period. Formally, I estimate

$$\begin{aligned} y_{st} =& \beta + \alpha \mathbf{1}\left(\text{Vote-in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \gamma_t \left(\mathbf{1}\left(\text{Civil Law}\right)_{\mathbf{s}} \times \gamma_t\right) \\ &+ \mathfrak{B}_t \left(\mathbf{1}\left(\text{Hist Default}\right)_{\mathbf{s}} \times \gamma_t\right) + \lambda X_{st} + \mu_s + \gamma_t + \varepsilon_{st} \end{aligned}$$

where all variables are defined as in the main text, and $\mathbf{1}\left(\text{Civil Law}\right)_{\mathbf{s}}$ and $\mathbf{1}\left(\text{Hist Default}\right)_{\mathbf{s}}$ are indicators taking a value of 1, respectively, if a state is has civil law origins or has experienced a historical default. The construction of these variables is described in the main text (Section 4.2). The inclusion of these time-paths marginally reduces the estimated coefficient on the upholding event, from 1.7 to 1.5 percentage points, which remains strongly statistically significant.

In Column 6, I investigate the possibility that my main control specification does not adequately capture differential convergence patterns in employment structure over time across states. Indeed, the IDB-implementation period was characterized by rapid structural convergence for the less-developed states in the US (Barro & Sala-i-Martin, 1992; Caselli and Coleman, 2001). To investigate this, I augment the main specification with an interaction between a linear time-trend and the cross-sectional level of state GDP per capita in 1940. That is, I estimate

$$y_{st} = \beta + \alpha \mathbf{1}\left(\text{Vote-in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \lambda \mathbf{X}_{st} + \mu_s + \gamma_t + \mathfrak{B}\left(\text{Initial\_Income}_{s1940} \times [t - 1940]\right) + \varepsilon_{st}$$

where variables are defined as in the main text, and where $\text{Initial\_Income}_{s1940}$ is the cross-section of initial GDP per capita in 1940, which is interacted with a linear time trend $[t - 1940]$. This leads to only a very

marginal reduction in the magnitude of the estimated impact of upholding.

Finally, I show that the results are robust to excluding the set of IDB states which were implementing the Tennessee Valley Authority program (Kline and Moretti, 2013), a concurrent Federal place-based development program. The joint IDB-TVA states are: Alabama, Kentucky, Mississippi, and Tennessee. I remove these states from the sample, and re-estimate the main specification on the reduced sample. Column 7 shows that this leads to no meaningful change in the estimated impact of the upholding event.

Table B.2: Alternative Measures of Employment and Income

	E-share of income	Avg Income (BEA)	Avg Income (Census)	Avg Income (IRS)
	(1)	(2)	(3)	(4)
1(Vote)	.000	23.369	38.807	-684.822
	(.004)	(29.475)	(43.544)	(978.986)
1(Uphold)	.013	.424	59.052	-624.991
	(.006)**	(33.114)	(55.131)	(1115.546)
Mean outcome variable	.707	1016	2003	1596
State FE	x	x	x	x
Year FE	x	x	x	x
State-year controls	x	x	x	x
States	28	28	28	28
State-year Obs	466	466	466	466

Notes: This table reports results from estimating the following regression

$$y_{st} = \beta + \alpha \mathbf{1}\left(\text{Vote-in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \beta \mathbf{1}\left(\text{Upheld}\right)_{st} \times \mathbf{1}\left(\text{EoI}\right)_{st} + \lambda \mathbf{X}_{st} + \mu_s + \gamma_t + \varepsilon_{st}$$

where s denotes state and t denotes time. $\mathbf{1}\left(\text{Vote- in}\right)_{st}$ indicates whether a vote has occurred in the state-House to allow issuance of IDB but the IDB has not yet been upheld, $\mathbf{1}\left(\text{Upheld}\right)_{st}$ indicates whether the State court system has upheld the constitutionality of the voted IDB statute or provision. The vote-in and upholding events are mutually exclusive. $\mathbf{1}\left(\text{EoI}\right)_{st}$ is an indicator variable taking a value of 1 when a State has passed an exchange of information agreement with the Federal Internal Revenue Service. In Columns 1, the outcome variable is the wage and salary share of of individual income, drawn from BEA historical data. In Column 2, the outcome variable is 'net earnings by place of residence', which is the BEA concept of personal income. It is equal to total earnings less contributions for government social insurance plans plus a residence adjustment. In Column 3, I use the measure of gross personal income from the decennial Census. This measure is interpolated between Census year, using a natural cubic spline. In Column 4, I use the IRS measure of income, adjusted gross income, which is drawn from the top income share series (Frank et al., 2015). The state-year controls, $\mathbf{X}_{st}$, are indicator variables for election year, and indicator variables for the existence of voting restrictions in the form of poll tax and literacy tests. These are the first stage instruments used by Besley et al. (2010) to study political competition and policy-making in US states. $*, **, ***$ denote significance at the 10%, 5%, 1% level. Robust standard errors clustered at the state level in parentheses. Source: Section 4.5 and Appendix Section B.3.

Table B.3: Alternative Specifications

	Employee-share of employment						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1(Vote)	.003		.004	.003	.004	.004	.004
	(.005)		(.005)	(.005)	(.006)	(.005)	(.006)
1(Uphold)	.017	.015	.017	.017	.015	.016	.018
	(.005)***	(.005)***	(.005)***	(.005)***	(.006)**	(.005)***	(.006)***
Specification	Main	Cross-sectional only	No controls	Extensive controls	Time path civil law states	Initial income time-trend	Exclude TVA states
Mean outcome variable	.707	.707	.707	.707	.707	.707	.777
State FE	x	x	x	x	x	x	x
Year FE	x	x	x	x	x	x	x
States	28	28	28	28	28	28	24
State-year Obs	466	466	466	466	466	466	409

Notes: this table reports results from estimating alternative specifications, described in detail in Section B.3. In all regressions, the outcome variable is the employee-share of the economically active workforce. Column 1 replicates the central finding from estimating the main specification (1). Column 2 removes the indicator variable for the vote event from the main specification. Column 3 removes the controls from the main specification. Column 4 augments the main specification with additional controls: log per capita income; an indicator for the existence of right to work laws; an indicator for the existence of a corporate income tax; and, a firm-size measure of the state's unemployment insurance. Column 5 augments the main specification with a full set of year indicator interactions with both the indicator for civil law origins and the indicator for historical rail default. Column 6 augments the main specification with an interaction between a linear time-trend and the cross-section of GDP per capita in 1940. Finally, Column 7 estimates the main specification, but on a reduced sample which excludes the four States (Alabama, Kentucky, Mississippi, Tennessee) which were part of the Tennessee Valley Authority development program. $*, **, ***$ denote significance at the 10%, 5%, 1% level. Robust standard errors clustered at the state level in parentheses. Source: Section 4.5 and Appendix Section B.3.

B.4 Local IDB impacts: evidence from county-level regressions

In this subsection, I study the impact of the IDB program at the county-level. While the effective start of the program is triggered at the level of the state supreme court (Sections 4.1- 4.2), the decision itself to issue IDB is predominantly made by counties within the state.[4] As such, a county-level analysis provides an assessment of the direct local economic impacts of the program. At the same time, the county-level analysis at the level of local implementation helps to shed light on the absence of an economically meaningful impact of IDB on non-employment outcomes.

In order to study the county-level impacts of IDB, I rely on two main data-sets. The first is the comprehensive county-level panel data-set ICPSR 2896 'Historical, Demographic, Economic, and Social Data: The United States, 1790-2002.' This data-set has been used in other studies of long-run impacts of local place-based development programs, including the Tennessee Valley Authority program (Kline & Moretti, 2013). I interpolate values between data-points in the ICPSR data-set. Note that during my period of interest, the primary source in ICPSR is the County Data Book. Since this source delivers data every five years, the interpolation period is smaller than between decennial Census which is used in the main analysis. Nonetheless, the data has the disadvantage that it does not contain a continuous measure of per capita income, which is the main object of interest in this county-level analysis. I therefore supplement it with a second county-year panel data-set. This data-set is the combination of the continuous BEA county-level per capita income data, which exists from 1969 onward; and, the 1959 Census module which measures per capita income in the cross-section of all counties. While I do have to interpolate per capita income between 1959 and 1969, this data-set nonetheless gives me a more naturally continuous measure of per capita income than the ICPSR data-set. The only disadvantage is that my sample only starts in 1959, while 7 states have voted in the IDB program before that date. The county-level analysis is therefore limited to the counties in the 21 states that vote in after 1959.

The aim is to investigate the impacts of the IDB by comparing counties with the program to counties without the program. The basis for this exercise remains the specification used in the main text, which assesses impacts by comparing changes before and after the upholding event, while controlling for any impact occurring during the vote in event. But without any additional modifications, this specification would rely on counties in different states as a counterfactual. Instead, I want to create a control county

[4]The decision could also be made by higher tiers of government, such as the the state government, or lower tiers of government, including cities. Data from Moody's Investor Service (1974) suggests that in practice, actual issuance was predominantly carried out by counties.

within the same state. This is a meaningful exercise since the IDB program was initiated at the county level and only a subset of counties in a given state would initiate IDB.

I assign treatment at the county-level within the state based on a list created by the federal government before IDB had become widely implemented. The Area Redevelopment Administration (ARA) is a federal agency that was created in 1961, with the aim of providing technical (data-driven) assistance to state and local governments to implement local development financing. For this purpose, the ARA created criteria that defined 'redevelopment areas.' These were predominantly rural geographical areas, characterized by "structural underemployment", where the encouragement of new industries was perceived as a solution to the stagnant levels of development (ARA, 1962). This characterization is effectively identical to the characterization of counties that IDB was targeted towards. In every state, the ARA compiled data from Census and the Departments of Health, Education, Welfare, and Agriculture, to establish a statistical profile of every county in 1961, and classify a subset of those as 'redevelopment areas.'

I digitize the list of 'redevelopment' counties based on the 'Statistical Profiles' in every state (ARA, 1961), and merge it with the main county panel data-set. This list has the appealing feature that it was created by a government entity which was not responsible for implementing IDB in the pre-IDB period. As such, the selection of counties into the list may be considered plausibly exogenous to unobservable county-time varying confounding determinants of local development.

I augment the empirical specification used in the main text with this list to create a difference-in-differences design. More specifically, I consider the ARA 'redevelopment' status to be a county-specific time-invariant assignment to program treatment. Since there exists counties on the ARA list that do not take up IDB, and there exists counties not on the ARA list that can take up IDB, this is an intent-to-treat design. The diff-in-diff evaluation will compare changes in outcome in ARA counties before and after the court upholding event to changes in outcomes in non-ARA counties within the same state, while controlling for any impacts that occurred during the vote-in event. Formally, I estimate

$$\begin{aligned} y_{cst} =& \beta + \alpha \mathbf{1}\left(\text{Vote-in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \pi\left(\mathbf{1}\left(\text{Vote in}\right)_{st} \times \mathbf{1}\left(\text{ARA}\right)_{c}\right) \\ &+ \phi\left(\mathbf{1}\left(\text{Upheld}\right)_{st} \times \mathbf{1}\left(\text{ARA}\right)_{c}\right) + \mu_c + \gamma_t + \varepsilon_{cst} \end{aligned} \tag{1}$$

where y_{cst} is the outcome of interest in county c, in state s, at time t, $\mathbf{1}\left(\text{Vote-in}\right)_{st}$ and $\mathbf{1}\left(\text{Upheld}\right)_{st}$ indicate whether a state has, respectively, voted in but not upheld or upheld the IDB program. $\mathbf{1}\left(\text{ARA}\right)_{c}$

is a county-specific, time-invariant indicator that takes value 1 if a county is on the ARA federal list of redevelopment areas, and μ_c and γ_t are county and year fixed effects, respectively. I cluster the standard error at the state level, to allow for spill-over between ARA and non-ARA counties within the same state. The time-window is identical to the one used in the main estimation (Tables 2-3): in every IDB state that upholds before 1971, I consider the time-period that ranges from 5 years before the vote-in event to 5 years after the upholding event.

The results are displayed in Table B.4. In the first column, I study the employee-share of the active workforce as the outcome variable. I find that the large, positive impact is concentrated in the ARA counties in the upholding period. In the following two columns, I find no overall impacts on the level of employment and urbanization. The absence of impacts on these two outcomes is consistent with the interpretation that the IDB program achieved its stated objective of reducing underemployment in specifically targeted rural areas. Issuance of IDB required documenting the specific local industrial needs of a county and a justification for why the size and characteristics of the proposed IDB facility would achieve this local need. In comparison to other place-based development programs, IDB was therefore highly targeted in nature and narrow in scope, aiming to finance industrial development commensurate with the specific local workforce needs. The increase in the employee-share and the absence of an impact on the size of the workforce suggests IDB primarily provided a transition into employee-jobs of 'underemployed', self-employed farmers. The absence of any change in urbanization suggests that workers did not migrate to the predominantly rural areas where the IDB facilities were being opened. Consistent with the absence of generalized economic impacts, the final two columns find no statistically significant impacts per capita income. The fourth column uses family income, measured in the ICPSR data-set, while the fifth column uses the continuous BEA per capita income measure. The impact on the continuous measure of per capita income is particularly insignificant, both statistically and economically.

Taken together, these county-level results provide additional evidence to support the absence of any meaningful non-employment development impacts, in the short-run 5-year window considered in this estimation strategy. In particular, the IDB program seems to have led to a significant transition from self-employment to employee-jobs but only locally in the specifically targeted IDB counties. The absence of any spill-over to work structure or workforce attachment in non-treated counties suggests sectoral re-allocation was limited. The absence of any change to levels of urbanization suggests migration from non-treated to treated counties was also limited. The compensation for migration costs

and the efficiency gains from sectoral re-allocation are two of the main mechanisms through which previous studies have found long-run positive income impacts of place-based programs. These mechanisms seem to not be significant forces in the IDB context in the 5-year short run.

Table B.4: County-Level Evidence on Local Impacts of IDB

	E-share	Employment	Urbanization	Log(Family Income)	Log(Personal Income)
	(1)	(2)	(3)	(4)	(5)
1(Vote)	-.007	-.003	-.003	-.016	.031
	(.006)	(.002)	(.003)	(.030)	(.043)
1(Vote)***1**(ARA)	.003	-.001	-.001	-.007	-.019
	(.007)	(.003)	(.009)	(.033)	(.038)
1(Uphold)	-.009	-.004	-.003	-.044	-.003
	(.008)	(.003)	(.005)	(.047)	(.072)
1(Uphold)***1**(ARA)	.031	-.002	.000	.044	-.008
	(.010)***	(.004)	(.010)	(.040)	(.051)
County FE	x	x	x	x	x
Year FE	x	x	x	x	x
States	21	21	21	21	21
County-year obs	5140	5140	5140	5140	5140

Notes: This table reports the results from estimating the following regression

$$\begin{aligned} y_{cst} =& \beta + \alpha \mathbf{1}\left(\text{Vote in}\right)_{st} + \theta \mathbf{1}\left(\text{Upheld}\right)_{st} + \pi\left(\mathbf{1}\left(\text{Vote in}\right)_{st} \times \mathbf{1}\left(\text{ARA}\right)_{c}\right) \\ &+ \phi\left(\mathbf{1}\left(\text{Upheld}\right)_{st} \times \mathbf{1}\left(\text{ARA}\right)_{c}\right) + \mu_c + \gamma_t + \varepsilon_{cst} \end{aligned}$$

where y_{cst} is the outcome of interest in county c, in state s, at time t, $\mathbf{1}\left(\text{Vote in}\right)_{st}$ and $\mathbf{1}\left(\text{Upheld}\right)_{st}$ indicate whether a state has, respectively, voted in but not upheld or upheld the IDB program. $\mathbf{1}\left(\text{ARA}\right)_{c}$ is a county-specific, time-invariant indicator that takes value 1 if a county is on the ARA federal list of redevelopment areas, and μ_c and γ_t are county and year fixed effects, respectively. The outcome variables are: employee-share of employment; economically active employment share of population; urbanization share of population; log family income; and, log personal income. The first 5 outcomes are drawn from ICPSR 2896, while the final outcome variable is constructed from historical BEA series and the 1959 Census. The sample is limited to counties in the 21 IDB States that vote in IDB statutes or provisions after 1959. $*, **, ***$ denote significance at the 10%, 5%, 1% level. Robust standard errors clustered at the state level in parentheses. Sources: Appendix Section B.4.

References

[1] Barro, R. J., X. Sala-i-Martin [1992]. "Convergence", *Journal of Political Economy*, 100(2).

[2] Caselli, F., J. Coleman [2001]. "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation", *Journal of Political Economy*, 109(3).

[3] Cobb, J. [1993]. *The Selling of the South*. University of Illinois Press, Chicago.

[4] de Mel, S., D. Mckenzie, C. Woodruff [2008]. "Measuring Micro-enterprise Profits: Musk We Ask How the Sausage is Made?", *Journal of Development Economics*, 88.

[5] Dusek, L., S. Bagchi [2017]. "Third-Party Reporting, Tax Collections, and the Size of Government: Evidence from Withholding", *Working Paper*.

[6] Economic Development Administration [1978]. *The Industrial Revenue Bond as a Financial Attraction Device*. Office of Economic Research, Washington D.C.

[7] Frank, M. E. Sommeiller, M. Price, E. Saez [2015]. "Series for Top Income Shares by US States since 1917", *WID Technical Note Series*, No.2015-7.

[8] Jordan, M., M. Grossman [2017]. *The Correlates of State Policy v.2.1*. East Lansing, MI: Institute for Public Policy and Social Research.

[9] Pinsky, D. [1972]. "State Constitutional Limitations on Public Industrial Financing: An Historical and Economic Approach", *University of Pennsylvania Law Review*, 111(3).

[10] Yale L.J. [1961]. "The Public Purpose of Municipal Financing for Industrial Development", *Yale Law Journal*, 70(5).