**From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2016

# ESSAYS ON AGRICULTURE AND INEQUALITY

Alvarez, Jorge Alejandro

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01wh246v556

*U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf*

# Essays on Agriculture and Inequality

Jorge Alejandro Álvarez García-Tuñón

A Dissertation

Presented to the Faculty

of Princeton University

in Candidacy for the Degree

of Doctor of Philosophy

Recommended for Acceptance

by the Department of

Economics

Adviser: Professor Richard Rogerson

May 2016

© Copyright by Jorge Alejandro Álvarez García-Tuñón, 2016.

All rights reserved.

# Abstract

In developing economies, wages in agriculture are significantly below those of other sectors and wages within each sector can be greatly unequal. This dissertation explores the forces behind these disparities.

Chapter 1 studies the large gap in wages between agriculture and the rest of the economy. Using a panel dataset on the universe of formal workers in Brazil, I use information on workers that switch sectors to decompose the drivers of this inter-sector gap. I find that most of the gap between sectors is explained by unobservable differences in the skill composition of workers, as opposed to differential pay of workers with similar skills. The evidence speaks against the existence of large short-term wage gains from the reallocation of workers out of agriculture and favors recently proposed Roy models of inter-sector sorting as drivers of lower average wages in agriculture.

Chapter 2, which is co-authored with Niklas Engbom and Christian Moser, studies the forces behind the inequality decline that took place in Brazil between 1996 and 2012. Using administrative linked employer-employee data, high-dimensional worker and firm fixed effects models are fitted to identify the sources of this decline. Compression in firm effects accounts for 45 percent of the total decline and compression in worker effects accounts for 24 percent. Half of the decrease in firm pay differences and a fifth of the decline in heterogeneity between workers are explained by observable characteristics. During this period, firm and worker characteristics became more dispersed, but a marked decline in returns to these variables lead to an overall inequality decline. This chapter concludes that changes in pay policies, rather than changes in firm and worker fundamentals, played the largest role in Brazil's inequality decline.

Chapter 3 studies the sources of cross-country differences in agricultural GDP per worker and average farm size, and suggests the lack of access to reliable savings assets as a potential driver of this variation. Motivated by recent micro evidence showing negative-return savings for the poor and large cross-country variation in intermediate input usage and farm size, a

model linking financial underdevelopment, land misallocation, and agricultural productivity is presented.

# Acknowledgements

I am grateful for the insightful advice and generous support from my adviser Richard Rogerson. I also want to thank Ilyana Kuziemko for her guidance and support in the completion of this dissertation.

Moreover, I appreciate the helpful comments from Mark Aguiar, Anne Case, Janet Currie, Angus Deaton, Thomas Fujiwara, Mike Golosov, Jeffrey Hammer, Oleg Itskhoki, Greg Kaplan, Nobu Kiyotaki, Alan Krueger, David Lagakos, Alex Mas, Ben Moll, Ezra Oberfield, Stephen Redding, Havery Rosen, and Tom Vogl, as well as participants at the Princeton Macroeconomics lunch seminar, the Princeton Center for Health and Wellbeing lunch seminar, and the Princeton Public Finance Working Group. In addition, I would like to thank Carlos Lessa and Luis Pinto at IBGE as well as Carlos Corseuil, Glaucia Ferreira, and Leandro Justino at IPEA for facilitating the work with Brazilian data. For its help in the completion of chapters 1 and 2 of this dissertation, I also gratefully acknowledge the financial support from the Private Enterprise Development in Low-Income Countries initiative from the Center for Economic Policy Research.

Finally, I am forever grateful to my fellow graduate students here at Princeton. I would like to especially thank Fernanda Marquez-Padilla and Carlos Sanz whose friendship, support and advice has been invaluable. I am also deeply appreciative of having collaborated with Christian Moser and Niklas Engbom, my co-authors for chapter 2 of this dissertation.

To my family and country.

# Contents


Abstract . . . . . iii

Acknowledgements . . . . . v

**1 The Agricultural Wage Gap: Evidence from Brazilian Micro-data 1**

1.1 Introduction . . . . . 1

1.2 Literature Review . . . . . 4

1.3 Data description . . . . . 7

1.4 The magnitude and evolution of the agricultural gap in Brazil . . . . . 10

1.5 Sources of the agricultural gap . . . . . 16

1.5.1 Human capital . . . . . 17

1.5.2 Unobservable characteristics . . . . . 22

1.6 How are compositional differences in characteristics generated? . . . . . 28

1.6.1 Motivating a Roy Model: The existence of sector-specific skills . . . 30

1.6.2 A Roy model of selection with mobility frictions . . . . . 33

1.6.2.1 Preferences and endowments . . . . . 34

1.6.2.2 Technology . . . . . 35

1.6.2.3 Competitive equilibrium . . . . . 35

1.6.2.4 Calibration strategy . . . . . 38

1.6.3 The explanatory power of the sorting mechanism . . . . . 41

1.7 Conclusion . . . . . 45

1.8 Appendix. The gap explained by observable differences after controlling for unobservable characteristics . . . . . 46
1.9 Appendix. Event-study of transitions into and out of agriculture . . . . . 48
1.10 Appendix. Within sector heterogeneity . . . . . 51

**2 Firms and the Decline of Earnings Inequality in Brazil 53**

2.1 Introduction . . . . . 53
2.2 Institutions and macroeconomic trends in Brazil . . . . . 58
2.3 Data . . . . . 60
2.3.1 Description of linked employer-employee data (RAIS) . . . . . 61
2.3.2 Description of firm characteristics data (PIA) . . . . . 65
2.4 Inequality trends in Brazil from 1988–2012 . . . . . 68
2.4.1 The evolution of earnings inequality . . . . . 69
2.4.2 Compression in different parts of the earnings distribution . . . . . 70
2.4.3 The (un)importance of worker observables . . . . . 71
2.4.4 The evolution of earnings inequality between and within firms . . . . 74
2.5 Empirical framework . . . . . 79
2.5.1 Introducing the two-way fixed effects model . . . . . 79
2.5.2 Determinants of the estimated firm effects . . . . . 83
2.5.3 Determinants of the estimated worker effects . . . . . 84
2.6 Results . . . . . 85
2.6.1 AKM decomposition . . . . . 85
2.6.2 The link between firm effects and firm characteristics . . . . . 88
2.6.3 The link between worker effects and worker characteristics . . . . . 92
2.6.4 Empirical support for the AKM model . . . . . 97
2.7 Conclusion . . . . . 101
2.8 Additional figures . . . . . 102
2.8.1 Evolution of real earnings levels . . . . . 102

2.8.2 Cross-section and time series comparison of alternative productivity measures . . . 103
2.8.3 Earnings inequality by education groups . . . 104

**3 Unproductive Land Savings 105**

3.1 Introduction . . . 105
3.2 Literature review . . . 107
3.3 Theoretical Model . . . 109
3.3.1 Preferences . . . 110
3.3.2 Endowments . . . 110
3.3.3 Technology . . . 111
3.3.3.1 Agricultural production . . . 111
3.3.3.2 Saving and borrowing . . . 112
3.3.4 Competitive equilibrium . . . 113
3.3.5 The agent's problem . . . 114
3.3.5.1 Consumption choice and risk aversion . . . 114
3.3.5.2 The farmer's problem . . . 116
3.3.5.3 Occupational choice . . . 120
3.4 A numerical example . . . 121
3.4.1 Allocations and aggregate productivity in competitive equilibrium . . . 122
3.4.2 Behavior of individual agents . . . 124
3.5 Empirical evidence . . . 128
3.5.1 Data . . . 128
3.5.2 Variation in land distribution, financial development and agricultural output . . . 129
3.5.3 Evidence on the impact of financial development on land misallocation 130
3.5.4 Evidence on the impact of land distribution on intermediate investments 132

3.5.5 Evidence on the impact of intermediate investments on aggregate productivity and the productivity gap . . . . . 133
3.6 Conclusion . . . . . 134

**Bibliography** **137**

# Chapter 1

# The Agricultural Wage Gap: Evidence from Brazilian Micro-data

## 1.1 Introduction

A key feature of developing economies is that wages in the agricultural sector are significantly below those of other sectors.[1] Additionally, these economies have most of their workforce in the agricultural sector. These two observations motivate a literature dating back to Lewis (1955) and Rostow (1960) that views the exit of workers out of agriculture as a fundamental mechanism of development. The body of work on agricultural development and inter-sector differences, however, has not completely settled the question of why so many workers stay in agriculture in spite of better wages being paid in other sectors. One possibility is that some barrier prevents the movement of workers across sectors, in which case wage gaps between agriculture and other sectors indicate unexploited potential gains from the reallocation of workers out of agriculture. A second possibility is that workers in agriculture are characteristically different from those in non-agriculture, in which case wage

[1] In a sample of developing countries studied by Vollrath (2014), the median average wage ratio between agriculture and manufacturing was 1.6. This is 1.9 when comparing agriculture against services. In the sample of countries studied by Herrendorf and Schoellman (2014), the median ratio between agriculture and the rest of the economy is 2.0.

gaps would not be evidence of potential wage gains. The objective of this paper is to shed light on which of these possibilities is a more likely explanation of the agricultural wage gap.

A challenge in exploring this question is assessing the role of unobserved worker characteristics. For instance, if an agricultural worker and a non-agricultural worker with the same observable characteristics (e.g. age and education) earn different wages, it is hard to distinguish whether the two sectors simply have differential pay for similar workers or the two workers are in fact different due to some unobserved characteristic. This paper assesses the role of unobserved characteristics by using panel micro-data on earnings and worker characteristics covering all sectors of the Brazilian economy from 1996 to 2013. The use of panel data is an improvement on the literature on agricultural wage gaps in developing countries, which has typically relied either on the estimation of structural models to match country-level moments or on the analysis of heterogenous cross-sectional surveys from a sample of countries. Specifically, the panel dimension of the data allows me to control for differences in both observable and fixed unobservable worker characteristics. Information on workers that switch between sectors (from now on referred to as 'sector-switchers') can therefore be used to distinguish whether the wage gap between agriculture and non-agriculture reflects differential pay of similar workers in the two sectors or, alternatively, whether the gap is due to differences in the composition of worker characteristics in each sector.

The main empirical finding of this study is that workers who transition out of agriculture experience limited compensation gains when compared to the overall gap in mean earnings between agriculture and other sectors. I conclude that the agricultural wage gap does not appear to be driven by differential pay of similar workers once fixed unobservable characteristics are controlled for. Instead, the largest share of the agricultural wage gap is explained by differences in the composition of worker characteristics in each sector. In addition, I find that the wage gap between agriculture and other sectors in Brazil declined significantly from 1996 to 2013 as the economy grew richer. This reduction is similar when comparing agriculture to both services and manufacturing, and it coincided with a decline in the share of

workers employed in agriculture — from 25% to 14%. Moreover, this decline does not appear to be driven by changes in educational attainment or country demographics. In fact, I find that age and education explain only a small share of the large wage gap in Brazil during the late 1990s, and that differences in the composition of these variables between sectors drove only a small share of the decline during this period. Most of the decline is driven by compositional changes in the distribution of fixed unobservable worker characteristics.

Both the limited wage gains experienced by workers who switch out of agriculture and the importance of worker composition differences between sectors pose a challenge for an agricultural wage gap model. Such a model must generate large declining wage gaps that do not result in large wage gains among sector-switchers. Building on the work of Roy (1951), a recent literature has proposed worker sorting as a possible explanation that is consistent with this pattern. In particular, Lagakos and Waugh (2013) and Young (2014) illustrate how workers with sector-specific skills can sort themselves into different sectors to generate large wage gaps. In this type of model, each worker faces a choice between two idiosyncratic wages in agriculture and non-agriculture. Workers with a comparative advantage in non-agriculture choose to work in that sector, and this generates a wage gap relative to workers who find it advantageous to stay in the agricultural sector.

To test the explanatory power of this mechanism, I build on the sorting model proposed by Lagakos and Waugh (2013) and test whether a calibrated model that targets micro-moments from sector-switchers can generate wage gaps of the magnitudes observed in Brazil in the late 1990s. I find that a large share of the gap level (66%) can be generated by this model. In addition, when calibrating the model to match moments of the data from 1996 Brazil and estimating the dynamics of the wage gap over time as the country grows, I ask whether growth can lead to a wage gap decline due to the exit of agricultural workers. I find that this mechanism can account for about a fourth of the wage gap reduction observed in Brazil during the period between 1996 and 2013.

The rest of the paper is structured as follows. Section 2 provides a literature review that relates this paper to the labor literature on inter-sector wage gaps and the macroeconomic literature on both wage and output per worker gaps between sectors. Section 3 describes the datasets to be used in the empirical estimations. Section 4 describes the magnitude and evolution of the wage and productivity gaps in Brazil as well as the decline in the share of workers employed in the agricultural sector. Section 5 assesses the role of observables, unobservables, and differential pay of similar workers in explaining the gap. Section 6 describes the mechanics and calibration of an economy where workers sort across sectors, as well as the power of the worker sorting mechanism in explaining the agricultural wage gap magnitude and its decline. Section 7 concludes.

## 1.2 Literature Review

Most studies show that large inter-sectoral wage gaps persist even after controlling for educational attainment and other worker observables. The remaining gap stems from either differential pay of similar workers or, alternatively, sector differences in the composition of both observable and unobservable worker characteristics.

U.S. labor studies have explored this distinction with mixed results. Using matched data from the Consumer Population Survey (CPS), Krueger and Summers (1988) argue that unobservable worker characteristics cannot explain much of the difference in wages between sectors. On the other hand, Murphy et al. (1987, 1990) also use the CPS and conclude that industry switchers receive only 27% to 36% of the total industry differential and thus nearly two-thirds of inter-sector wage gaps can be attributed to differences in the composition of worker characteristics in each sector. Also using US data, Gibbons and Katz (1992) find limited evidence for differential pay of similarly-skilled workers between sectors and instead highlight the role of differences in the composition of observable and unobservable characteristics.

International studies on developing countries have also highlighted the role of differences in observable and unobservable worker characteristics in explaining the gap. Vollrath (2014) finds that large wage differences exist between workers after controlling for observed human capital in a set of 14 countries. He explores whether these wage gaps could be the result of distortions that prevent workers from being paid the value of their marginal product in each sector. Using a misallocation framework similar to Hsieh and Klenow (2009), Vollrath (2014) estimates that potential gains from eliminating distortions and eradicating human capital misallocation are less than 5% in developing countries. If misallocation is not important, this implies that differences in the composition of worker productivity are likely to be more important drivers of the gap. Similarly, using a different sample of countries, Herrendorf and Schoellman (2014) regress wages on observables allowing for returns on observables to vary by sector. They conclude that most of the wage gap between agriculture and other sectors can be accounted for by differences in workers' human capital — and sector-specific differential returns— present in each sector.

However, because of data constraints, these studies are limited to the comparison of a diverse collection of cross-sectional surveys. This prevents rigorous empirical testing of whether differences attributed to unobservable characteristics or differential human capital returns could in fact be the result of other forces producing differential pay of similar workers. Mobility frictions and compensating differentials, for instance, are two alternative explanations consistent with both the differential returns on observables estimated by Herrendorf and Schoellman (2014) and the residual wage differences reported by Vollrath (2014). By using a panel dataset where workers are observed as they switch across sectors, the current study overcomes the limitations of cross-sectional data and distinguishes the role of fixed unobservable characteristics from alternative stories of differential pay.

The study of wage gaps is also closely related to the study of output per worker gaps between agriculture and other sectors. Kuznets (1971), Caselli (2005), Restuccia, Yang and Zhu (2008), among others, have argued that a large share of income differences across

countries is explained by labor productivity gaps between agriculture and other sectors. However, focusing on output per worker, even in advanced countries, risks exposure to important sources of measurent errors. For instance, Gollin, Parente, and Rogerson (2004) suggest that unaccounted home production understates agricultural output and Herrendorf and Schoellman (2014) point out that errors in value added measurement muddy comparisons of worker productivity across US States. Partially as a result of this, the role that both observed and unobserved human capital play in explaining these output per worker gaps is still an open debate. Herrendorf and Schoellman (2014) argue that human capital accounts for most of the output per worker gap between agriculture and other sectors in the US and other selected countries. In contrast, Gollin, Lagakos, and Waugh (2013) argue that human capital —along with adjustments to labor supply— account for only about a half of the gap in the developing countries they study. Focusing on wages avoids many of the problems with the measurement of differences between agriculture and the rest of the economy. Though wages and output per worker are not equivalent measures of labor productivity, the results of this paper can speak to some of the debates about the role of differences in worker composition on inter-sector gaps explored by this literature.

Beyond establishing the role of worker characteristics in explaining the inter-sector gaps, a second objective of the literature is to establish which mechanisms are behind compensation and output per worker differences. Two main types of mechanisms are relevant to this study. The first are distortions that create wedges in marginal productivity of labor between sectors. These distortions can include scale effects that impact the allocation of resources across agricultural firms (Adamopoulos and Restuccia (2011), Donovan (2012)) or barriers that prevent the free flow of capital and workers (Restuccia et al. (2008), Herrendorf and Teixeira (2011)). Distortions that prevent marginal labor products to equalize have also been studied at the firm level by Restuccia and Rogerson (2008) and Hsieh and Klenow (2009), who highlight their greater importance in developing countries. To the extent that

these distortions are also present between sectors —and workers are not freely mobile— the mechanisms generating productivity gaps can be related to the agricultural wage gap.

A second type of mechanism highlighted by Young (2014), and Lagakos and Waugh (2013) portrays wage gaps as the result of sector differences in worker skill composition. Lagakos and Waugh (2013) illustrate how such skill differences can be the result of an equilibrium outcome. In their model, workers sort themselves to the sector where they are most productive. This process induces differences in the composition of worker skills employed by each sector, and this in turn generates a gap in mean wages paid in agriculture relative to non-agriculture. Importantly, the agricultural gap in this context isn't the result of any additional distortions that induces differential pay of similar workers. Building on this idea, Young (2014) uses cross-sectional surveys from developing countries to show how migration is consistent with rural-urban consumption driven by the sorting of workers. Though his focus is on consumption, his findings are also consistent with agricultural wage gaps generated by the sorting of workers with different unobservable skills. The mechanism proposed by this paper —which is also supported by the empirical results to be presented— belongs to this family of sorting models, where the agricultural wage gap is ultimately driven by compositional differences in worker characteristics.

## 1.3 Data description

Two main databases are used. The first is the set of Brazilian household surveys from the *Pesquisa Nacional por Amostra de Domicílios* (PNAD) from 1996 to 2013. This contains a representative sample of households covering all of Brazil. The survey includes both formal and informal workers and records demographic and employment-status characteristics as well as monthly earnings for all members of a household. In this paper, this data is used to show trends in sector employment shares among all workers, including both formal and informal, during the period of study. In particular, I establish that the trends and magnitudes in inter-

sector pay differences in Brazil among all workers are similar to the ones observed among formal workers. Data from PNAD is also used to compute the total number of workers in each sector and —in combination with the national accounts recorded by the *Instituto Brasileiro de Geografía e Estadística* (IBGE)— value added per worker for each year and sector. Due to the cross-sectional nature of the surveys, however, individuals cannot be followed over time in the PNAD and I am therefore unable to control for worker unobservable characteristics using data on both formal and informal workers. For this reason, most empirical decompositions in this paper focus on formal sector data which is now described in greater detail.

Data on formal workers comes from the *Relação Anual de Informações Sociais* (RAIS), which is administered by the Brazilian Ministry of Labor and Employment. This database is constructed from a mandatory annual survey filed by all formally registered firms in Brazil and contains earnings, occupation and demographic characteristics of workers as reported annually by their employers.[2] Importantly, each worker in the data has a unique and time-invariant worker ID that does not change as a worker switches employers. This feature of the data allows me to follow individuals over time and create a panel of the universe of employed formal workers across all sectors. In addition, each worker is linked to their employing firm, which also has a unique and time-invariant ID. This allows me to link workers to their respective sectors, and identify transitions between sectors.[3] The data covers the period from 1996 to 2013.[4]

The RAIS dataset reports average monthly gross labor earnings including regular salary payments, holiday bonuses, performance-based and commission bonuses, tips, and profit-sharing agreements as well as the start and end month of the job. To account for heterogeneity in the duration of job-spells, I divide annual earnings by the number of months

[2]It is common practice for businesses to hire a specialized accountant to help with the completion of the RAIS survey to avoid fines levied on late, incomplete, or inaccurate reports, which makes the quality of the data superior to household surveys.

[3]IDs available are anonymized to protect the identity of both workers and firms.

[4]Though earlier years are available for a large subset of Brazilian workers, the lack of universal coverage in earlier periods can be particularly problematic in studying transitions out of agriculture. Hence, the analysis is restricted to this later period.

worked at each job within a particular firm. A worker might have multiple spells in a year if he or she switched employers during the year or worked multiple jobs, but on-the-job earnings changes within a year are not recorded. To standardize the dataset at an annual level, I restrict attention to a unique observation per worker-year by choosing the highest-paying among all employment spells in any given year. One limitation of the data is that it does not contain hours worked for all of the period of study. Therefore, I cannot transform the measure of monthly earnings into a measure of hourly wages. The implications of interpreting results from sector earnings differences as wage gaps is discussed in section 1.5.2. For the rest of the paper, differences in the measure of monthly earnings described between sectors will be referred to as wage gaps.

The dataset also contains the age and educational attainment of each worker. Both of these variables are classified by groups. In particular, age groups of 18-24, 25-29, 30-39, 40-49, and 50-54 years are used.[5] Similarly, educational levels are classified into less than high school, high school, some college education, and completed college education. In all regression specifications utilizing age and education as explanatory variables of the wage gap, a full set of age and education interacted dummies is used.

Finally, to identify the employment sector and occupation of workers, classification is based on categories from the IBGE. Both the industry and occupation classification system changed during the period of study. Here, I use conversion tables provided by IBGE to standardize classification between different years and choose categories for both occupations and sectors coarse enough in order to avoid potential biases arising from mechanical changes in the classification system over time. The three sectors used are Agriculture, Manufacturing (including energy and mining), and Services. Occupation categories used are at the 3-digit disaggregation level.

Due to imperfect matching of all categories within a sector and occupation classification system, I exclude firms with inconsistent sector classifications so that sector switchers are

[5]Workers over 55 years old are excluded to avoid biases stemming from workers close to retirement age.

not incorrectly specified. I also exclude individual observations that have either firm IDs or worker IDs reported as invalid as well as data points with missing earnings, dates of employment, educational attainment or age. Together, these basic cleaning procedures drop less than 0.5% of the original population. The low percentage of missing responses and inconsistencies are indicative of the high quality of this administrative dataset.

Table 2.1 provides key summary statistics for the RAIS data for three sub-periods: 1996-2001, 2002-2007, and 2008-2013. Some features of the data are worth noting. The first is that the number of workers increases substantially over time from 49 million workers in 1996-2001 to 80 million in 2008-2013. This rise is mainly the result of two forces: population growth and an increase in formality in Brazil. A second observation is that education is quite different in agriculture in Brazil relative to other sectors. In 1996-2001, for instance, only 5% of formal workers in the agricultural sector had a high school degree and 1% had completed college, relative to 34% and 10% in other sectors. During 1996-2013, educational attainment substantially improved partially as a result of educational reforms in the late 1990s and the rise of social programs in the 2000s. In contrast, the age distribution in each sector did not change substantially. The explanatory power of age and education will be one of the focal points of the analysis. Finally, though earnings between agriculture and other sectors are quite different, there are virtually no gaps in earnings when comparing services against manufacturing in all periods. This motivates the dual economy focus of this paper: explaining the gaps between agriculture and all other sectors in the economy.

## 1.4 The magnitude and evolution of the agricultural gap in Brazil

Differences in pay between agriculture and other sectors were significantly reduced in Brazil during the last two decades. The ratio of mean earnings between non-agriculture and agriculture among all workers (both formal and informal) in the economy —as measured by the

Table 1.1: RAIS Summary Statistics

| | | | | log(Earnings) | | Education | | Age | |
|---|---|---|---|---|---|---|---|---|---|
| | | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Period | Sector | # Worker-years | # Unique Workers | Mean | Std. dev. | Mean | Std. dev. | Mean | Std. dev. |
| 1996-2001 | Agriculture | 9.5 | 3.1 | 6.22 | 0.60 | 1.10 | 0.43 | 34.60 | 11.33 |
| | Manufacturing | 34.3 | 9.4 | 6.85 | 0.82 | 1.43 | 0.75 | 32.51 | 10.12 |
| | Services | 131.4 | 36.5 | 6.84 | 0.84 | 1.74 | 0.97 | 34.82 | 10.92 |
| | All | 176.0 | 49.2 | 6.81 | 0.84 | 1.64 | 0.93 | 34.36 | 10.83 |
| 2002-2007 | Agriculture | 13.0 | 3.9 | 6.38 | 0.51 | 1.10 | 0.53 | 34.74 | 11.32 |
| | Manufacturing | 43.3 | 11.0 | 6.84 | 0.75 | 1.60 | 0.79 | 32.78 | 10.36 |
| | Services | 172.3 | 45.0 | 6.84 | 0.77 | 1.94 | 1.00 | 35.11 | 11.22 |
| | All | 229.6 | 60.2 | 6.82 | 0.77 | 1.83 | 0.97 | 34.66 | 11.11 |
| 2008-2013 | Agriculture | 13.8 | 3.8 | 6.66 | 0.47 | 1.24 | 0.61 | 36.01 | 11.48 |
| | Manufacturing | 56.3 | 13.6 | 7.01 | 0.68 | 1.80 | 0.80 | 33.57 | 10.71 |
| | Services | 246.0 | 62.4 | 6.98 | 0.71 | 2.09 | 0.97 | 35.49 | 11.47 |
| | All | 317.7 | 80.2 | 6.98 | 0.71 | 2.00 | 0.95 | 35.18 | 11.36 |

Note: Number of workers and worker-years are in millions. Earnings refer to average monthly earnings in real terms (Using 2013 Reais). Education levels are defined as 1= Primary or middle school or no education, 2= high school 3= some college education and 4= college completed. Age is in years.

PNAD household surveys— declined from 2.1 in 1996 to 1.7 by 2013. As discussed above, the main contributions of this paper hinge on the use of the panel structure of the data so that workers can be followed over time. Since this feature is only available for formal workers, the rest of the paper will focus on formal sector data. Similarly to the overall economy, formal workers exhibit a very similar decline in the ratio of mean earnings between non-agriculture and agriculture from 2.3 in 1996 to 1.6 in 2013 (Figure 1.1). Moreover, the magnitude of the gap and its decline has been similar when comparing agriculture to both services and manufacturing individually. In contrast to the differences between agriculture and non-agriculture, mean earnings in the two non-agricultural sectors were close to equal throughout this period.

Another feature of the data is that the agricultural wage gap is present throughout the earnings distribution. Figure 1.3 shows the ratio of earnings percentiles in agriculture and non-agriculture. Percentiles are here defined by the ranking of workers within each sector. There is a pattern, with the top earners in the agricultural and non-agricultural sectors being further apart than the bottom earners in the two sectors. The differences, however, are still significant across all percentiles and it is not the case that wage gaps are a phenomenon that is only applicable to certain parts of the earnings distribution. Furthermore, when looking at the evolution of these ratios over time, the decline in compensation differences does not appear to be driven by the catch up of only the poorest or richest parts of the distribution of agricultural workers.

In addition, the wage gap decline was accompanied by a similar decline in the value added per worker gap. Figure 1.2 shows how the between-sector difference in gross domestic product per worker as measured by the national accounts declines over the 1996-2013 period. Similar to the earnings pattern, the decline is large when comparing agriculture against both manufacturing and services. Unlike earnings, however, the differences and the magnitude of the decline is much larger when looking at the agriculture-manufacturing gap than when looking at the agriculture-services one. This is expected due in part to the natural differences

in capital intensities between services and agriculture. These differences notwithstanding, the qualitative pattern of pay and value added per worker gaps is qualitatively similar. Importantly, this reduction in sectoral inequities occurred during a period where yearly real GDP growth averaged 2.7%, as the country transitioned out of a period of macroeconomic instability and hyperinflation into a period of technology modernization and growth.[6] The interrelation of growth, productivity, and the decline in inter-sector gaps will be central to our analysis of mechanisms in section 1.6.

The magnitudes of both the wage and value added per worker gaps between agriculture and other sectors are large when compared with other estimates in the literature. In 1996, the magnitude of the value added per worker gap between agriculture and other sectors is 5.3, which is greater than the maximum found by Herrendorf and Schoellman (2014) in their 12 country sample and just below the mean gap reported in Gollin, Lagakos and Waugh (2013) for the poorest quartile of countries in their 151 country sample. By 2013, after a cumulative real output growth of 61%, the value added per worker gap is 2.4. This estimate is similar to the median of 2.3 in the Herrendorf and Schoellman (2014) sample and closer to the 2.0 mean of the richest 25% of countries in the Gollin, Lagakos and Waugh (2013) sample. When compared to the cross-sectional evidence, Brazil appears to have endured a significant transformation during the period of study.

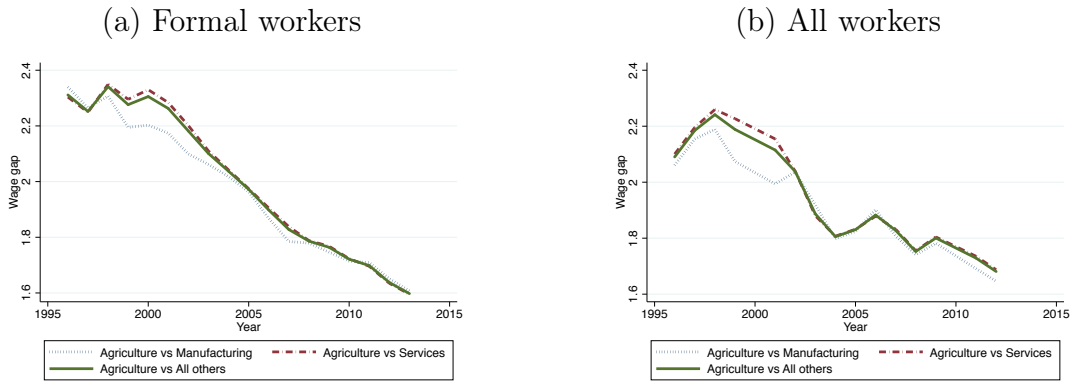
In terms of the wage gap between agriculture and other sectors, Brazil's 1996 wage gap of 2.3 is above the median of 2.0 from the Herrendorf and Schoellman (2014) sample. By 2013, it falls below the sample's mean to 1.6. Figure 1.4 shows how these gaps compare to the list of 15 developing countries studied by Vollrath (2014). Brazil's 1996 gap between agriculture and manufacturing would rank third highest, just below Ecuador in that sample. When comparing agriculture vs services, the rank would be 5th, just above Indonesia. In contrast, Brazil's 2013 gap levels with respect to manufacturing and services would rank 8th and 11th, respectively. Though the data on Brazil is not entirely comparable to the

[6]Bustos and Bruno (2015) explain some of the agricultural modernization of the agricultural sector in Brazil.

wage data from other countries' surveys, the significant move down the ranking of countries suggest that Brazil's decline cannot be described as an insignificant change.

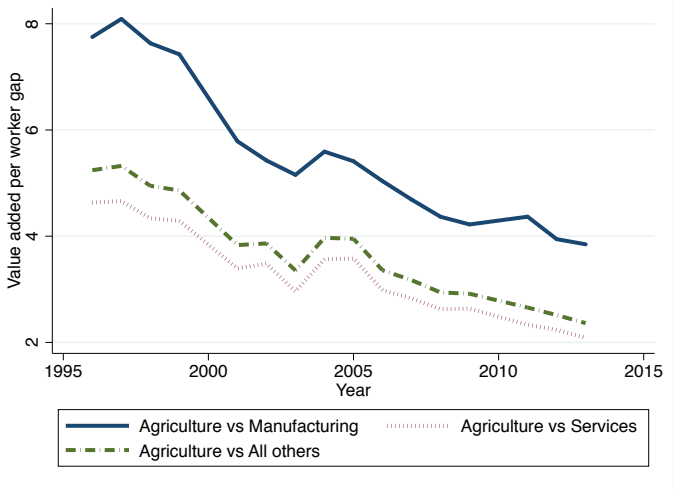
In parallel to the closing of both output per worker and wage gaps, Brazil also endured a substantial transformation of the employment structure. The workforce composition based on household surveys is shown in Figure 1.5. The economy employed 25% of the labor force in agriculture in 1996, which declined to 14% by 2013. Manufacturing employed 13-15% throughout this same period, and services increased from 61% to 72%. Among formal workers, a similar pattern is observed and the share of workers in agriculture has declined from 5% to 3.6% since 1996. Though the population of formal workers is much smaller than the universe of workers in agriculture, the magnitude of the wage gap also shows a declining pattern in the share of labor employed in agriculture. The interrelation between the movement of workers out of agriculture and the agricultural wage gap will be considered in section 1.6, when mechanisms behind the gap's decline are discussed. First, a statistical decomposition of the agricultural wage gap is conducted using the panel structure of the data on formal workers.

Figure 1.1: Wage gap in Brazil




Note: The wage gap is calculated as the ratio in average labor monthly earnings between agriculture, manufacturing and services as classified by the IBGE. Data on formal workers comes from the *Relação Anual de Informações (RAIS).* Data on all workers (both formal and informal) comes from the PNAD household surveys.

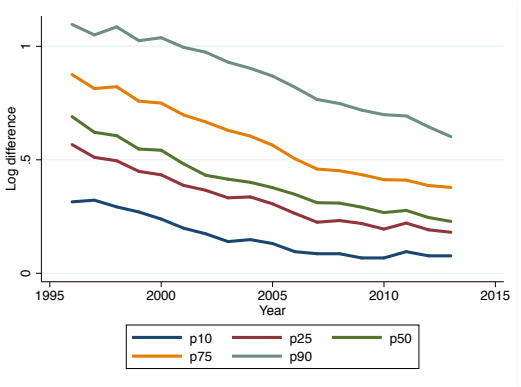
Figure 1.2: Value added per worker gap in Brazil



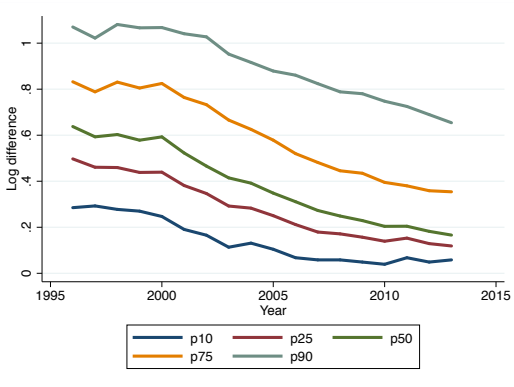

Note: Value added per worker gaps are constructed from national accounts available from IBGE and labor statistics from the Pesquisa Nacional por Amostra de Domicílios (PNAD).

Figure 1.3: Gaps in Brazil by percentile

(a) Agriculture vs manufacturing

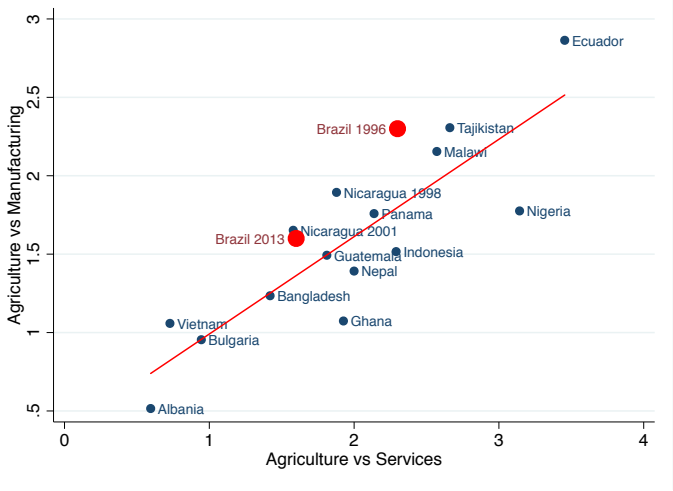



(b) Agriculture vs services




Note: Difference in the means of log earnings between sectors for formal workers are presented. Each line corresponds to the difference between each percentile group in the two sectors.

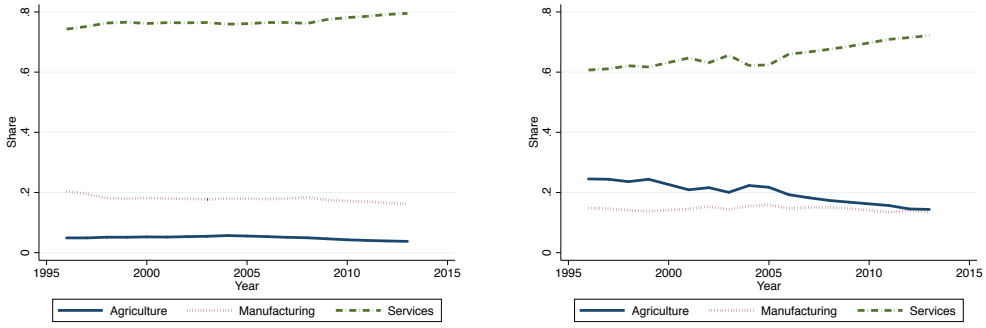
Figure 1.4: Wage gaps in Brazil vs other countries




Note: Data for Brazil comes from PNAD and national accounts from IBGE. For other countries, estimates are constructed based on cross-country data from Vollrath (2014).

Figure 1.5: Workers by sector

(a) Formal workers (b) All workers




Note: Share of total employed workers. Formal worker data is from RAIS. Data on all workers is from PNAD household surveys.

## 1.5 Sources of the agricultural gap

We now turn to explore what drives the wage gap between agriculture and other sectors. Three possible alternatives are considered. The first are differences in the composition of

observable human capital as measured by age and education. The second are differences in the distribution of fixed unobserved worker characteristics between sectors. Finally, the third alternative is the presence of mechanisms that induce differential pay of similar workers employed by different sectors. Inter-sector mobility frictions, sector-specific rent-sharing agreements, and compensating differentials are some of the mechanisms that fit this third category. This section argues that the first two alternatives, where the gap is driven by compositional differences in worker characteristics, explain most of the agricultural wage gap and its decline.

### 1.5.1 Human capital

Differences in human capital introduce heterogeneity in the productivity of workers which, in a standard competitive environment, should translate into wage differences. Table 2.1 indeed shows differences in education between sectors, with agriculture workers being on average less educated than their peers in services and manufacturing. To the extent that these characteristics determine human capital, these differences can potentially explain part of the agricultural wage gap.

There are two margins on which human capital influences the wage gap. On one hand, human capital can be lower in one sector than the other. On the other hand, even if the *composition* of human capital is the same in the two sectors, the *returns* to human capital might be different in the two sectors. I first assess whether compositional differences in human capital, as measured by age and education, can account for a substantial share of the gap by estimating the following model for each year $t$.

$$log(w_{ist}) = F_t(education_{ist}, age_{ist}) + \epsilon_{ist}$$

Here, $education_{ist}, age_{ist}, w_{ist}$ are the wage, education level, and age of worker $i$ in sector $s$ in year $t$. To impose minimal restrictions on how age and education influence wages, the map-

ping of education and age to wages is specified as $F_t(education_{ist}, age_{ist}) = \sum_{a,e} \mathbf{1}(age_{ist} = a, edu_{ist} = e) \times \beta_{aet}$. Thus, the specification allows full flexibility in terms of both age and education, and this relationship can vary in every year of the sample.

For the rest of the paper, I will define the wage gap as the mean difference of log earnings with respect to agriculture. Specifically, the gap between sector $s'$ and agriculture is defined as

$$\Delta_{s'} E(log(w_{ist})) \equiv E(log(w_{ist})|s = s') - E(log(w_{ist})|s = a)$$

where the possible values for sector $s'$, $\{a, m, s\}$, refer to agriculture, manufacturing and services respectively. The focus on additively separable mean log-wage gaps is done to simplify the presentation of the log-linear models to be studied.

Figure 1.6 shows the decomposition of the mean log difference into two parts: a component due to age and education and another due to the residual. There, we can see that the effect on wages from age and education differences between agriculture and other sectors have remained roughly constant throughout 1996-2013. When comparing agriculture vs manufacturing, these observable characteristics explain a nearly constant .09 log points of the gap. When comparing agriculture vs services, observables matter more and wage gaps due to age and education have averaged .24 log points. Overall, age and education differences accounted for 10% to 26% of the wage gap level during the period. The results show that most of the wage gap level is largely driven by factors not accounted by compositional differences in age and education alone.

Moreover, the decline in the wage gap cannot be entirely attributed to changes in education and the distribution of age in each sector. When comparing manufacturing and agriculture, the stability of the gap due to age and education shown in Figure 1.6 contrasts the decline in the overall wage gap. When comparing services and agriculture, age and education explain some of the decline, but the flatter pattern of this component relative to the total gap decline indicates that this reduction is not sufficient to explain the entire decline

in the gap between these two sectors. Similarly to the wage gap level, the decline does not appear to be explained by either demographic changes or educational attainment.

Figure 1.6: Gap in mean log earnings between agriculture and other sectors due to age and education

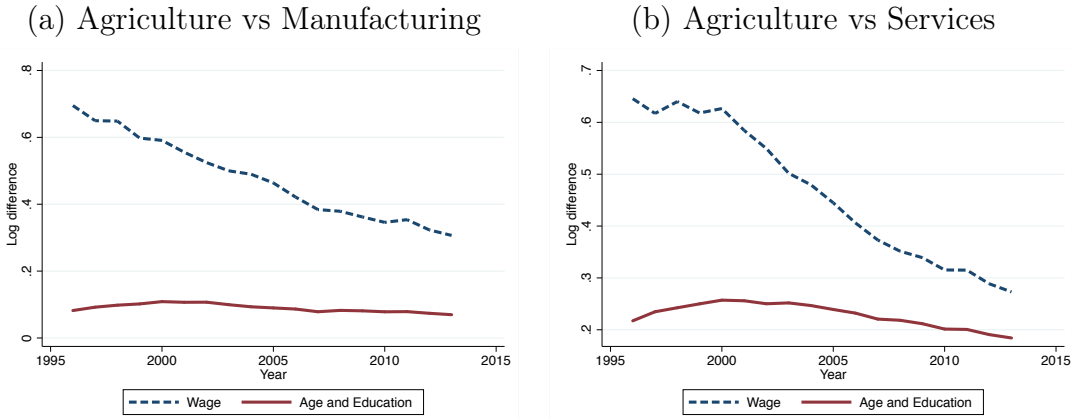
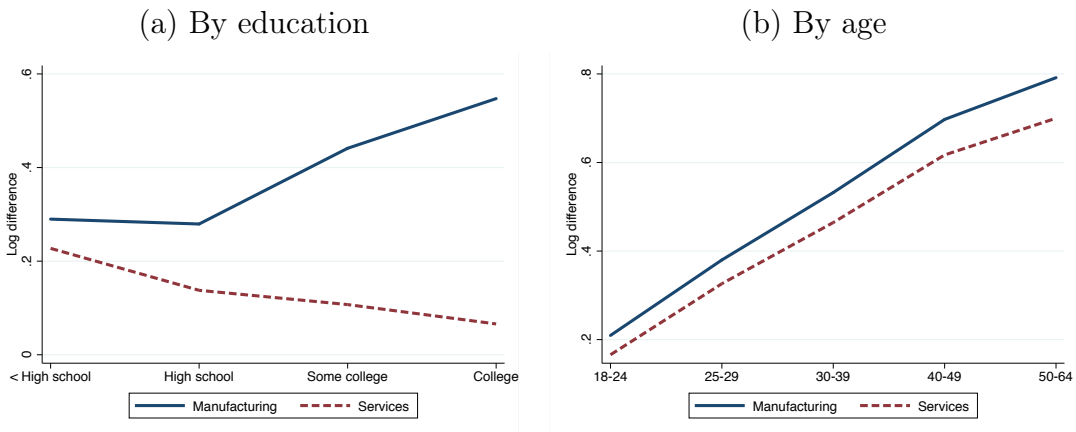



Figure 1.7: Mean difference in log earnings relative to agriculture by educational attainment and age




The approach above, however, does not allow for differences in pay within each education-age group across sectors. Figure 1.7 shows that large average wage differences by education and age groups exist, with older workers gaining significantly less in agriculture relative

to other sectors and workers in each age and education group being paid less than their comparable peers in non-agriculture. The difference in average pay for worker characteristics in each sector may reflect differential returns to education and experience by sector. For instance, worker with a high school degree might be more productive in manufacturing and services than in agriculture due to the availability of jobs that require this level of educational attainment.

The question is then to what extent do composition vs differential pay of each education-age group can explain the overall gap. In order to separate these components, I conduct a Oaxaca decomposition with agricultural workers as the reference group (Oaxaca (1973)). For notation simplicity, let $F_{st}(education_{ist}, age_{ist}) = \beta_t^s X_{it}^s$, where $X_{it}^s$ is a vector of dummies for each age-education group in sector $s$. We can then decompose the wage gap in each year as follows:

$$\begin{aligned}\Delta_{s'}(E(log(w_{ist}))) &= \beta_t^a(E(X_{it}^{s'}) - E(X_{it}^a)) + (\beta_t^{s'} - \beta_t^a)E(X_{it}^a) \\ &+ (E(X_{it}^{s'}) - E(X_{it}^a))(\beta_t^{s'} - \beta_t^a)\end{aligned}$$

The first term is entirely due to composition effects due to age and education differences in workers employed by sector $s$' relative agriculture. In other words, this component reflects the mean wage gap if all education-age groups were equally paid in both agriculture and sector $s'$. The second term reflects the wage gap due to differential pay of each age and education pair, weighted by the distribution of observable characteristics present in agricultural workers. Unlike the first term, this second component is solely affected by differential returns to age and education, and not by differences in composition. The third term accounts for the interaction between the the composition and return effects.

Figure 1.8 shows the result of this decomposition. In 1996, composition effects explain only a small share of the agriculture vs manufacturing gap throughout the sample period,

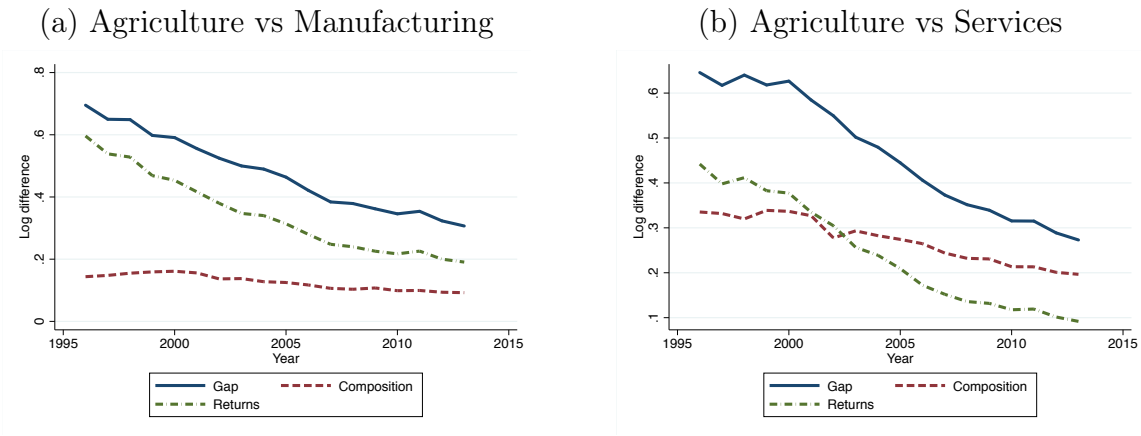
and they explain a larger share, but not all, of the services vs agriculture gap. This means differences in the age composition and educational attainment in each sector cannot account for most of the agricultural wage gap in the earlier period. In contrast, differential pay to these characteristics in each sector explains a larger share of the total gap level. Moreover, when looking at the evolution of this decomposition over time, most of the decline in the gap between agriculture and both manufacturing and services is driven by the steeper decline in the gap due to the estimated return coefficients. This implies that changes in the age distribution of the population or educational attainment do not drive the decline in the gap in Brazil. As seen in Table 2.1, educational attainment differences between agriculture and the rest of the economy were reduced over 1996-2013 in Brazil; nevertheless, this reduction alone explains little of the overall decline.

This limited role of age and education is present in spite of the lack of control for unobservable skill differences between education-age groups. It is likely that this omission overstates the role of compositional differences implied by the decomposition described. For instance, if workers with higher education are payed more not because of their education, but rather because of unobservable skills that are correlated with their education level, this correlation biases upward the share of the wage gap explained by these observable characteristics. Hence, to the extent that more highly payed age-education groups possess more highly valued unobservable skills, the share of the gap explained by observables above is an upper bound on the role of these characteristics. In appendix 1.8, the role of observables after controlling for worker fixed effects is estimated. Since individual workers' changes in age and education have little impact on their wages, controlling for unobservable fixed characteristics erases most of the role of observables in explaining the gap.

Therefore, to understand the magnitude of the wage gap and its decline, one must go beyond measuring observable differences in worker characteristics. That is, one must ask why workers with similar age and education are paid differently across sectors, why this

difference is part of an equilibrium, and why has this difference declined over time as the economy has developed.

Figure 1.8: Oaxaca decomposition




Note: Gap refers to the difference in mean log earnings between two sectors. Returns refer to the term $(\beta_t^{s'} - \beta_t^a)E(X_{it}^a)$ and composition refers to term $\beta_t^a(E(X_{it}^{s'}) - E(X_{it}^a))$ of the Oaxaca decomposition.

### 1.5.2 Unobservable characteristics

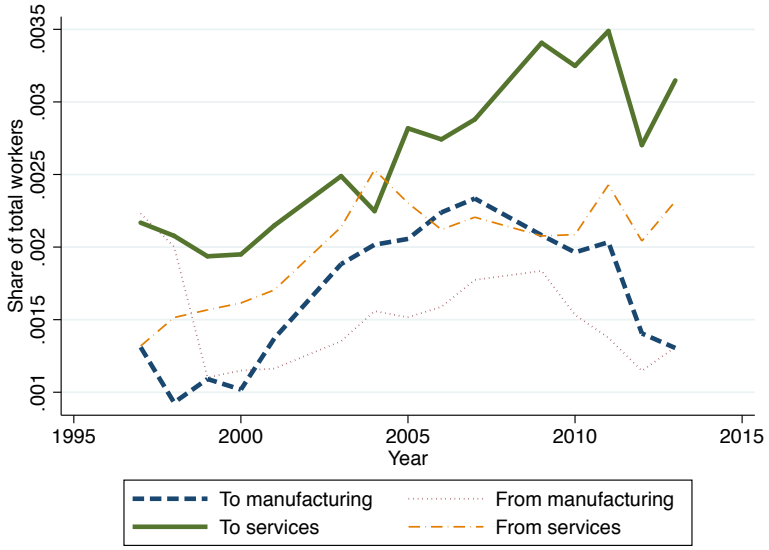
The role of differential returns emphasized above does not necessarily imply that workers in agriculture are intrinsically less productive or skilled. It might still be the case that their unobservable characteristics are similar to that of workers in other sectors but that these traits are simply paid differently in each sector.

Under perfectly competitive labor markets with fully mobile workers, every worker should move to the sector where he or she is paid the most. This process would eliminate any differences in pay among workers with similar —observed and unobserved— characteristics. This result is independent of any capital or technological limitations that are particular to each sector. However, compensating differential stories —where workers value sector-specific non-pay characteristics and are therefore willing to receive lower pay in some sectors— or mobility frictions can break this pattern. For instance, one can imagine a situation in which

workers are unwilling to pay a mobility cost from moving to industrial areas or one in which workers are unwilling to sacrifice the perks of employment conditions in agriculture. These stories are able to generate wage gaps within each age-education groups that are consistent with the differential returns observed in the previous section. There are therefore two types of competing types of stories in explaining the Oaxaca decomposition above. On one hand, agricultural workers may have a different composition of unobservable characteristics which makes them less valuable in the market. On the other hand, workers may be similar in the two sectors, but mobility frictions or compensating differentials may induce differential pay for each worker type.

In order to distinguish differential pay from compositional differences in unobservable characteristics, it is necessary to use the panel dimension of the dataset. In particular, I analyze the inter-sector flows of workers and test whether wages change substantially when a worker switches sectors. The two types of stories described above have different hypotheses about the behavior of wages among sector-switchers. If the wage gap is entirely due to differences in unobservable characteristics, then we should not expect wages to change substantially among workers that switch sectors. In contrast, if mobility frictions or compensating differentials are driving most of the wage gap between agriculture and the rest of the economy, we would expect workers that switch out of agriculture to experience large gains. In addition, since we do not observe hours, but only monthly earnings in each sector, a third possibility is that differences in hours worked are driving the gap. This can be considered as an extension of the compensating differential story where the hours worked are an additional non-pay component of work that is specific to each sector. Similarly to the compensating differential story, if the measured earnings gap in Brazil is driven by non-agriculture workers laboring for longer hours, we would expect these to be reflected in large earnings jumps among the workers that switch out of agriculture into other sectors.

Figure 1.9: Number of sector-switchers from and into agriculture




Note: Share of total employed workers that switch out of or into agriculture in any given year.

Using information on sector-switchers, I estimate the magnitude of wage changes from sector transitions, after controlling for time trends and changes in workers' observable characteristics. In order to study these switches, however, enough sector-switchers are needed to estimate these changes precisely. Figure 1.9 shows the share of workers that switch across sectors throughout the sample period. The small share of sector switchers (<1% of employed workers in a given year) would usually complicate the study of sector wage jumps using a small-sample panel dataset. However, because of the large number of workers covered the Brazilian data, this is not a problem. In any given year, there are over 180,000 formal workers who switch into and out of agriculture in Brazil.

To assess the magnitude of wage changes after controlling for differences in unobserved characteristics, the following worker fixed effect model is estimated

$$
\begin{aligned}
log(w_{it}) &= \beta_m^t * M_{it}\phi_t + \beta_s^t * S_{it}\phi_t \\
&+ \phi_t + \phi_i + F_t(education_{it}, age_{it}) + \varepsilon_{it} \qquad (1.1)
\end{aligned}
$$

where $M_{it}$ and $S_{it}$ are indicators for working in the Manufacturing and Services sectors, respectively; $\phi_t$ and $\phi_i$ are time and individual fixed effects; and $F_t(education_{it}, age_{it})$ is a set of interacted year, age and education dummies. Most importantly, sector indicators are interacted with time; therefore, the coefficients $\beta_s^t$ and $\beta_m^t$ reflect average earnings change from switching sectors from agriculture to both manufacturing and services in each year $t$. I will refer to these coefficients as sector premiums with respect to agriculture, of which there are $2*T$ in the model, where $T$ is the number of years in the sample. The model is estimated using all formal workers in Brazil from 1996 to 2013. In the baseline estimation of the model, the sector premiums are identified by workers who switch sectors during this period, and controls are estimated using information from all formal workers in the data.

The time series of both services premiums ($\beta_m^t$) and manufacturing premiums ($\beta_m^t$) are shown in Figure 1.10. A first takeaway from the figure is that wage differences estimated from switchers are much smaller than the overall wage gap. This is true throughout 1996-2013. For manufacturing, the average sector premium during 1996-2013 is 0.11 compared to the overall wage gap of 0.48 log points relative to agriculture. Similarly, for services, the average jump in wages is 0.08 compared to the mean total gap of 0.46 log points. Hence, sector premiums as a percentage of the total gap in a given year averaged 21% when comparing agriculture vs manufacturing and 16% when comparing agriculture to services. Potential gains from switching out of agriculture are significantly smaller than the overall agricultural gap. These shares suggests that the potential role of theories producing differential pay of similar workers across sectors is limited.

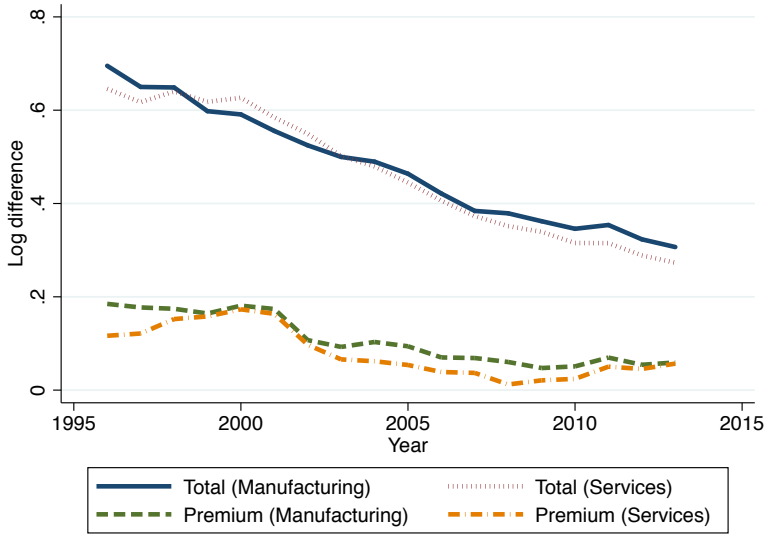
A key identification assumption of the model is that the error term must be orthogonal to the manufacturing and services dummies. This is violated if workers that switch out of agriculture are precisely the ones who would experience the largest wage jump from switching out of agriculture, which may certainly be the case. In a mobility frictions story, for example, it is precisely the workers who stand to gain the most from transitioning the ones who are willing to overcome this friction and move out of agriculture. Similarly, in a compensation differential story, workers only accept to move out of agriculture if compensated for the loss of non-pay benefits enjoyed in their original sector. These mechanisms, however, would bias our sector premium estimates upwards, so that $\beta_m^t$ and $\beta_s^t$ are upper bounds on the potential wage gains to be obtained from switching out of agriculture. To the extent that sector-switchers are the ones who stand to gain the most, this further depresses the role of differential pay stories in explaining the overall wage gap.

Another related concern is that the estimates might be affected by the inclusion of all workers in the estimation rather than just sector-switchers. Table 1.2 shows the average sector premium coefficients by period when the model in equation (1) is estimated using only sector-switchers and only transitions out of agriculture. If anything, a focus on switchers further lowers the estimates of sector premiums estimated in the baseline. Moreover, results do not appear to be driven by asymmetries from sector-switches. This might be a concern if switchers into agriculture are solely driven by improving job offers and these positive job changes counterweight large potential premiums from workers switching out of agriculture. This is no the case, as the model estimated solely on workers who switch out of agriculture yields lower results relative to the baseline. In fact, when performing an event-study of workers that exit agriculture (appendix 1.9), it is not entirely clear whether the wage jump from exiting agriculture is drastically different that the average gain expected from an extra year of experience working in any given sector. Potential gains from switching sectors appear to be even smaller when focusing exclusively on transitions out of agriculture.

The small sector-premiums described above may seem surprising. At first glance, they appear at odds with a literature that highlights the role of working environments —and firms in particular— in determining wage differences. For example, Card, Devicienti and Maida (2013), Menezes-Filho, Muendler and Ramey (2008), Barth et al. (2014) and others have highlighted the existence of a firm component of wages that is orthogonal to both observable and unobservable worker characteristics. Alvarez et al. (2016) analyze this pattern in Brazil and conclude that a reduction in wage dispersion has been accompanied by a compression of firm-specific components. The reason for this apparent discrepancy is that firm premium differences do exist between firms in Brazil, but the differences in average premiums between agriculture and other sectors are small when compared to the overall gap. Appendix 1.10 addresses the role of within sector heterogeneity in this context.

Altogether, the limited pay gains associated with the exit out of agriculture into other sectors limits the scope of stories that predict large wage jumps among workers who switch sectors. Stories generating differential pay of similar workers do not appear to be the main drivers of the gap. On the contrary, most of the wage gap stems from differences in the composition of unobservable —and to a lesser degree observable— worker characteristics in each sector.

Figure 1.10: Sector gaps relative to agriculture controlling for individual fixed effects




Note: Total refers to the difference in mean log earnings between each non-agricultural sector and agricultures. Sector premiums for services ($\beta_s^t$) and manufacturing($\beta_m^t$) are defined by equation (1). With the exception of 2008, coefficients are all statistically different from zero ($p < .01$).

## 1.6 How are compositional differences in characteristics generated?

The analysis above suggests that most of the wage gap —particularly in the late 1990s— is not due to differential pay of equally skilled workers between agriculture and non-agriculture. Instead, the wage gap appears to be largely driven by compositional differences in educational attainment and fixed unobservable characteristics between sectors. According to the results presented, a plausible mechanism for generating wage gaps must therefore achieve a very particular goal. It must generate wage gaps driven by large differences in worker char-

Table 1.2: Sector premiums relative to agriculture estimated using sector-switchers and all workers

| | 1996-2001 | 2002-2007 | 2008-2013 |
|---|---|---|---|
| **Services ($\bar{\beta}_s^t$)** | | | |
| All workers | 0.14 | 0.06 | 0.03 |
| Sector-switchers | 0.08 | 0.02 | 0.00 |
| Exiters from agriculture | 0.09 | 0.03 | -0.01 |
| **Manufacturing ($\bar{\beta}_m^t$)** | | | |
| All workers | 0.17 | 0.09 | 0.06 |
| Sector-switchers | 0.17 | 0.10 | 0.06 |
| Exiters from agriculture | 0.13 | 0.07 | 0.07 |

Note: Average of sector premiums, $\bar{\beta}_s^t$ and $\bar{\beta}_m^t$, over each 6-year period are presented. These are defined by equation (1), and they are estimated for three samples . All workers category comprise all formal workers between 18 and 55 years old in the RAIS. Sector switchers restrict the sample to workers that have switched into or out of agriculture at least once in each 6-year interval. Exiters from agriculture are defined as workers that have switched from agriculture to another sector at least once in each 6-year interval.

acteristics in each sector without giving rise to large differences in pay for similar workers in the two sectors.

Following the work of Roy (1951), recent papers have proposed the sorting of workers with sector-specific skills as a possible explanation of wage and productivity differences between countries, urban vs rural areas, and sectors.[7] This mechanism can generate inter-sector gaps driven by compositional differences in worker characteristics in a manner that is consistent with the empirical observations described. In this section, I test the explanatory power of worker sorting in explaining the wage gap level and its decline. I first assess the existence of sector-specific skills, which is a key assumption of these models. Motivated by this exercise, I then describe and calibrate a sorting model to show how large differences in mean wages between sectors can be generated as an equilibrium outcome of heterogenous workers with sector-specific skills freely choosing sectors. Finally, I show how growth in such an economy can explain part of the wage gap decline.

[7] e.g. Lagakos and Waugh (2013) and Young (2014).

### 1.6.1 Motivating a Roy Model: The existence of sector-specific skills

A basic premise of Roy models is the existence of occupation or sector-specific skills. In the context of the agricultural wage gap, a worker under this view has agriculture-specific skills and non-agriculture-specific skills, which determine the productivity of the worker when performing sector-specific tasks. By influencing labor productivity, sector-specific skills also determine the potential wage of the worker in each sector and influence his labor allocation decision between sectors. It is not clear, however, whether sector-specific skills exist at all.

Though I cannot test the presence of sector-specific skills directly, I can test an implication of sector-specific skills on wage changes among sector-switchers and the sector premiums described in section 1.5.2. In particular, I study to what extent are sector premiums driven by workers performing different occupations after switching sectors. In a world where workers have sector-specific skills, wage gains from transitioning out of agriculture into another sector should be more prominent when workers perform a different task in their new sector of employment. If, on the contrary, workers are equally productive regardless of the task performed, then wage changes from transitioning out of agriculture must be driven by other forces that are not necessarily related to an increase in labor productivity.

For example, consider a member of the cleaning staff of an agricultural firm who is considering switching out of agriculture. In a world where sector-specific skills exist, he has the potential to achieve a different level of productivity in the non-agricultural sector. That is, the possibility of performing new tasks (e.g. machinery operation, human-capital intensive tasks) that are fundamentally different from the ones performed originally can enable the worker to exhibit sector-specific skills, and therefore improve his productivity of his labor. This change in productivity can in turn induce a wage gain from transitioning sectors. In contrast, if the worker transitions out of agriculture but performs the same set of tasks related to his original cleaning job, we would expect gains to be more limited. Switching sectors

without switching occupations limit the realization of sector-specific skills and, therefore, potential wage gains under this view.

A measurable implication of this is that, to the extent that sector-specific skills are important in explaining wage changes among sector-switchers, we would expect sector premiums from 1.5.2 to be significantly reduced once we control for changes in occupation. If, on the contrary, wage changes are driven by differential pay for similar work, a cleaning worker switching between sectors but remaining a cleaning worker, would experience a gain in wages that is similar to the overall sector premium. In this second case, sector premiums should not be significantly reduced by controlling for occupations. To test which story is a more accurate description of sector wage premiums, I calculate premiums before and after controlling for occupations. The following equation is estimated

$$\begin{aligned} log(w_{it}) &= \gamma_m^t * M_{it}\phi_t + \gamma_s^t * S_{it}\phi_t + \phi_{occupation} \\ &+ \phi_t + \phi_i + F(education_{it}, age_{it}) + \varepsilon_{it} \end{aligned} \tag{1.2}$$

where coefficients $\gamma_m^t$ and $\gamma_s^t$ reflect the average differential pay of workers performing the same occupation in both pre and post-transition sectors, $\phi_{occupation}$ are occupation fixed effects at the three-digit classification level, and the rest of variables are defined as described in section 1.5.2. Similarly to the model outlined in the previous section, this model is identified by workers who switch sectors. The main difference of this approach, however, is that the coefficients $\gamma_m^t$ and $\gamma_s^t$ are identified using sector-switchers that do not switch occupations after they transition. In the data, several occupations are common to all sectors (e.g. cleaning, security services, drivers/messengers) and the model is therefore identified.

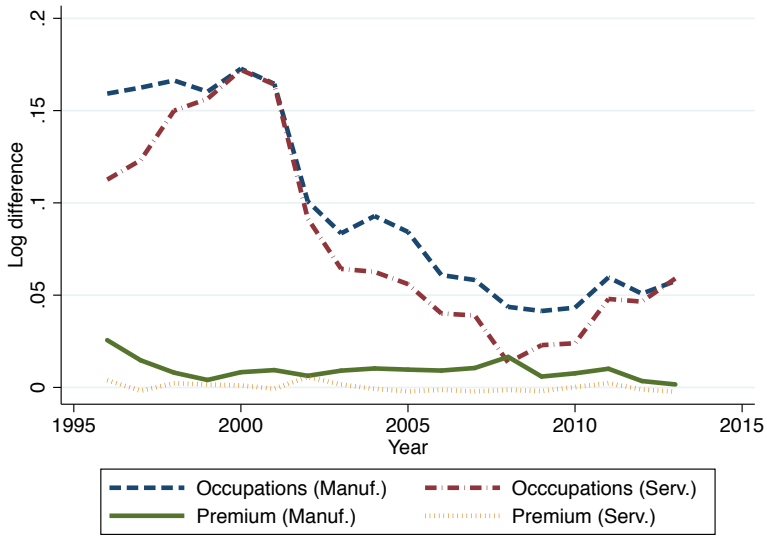
The specification above allows the decomposition of the sector premiums estimated in section 1.5.2 into two components: one stemming from occupation switches and another

reflecting differential pay for workers performing the same occupation in the two sectors.[8] Figure 1.11 shows the evolution of this decomposition over time. The decomposition is depicted comparing agriculture to both services and manufacturing separately with similar results. Throughout the period, most of the sector premium was explained by occupation differentials across sectors. In contrast, the differential pay component was relatively low throughout the period. In fact, the sector premium practically disappears by the end of the period when comparing services vs agriculture, and it is below .03 log points when looking at manufacturing vs agriculture.

The fact that sector premiums are much diminished once occupation differences are controlled for imply that sector premiums are attributable to the changes in occupation when transitioning sectors. This favors the existence of sector-specific skills, for the difference between pre and post transition can be nearly entirely attributed to the change in tasks performed by the worker after switching sectors. In terms of the example described, the results suggest that a worker in cleaning services in the agricultural sector is not expected to experience a significant gain in earnings when switching to a different sector unless he performs a different occupation. Instead, workers switching out of agriculture appear to get access to new types of jobs, and this accounts for most of the average jump in their wages. To the extent that sector-specific jobs imply the demonstration of sector-specific skills, this jump in wages among sector-switchers supports a Roy view of the world where workers have sector-specific abilities.

[8]The sector premiums estimated in section 1.5.2 can be written as $\beta^t_{s'} = \gamma^t_{s'} + (E_{s'}(\phi_{occupation}) - E_{agriculture}(\phi_{occupation}))$ where $s$ is each non-agricultural sector and $E_{s'}(\phi_{occupation})$ is the average of occupation fixed effects in each sector.

Figure 1.11: Sector premiums after controlling for occupational changes




Note: Occupations refer to the difference in mean occupation fixed effects between sectors.

Sector premiums for services ($\gamma_s^t$) and manufacturing($\gamma_m^t$) are defined by equation (2).

Premium coefficients are all statistically different from zero ($p < .01$).

### 1.6.2 A Roy model of selection with mobility frictions

Motivated by the empirical results shown above, I construct a Roy model to assess the explanatory power of worker sorting in explaining the agricultural wage gap. The following model borrows heavily from Lagakos and Waugh's (2012) framework, deviating from it in two ways. The first is the introduction of a friction that allows for differential pay of of similarly skilled workers. The second is the use of earnings information from workers that switch sectors to calibrate the distribution of idiosyncratic and unobservable productivity parameters of workers. In contrast to their framework, the level of the inter sector wage gap is not a calibration target of the model but rather an outcome that I evaluate the model against. In particular, I ask whether a wage gap driven by compositional differences in worker characteristics of the magnitude seen in Brazil can be explained by the worker sorting mechanism. I then assess whether growth in this environment can generate a wage

gap decline similar to one observed in Brazil during 1996-2013. The model's components are now described in detail.

#### 1.6.2.1 Preferences and endowments

There is a unit continuum of workers with unit mass, where each worker has identical preferences over agricultural and non-agricultural goods. Following other studies on structural change and the dual economy, workers have a subsistence requirement of $\bar{a}$ agricultural goods. Hence, preferences are non-homothetic and the share of expenditure on agricultural goods increases as income grows.[9] Preferences for each worker $i$ are given by

$$U(c_{a,}^{i}c_{n}^{i}) = log(c_{a}^{i} - \bar{a}) + \phi log(c_{n}^{i})$$

where $c_a^i$ and $c_n^i$ refer to consumption of agricultural and non-agricultural goods respectively and $\phi$ is a weight parameter that determines the relative importance of non-agricultural goods in consumption.

In addition, each worker is endowed with one unit of labor and sector-specific individual productivities $\{z_a^i, z_n^i\}$ drawn from a distribution $G(z_a, z_n)$ with support $[\underline{z}, \infty]^2$. A worker can freely choose to work in one of the two sectors but faces a cost $k$ if he decides to work in the non-agriculture sector.[10] This distortionary mobility friction will introduce compensation differences for workers that switch sectors as observed in the data. In this environment, workers maximize their income $y^i$, so that $y^i = max\{w_a(z_a^i), w_n(z_n^i) - k\}$ where $w_a(z_a^i)$ and $w_n(z_n^i)$ are the wages offered to worker $i$ in the two sectors. Thus, each worker faces the following budget constraint

$$p_a c_a^i + c_n^i \leq y^i$$

[9] See Gollin, Parente and Rogerson (2002) and Gollin, Parente and Rogerson (2004*a*) for a more detailed discussion on the role of the subsistence requirement in explaining decreasing agricultural output shares and its role in sector allocations.

[10] Since the focus of the quantitative exercise is analyzing the growth-induced exit of agricultural workers into other sectors, I focus on the mobility cost of going from agriculture to other sectors and not vice-versa.

where the non-agricultural good is set to be the numeraire and $p_a$ is the relative price of the agricultural good.

#### 1.6.2.2 Technology

There is a profit-maximizing representative firm in each sector with production functions given by

$$Y_a = AZ_a, Y_n = AZ_n$$

where $A$ is an economy-wide productivity parameter and $Z_s$ are the total effective units of labor employed by sector $s$. This second term is equivalent to the sum of individual worker productivities hired by the firm, or $Z_s = \int_{i \epsilon \Gamma^s} z_s^i dG$, where $\Gamma^s$ is the set of workers hired by sector $s$. Similarly, the number of workers employed by sector $s$ is given by $L_s = \int_{i \epsilon \Gamma^s} dG$. Labor productivity of a sector is therefore a function of the integral over the individual productivity parameters of workers employed in that sector.

#### 1.6.2.3 Competitive equilibrium

An equilibrium is determined by a relative price of the agricultural good $p_a$, wage functions $w_a(z_a), w_n(z_n)$, consumption decisions $c_a^i, c_n^i$ and labor allocations such that

1. Firms maximize profits in the two sectors given their technology.
2. Workers' labor allocations maximize their income.
3. Consumption allocations maximize utility subject to the budget constraint.
4. Labor, agricultural, and non-agricultural goods markets clear.

In a competitive market, the first condition requires firms to offer workers a wage equal to the value of their marginal product in their respective sector. These wages vary by both the

aggregate productivity factor, relative prices, and the idiosyncratic productivity parameter of each worker.

$$w_a(z_a) = p_a A z_a, w_n(z_n) = A z_n$$

Taking these wages as given, each worker decides to allocate their unit of labor to one of the two sectors. The second condition implies that a worker will chose to work in agriculture as long as the value of their marginal product in agriculture is more than his potential production in non-agriculture minus the mobility cost, or

$$A p_a z_a^i \geq A z_n^i - k$$

An implication of this is that higher relative agricultural prices, for a given non-agriculture productivity, lowers the minimum agricultural productivity required to stay in agriculture. Moreover, income is given by $y^i = max\{A p_a z_a^i, A z_n^i - k\}$ and consumption demand for both goods are given by,

$$c_a^i = \frac{y^i}{p_a(1+\phi)} + \bar{a}\frac{\phi}{1+\phi}, c_n^i = \frac{\phi}{(1+\phi)}(y^i - \bar{a}p_a)$$

The above holds as long as $y^i \geq \bar{a}p_a$. Otherwise, $c_a^i = y^i/p_a$ and $c_n^i = 0$. Intuitively, the consumption rule consists of allocating resources on agricultural goods until the minimum subsistence requirement is met and then distributing the remainder among the two goods according to the weight parameter $\phi$. Thus, as income grows, a lower proportion of income is allocated to the consumption agricultural goods.

Finally, the market clearing conditions require that good markets clear and that the labor employed in each sector is consistent with workers' labor allocations. This is

$$
\begin{aligned}
\int c_n^i dG &= \int_{i\epsilon\Gamma^n} (Az_n^i - k) dG \\
\int c_a^i dG &= \int_{i\epsilon\Gamma^a} Az_a^i dG \\
\Gamma^n &= \{i : Ap_a z_a^i \leq Az_n^i - k\} \\
\Gamma^a &= \{i : Ap_a z_a^i > Az_n^i - k\}
\end{aligned}
$$

Our main subject of study is the wage gap between agriculture and non-agriculture. In the model, this is determined by

$$
\frac{E(w_n)}{E(w_a)} = \frac{\int_{i\epsilon\Gamma^n} w_n(z_n^i) dG}{\int_{i\epsilon\Gamma^a} w_a(z_a^i) dG} = \frac{\int_{i\epsilon\Gamma^n} z_n^i dG}{p_a \int_{i\epsilon\Gamma^a} z_a^i dG}
$$

The wage gap is therefore the result of two main mechanisms. The first is the direct effect of the relative price, which affects the relative valuation of efficiency units, $(z_a^i, z_n^i)$, for the output produced in the two sectors. The lower is the relative price of agriculture, the lower is the relative value of agricultural output and hence the greater is the wage gap, holding composition of workers constant. The second mechanism is selection, which affects the productivity distribution of the sets of workers working in the two sectors $(\Gamma^a, \Gamma^n)$. The lower is the relative price of the agricultural good, the more people exit agriculture to work in the other sector. This process changes the composition of workers in each sector which can increase or decrease the gap in mean worker productivity between sectors. The wage gap level is therefore larger or smaller depending on the equilibrium effect of these two mechanisms. Moreover, as a country grows richer and the relative price of agriculture declines[11], the net effect of both price and composition effects on the wage gap over time

[11] As a country grows richer, non-homothetic preferences imply that a lower share of income is allocated to agricultural goods consumption and a lower relative price of agricultural goods. See Lagakos and Waugh (2013) for a detailed discussion.

is undetermined. Whether wage gaps decline or rise in this environment depends on the parameters of the economy.

#### 1.6.2.4 Calibration strategy

I now proceed to calibrate the model to answer two questions. The first is whether worker sorting can generate wage gaps that are of the same order of magnitude of the ones seen in the data. The second is whether sector-neutral growth can reduce the gap over time.

To do this, I calibrate the economy to Brazil in the earliest period of 1996-1997, when the gap was the largest. I then introduce growth in parameter $A$ that is consistent with real output growth rates observed in 1996-2013. Preference, production, friction and productivity distribution parameters are jointly estimated to match different moments of the data. Although all of these parameters interact in the model, each of them has stronger implications for particular moments. Below, I describe the relationship of each parameter to each moment and how these are calibrated.

**Preference parameters** Consistent with the literature using dual economy models with minimum subsistence requirements, preference parameters $\phi$ and $\bar{a}$ are calibrated to match two moments of the data that that relate to labor and output shares. The first is the share of workers in agriculture of 25% observed in 1996. The second is a long-run agriculture output share of 0.5%, which is the standard parameter used by Lagakos and Waugh (2013), Restuccia, Yang and Zhu (2008$a$) and other studies of structural change. Once calibrated, the minimum subsistence requirement is 23% of the average wage.

**Production and friction parameters** The technology parameter $A$ is set to 1 for the initial calibration in 1996. Later, when studying the effect of growth on the wage gap, changes in $A$ are calibrated to match total yearly real output growth in Brazil during 1996-2013. This implies a cumulative growth in $A$ of 105% by 2013. The sector-neutral nature of

the productivity parameter in this economy implies that the wage gap is solely dependent on endogenous price and selection effects in the model.[12]

In addition, the mobility friction $k$ is calibrated to match the average wage gain, in logs, from workers that switch out of agriculture. This was at its highest point (.13 log points) in 1996. The relative small transition wage gains relative to the overall gap described in section 1.5.2 imply that the magnitude of $k$ is relatively small at 20% of the minimum subsistence requirement and 4.6% of the average wage in the economy. This confirms the conclusion in the previous section: micro-data is not consistent with mobility frictions generating large pay differentials in sector-switching workers.

**The joint distribution of sector-specific worker productivities** The wage gap is affected by the exogenous distribution of individual productivity parameters. I calibrate these parameters using information from wage dispersion within each sector as well as information from workers that have worked in two sectors. First, I restrict $\underline{z}$ to match $\bar{a}$, so that every worker in the economy has a labor productivity endowment that is sufficient to both afford minimum subsistence and transition out of agriculture. This is, workers close to $\underline{z}$ will freely choose to stay in agriculture because of preferences, and not because of the lack of sufficient income potential to pay the cost $k$. In this way, the wage gap is entirely due to endogenous selection and prices, and not to the distribution's support parameter.

Since I do not observe all workers at all sectors, it is impossible to calibrate the distribution $G$ without imposing some structure. Non-parametric estimation is therefore not an option. Instead, following Lagakos and Waugh (2013), I allow workers to have dependent draws from sector-specific Fréchet distributions $X(z_a)$ and $Y(z_n)$ and restrict the joint distribution $G(z_a, z_n)$ to be a Frank copula resulting from the two primary distributions. This is

[12] Alternatively, one could introduce distinct growth rates by sector. The impact of differential growth on the wage gap, however, is similar to the one described in this paper. Sector-specific growth rates change the relative price of agriculture which induces the exit of workers out of agriculture and changes worker productivity composition in each sector. Regardless of whether growth is sector-neutral or not, wage gaps are driven solely by the sorting of workers across sectors.

$$
\begin{aligned}
G(z_a, z_n) &= C[X(z_a), Y(z_n)] \\
C[u, v] &= -1/\rho * log(1 + \frac{(e^{-\rho u} - 1)(e^{-\rho v} - 1)}{e^{-\rho} - 1}) \\
X(z_a) &= e^{-z_a^{-\theta_a}} \\
Y(z_n) &= e^{-z_n^{-\theta_n}}
\end{aligned}
$$

The sector-specific distributions have dispersion parameters $\theta_a$ and $\theta_n$, which control the within sector variance of the productivity distribution in agriculture and non-agriculture. These are calibrated to match the variances of log earnings in agriculture and non-agriculture of .33 and .63 respectively. The calibrated parameters that match these log variances are $\theta_a = 1.2$ and $\theta_n = .51$.

Besides the transparent mapping that exists between these parameters and wage dispersion, there are two other reasons why a Fréchet shape is a sensible choice to model sector-specific distributions. First, the Fréchet distribution is a special case of the extreme value distribution; therefore, the marginal Fréchet distribution of a particular sector can be interpreted as the distribution of the maximum draw from a set of productivity distributions within that sector. For example, this may represent the maximum productivity draw out of a series of jobs that are available within the manufacturing or agricultural sectors.[13] Second, the shape of of the distribution, with greater mass at lower productivity parameters and fat tails, resembles the within sector distribution of both raw wages and unobservable worker characteristics observed in the data.

To form a joint distribution out of the two sector-specific marginal distributions, a Frank copula is used. The advantage of using this copula is that it allows the degree of dependence in the two distributions to be controlled by a single parameter $\rho$. Along with $\theta_a$ and $\theta_n$, this parameter is calibrated to match the fraction of workers that exit agriculture during 1996-

[13]By the extreme value theorem, the maximum of independent draws from any distribution converges to an extreme value distribution. The Fréchet is an example of these distributions.

1997 (4% of agricultural workers) when the economy grew 2.2% in real terms. Intuitively, for given dispersion parameters $(\theta_a, \theta_n)$, $\rho$ controls the amount of workers close to the labor allocation indifference condition ($Ap_a z_a^i = Az_n^i - k$). A growth-induced change in prices pushes a larger or smaller share of workers out of agriculture depending on the mass of workers that are close to indifferent in the base year. The resulting parameter from this calibration is $\rho = 12$, which implies a linear correlation of 35% between the sector-specific productivity parameters $z_a$ and $z_n$. Importantly, no difference in mean productivity between sectors is assumed in the calibration of the joint distribution. The agricultural wage gap is therefore not a calibration target but an outcome of the model.

### 1.6.3 The explanatory power of the sorting mechanism

The model generates both cross-sectional and inter-temporal predictions. In the 1996 cross-section, the predicted wage gap by the calibrated model is .43 log points. This magnitude is 66% of the wage gap observed. The calibration using data on sector switchers and within sector variances is, therefore, consistent with large compositional differences in worker skills arising from sorting, but working sorting by itself does not explain all of the gap.

The reason behind this shortcoming is related to the magnitude of $\rho$. In the model, a higher correlation in the sector-specific productivity parameters generates larger wage gaps between agriculture and non-agriculture. Intuitively, under perfect correlation of productivity draws $z_n$ and $z_a$, the most skilled workers would work in non-agriculture, while the worst skilled workers would work in non-agriculture. In that setting, the lowest paid worker in non-agriculture earns more than the highest paid worker in agriculture. Thus, the perfect correlation produces a large inter-sector wage gap. In contrast, with imperfect correlation, some workers are much more productive in agriculture than in non-agriculture. These workers sort into agriculture. In equilibrium, imperfect correlation produces some agricultural workers that earn more than their peers in non-agriculture. A lower correlation parameter $\rho$, therefore, reduces the wage gap.

In the model, the correlation parameter $\rho$ required to match the total wage gap between sectors would generate too few workers that are close to indifferent between the two sectors and, therefore, would underestimate the amount of workers who choose to switch. In order to match the wage gap perfectly, an additional source of wage heterogeneity would be required. For instance, the inclusion of transitory shocks in a dynamic setting would increase the share of sector switchers and ease the effects required from productivity dispersion parameters. Larger wage gaps can be generated if the target share of switchers to be produced by the sorting mechanism is lowered. This additional uncertainty is beyond the scope of this paper, as we are focused on sorting stemming from fixed unobservable skill differences between workers. However, it is worth noticing that, even when restricting the sources of wage heterogeneity in the model, this mechanism can generate a wage gap that is of a similar order of magnitude to the one observed in the data.

Once the economy is calibrated to 1996, I examine whether growth —here introduced as an increase in the sector-neutral productivity parameter $A$— generates a wage gap decline. Growth has two effects. On one hand, the increased income decreases the demand for agricultural goods consumed in the economy relative to non-agricultural goods. This is a direct consequence of the subsistence requirement present in preferences. The reduction in relative demand induces a lowering of the relative price of the agricultural good, which decreases the market value of the marginal product of agricultural workers. This in turn depresses relative wages in agriculture which widens the gap. On the other hand, the price changes cause the exit of workers out of agriculture into the manufacturing sector. The transitioning workers lower the average productivity of the non-agriculture sector, which reduces the gap in average wages between the two sectors. As mentioned before, the net effect of these two forces depends on the economy's parameters.

Figure 1.12 shows the comparison of the wage gap decline, in absolute value and relative to the 1996 wage gap level, produced by the model and the data. Qualitatively, the model is indeed successful in generating a decline. Over the 18 year period, the wage gap in the model

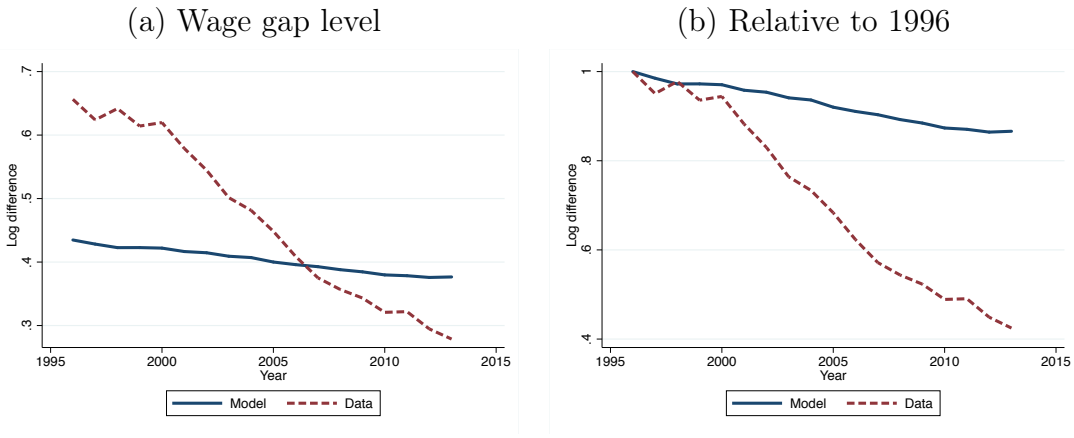
declines by 15% relative to the the start of the period. This relative decline is 58% in the data. In absolute magnitude, the understated gap in 1996 declines .06 log points, which is small relative to the overall decline in the data of .38 log points. The worker sorting mechanism is therefore consistent with the wage gap decline; nonetheless, the limited magnitude of the predicted wage gap reduction leaves a substantial share of the decline unexplained.

One concern is whether the magnitudes of the results are driven by the parametrization assumptions behind the joint distribution of sector-specific productivities. Parametrization is necessary since we are not able to observe the productivity of all workers across all sectors, but it is hard to determine whether this is a reasonable approach to study the sorting of agricultural workers into other sectors as an economy grows. Though far from being a test of validity from the parametrization assumptions, one can assess whether the shape of the distribution is reasonable by observing the rate of exit of workers out of agriculture. For instance, if the joint distribution worker productivities $G$ contains discontinuities or a radically different shape from the calibrated distribution, we would expect the relationship between growth and the exit of workers out of agriculture to behave different in the data and the model. The evolution of the share of workers employed in agriculture in the data and in the model are shown in Figure 1.13. Similar to the data, the model generates a gradual exit of agricultural workers that has a similar pattern to the one observed in Brazil. The model predicts a decrease of 14% in the share of workers in agriculture so that it reaches 11% in 2013. This is slightly lower than the 14% share observed in the data by the end of the period. Remarkably, however, the qualitative direction and linear relationship between growth and the share of workers in agriculture is similar in both the data and the model. The pattern resemblance is partially due to the calibration strategy used to match growth and the exit of workers, since the parameters are matched to the share of workers exiting agriculture in 1996-1997. However, the relatively good fit of this relationship over the entire 18 year pattern is somewhat surprising. The sensible magnitudes in the number of workers

that exit agriculture speak in favor of the validity of the shape assumptions and calibration of the joint productivity distribution.

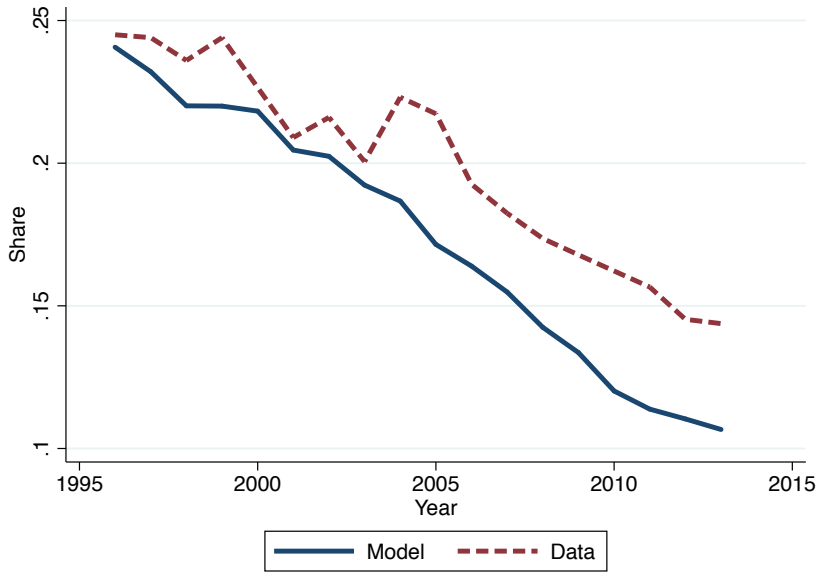
Overall, the sorting mechanism of the model captures relatively well the movement of workers out of agriculture, the wage gap levels, and the qualitative decline in the agricultural wage gap. However, the calibration using micro-data suggests that other complimentary mechanisms must also be at play to fully explain the magnitude of these patterns.

Figure 1.12: Wage gap




Note: Panel (a) shows the difference in the mean of log wages between sectors . Panel (b) shows the relative decline of in the wage gap relative to the level observed in 1996 in both the data and model output.

Figure 1.13: Share of workers in agriculture




Note: Share of workers from data includes formal and informal workers.

## 1.7 Conclusion

The large wage gaps between agriculture and other sectors are mostly driven by compositional differences in unobservable worker characteristics. In accordance with cross-country patterns on inter-sector productivity and wage gaps, these differences declined gradually in Brazil since the late 1990s as the country became richer. This paper argues that a Roy model of worker sorting based on sector-specific skills is consistent both with the magnitude of the gap and its decline as a country grows. Moreover, the mechanism suggests that the fall in the wage gap may be an integral part of the development process, where compositional differences between sectors are reduced over time as workers transition out of agriculture.

Nonetheless, a significant share of the gap remains unexplained, and the model does not speak to the steeper decline in value added per worker observed both in Brazil and in cross-country studies. Exploration of additional mechanisms is therefore needed, and the evidence presented in this paper can guide the design of future wage gap models. Specifically,

the results show that pay differences for workers with similar skills are relatively small when compared with the total wage gap. This finding discourages models that generate large pay differences for similar work in different sectors. For instance, mobility frictions or compensation differentials that induce large gaps in wages per efficiency units between agriculture and non-agriculture would predict large wage gains from switching sectors that are at odds with the data. Therefore, complementary mechanisms to worker sorting that attempt to rationalize the wage gap level and its decline must, at the very least, produce large differences in average pay between sectors without producing relatively large wage gains for workers that switch out of agriculture.

## 1.8 Appendix. The gap explained by observable differences after controlling for unobservable characteristics

Differences in pay across sectors for each age and education group can potentially be explained by differences in the composition of unobservables. It might be the case that, for a a given age and education level, workers in one sector are fundamentally more productive than others. Workers certainly differ in dimensions other than age and education, and the literature has emphasized the importance of unobservable skill heterogeneity in explaining wage differences. Thus, an added challenge when comparing wages in different sectors is understanding how these unobserved individual skills and human capital play out differently in each sector. In order to analyze the relative importance of observed and unobserved worker characteristics, I estimate the following model using the panel dimension of the dataset.

$$log(w_{ist}) = \phi_{is} + F_{st}(education_{ist}, age_{ist}) + \varepsilon_{ist} \tag{1.3}$$

where $\phi_{is}$ is the fixed effect of a worker in sector $s$.[14] This specification controls for worker unobservables that are fixed over time when estimating differential compensation of age and education levels. As before, $F_{st}(education_{ist}, age_{ist})$ is a sum of education, age and year interacted dummies. Unlike the baseline specification in section 1.5, the identification of age and education returns comes from the few workers that switch age and education categories throughout the period of study. I then conduct a modified Oaxaca decomposition[15] using this age and education compensation function. This decomposition is given by

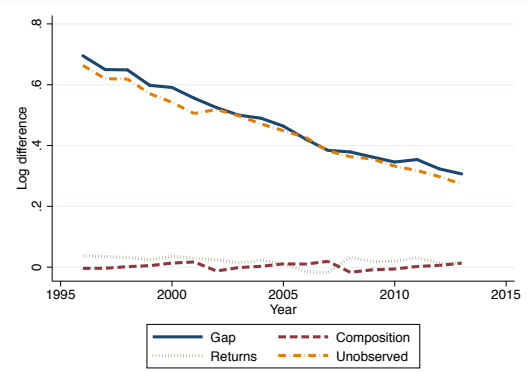
$$\begin{aligned}
\Delta(E(log(w_{ist}))) &= \Delta E(\phi_{is} + \varepsilon_{ist}) \\
&+ \beta_t^s(E(X_{it}^s) - E(X_{it}^a)) + (\beta_t^s - \beta_t^a)E(X_{it}^a) \\
&+ (E(X_{it}^s) - E(X_{it}^a))(\beta_t^s - \beta_t^a)
\end{aligned}$$

The first term reflects the wage gap between sectors due to unobservables, the second term accounts for compositional effects, the third term measures wage gaps due to differential pay to education characteristics, and the last term is the interaction effect. Figure 1.14 shows the results of this decomposition for both manufacturing and services. Once we control for fixed unobserved characteristics, differential pay across sectors of each education-age group explains little of the gap levels.

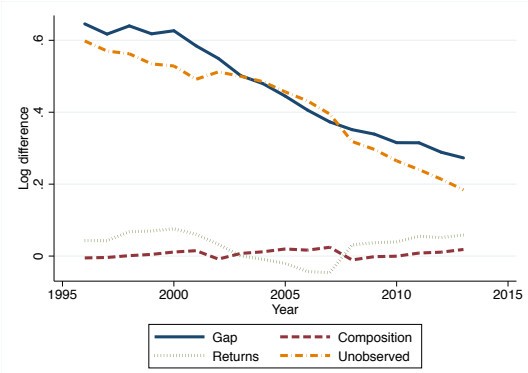
[14] Since there are relatively few sector switchers these results hold even when restricting $\phi_{is} = \phi_i$.
[15] From Oaxaca (1973).

Figure 1.14: Oaxaca decomposition controlling for unobservables

(a) Agriculture vs manufacturing




(b) Agriculture vs services




Note: Gap reference to the difference in mean log earnings between two sectors. Returns refer to the term $(\beta_t^{s'} - \beta_t^a)E(X_{it}^a)$ and composition refers to term $\beta_t^a(E(X_{it}^{s'}) - E(X_{it}^a))$ of the Oaxaca decomposition. Unobservables refer to the difference in mean worker fixed effects,$\Delta E(\phi_{is})$.

## 1.9 Appendix. Event-study of transitions into and out of agriculture

This section adopts an event-study framework focusing only on workers that switch out of agriculture, restricting the analysis to workers that were employed in the pre and post transition sector for at least two years. The following equation is estimated for transitions out agriculture into both services and manufacturing.

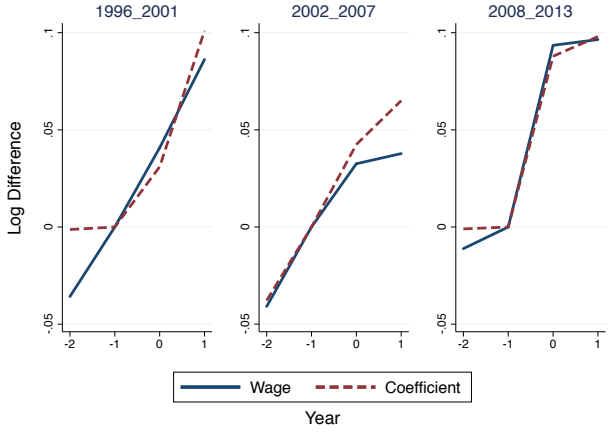
$$log(w_{ist}) = \sum_{j=-2}^{1} \gamma_j + \phi_t + F_t(education_{ist}, age_{ist}) + \varepsilon_{ist}$$

Figure 1.15 shows the results of the transition impact coefficients, $\gamma_j$, as well as the pre and post-transition raw wages. As before, the wage jumps from transitions are much smaller than the magnitude of the overall gap. The average wage boost in the year after switching out of agriculture ranges from 0.01 to 0.1 compared to the overall gap of 0.3-0.7. These small

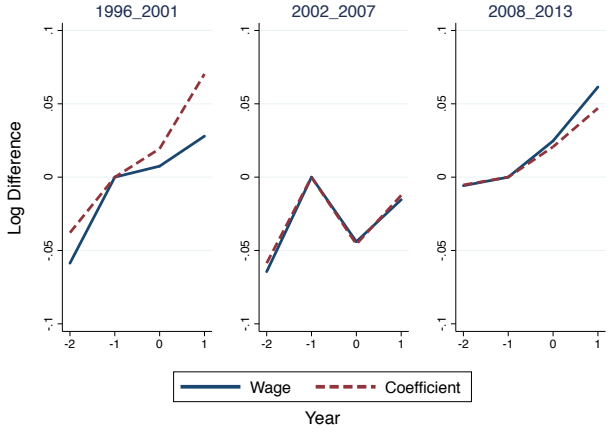
magnitudes are persistent for workers that switch out of agriculture into both services and manufacturing. Moreover, when looking at the yearly changes in the pre and post-transition years, it is not entirely clear the whether wage jumps in transition years are of substantially different magnitudes than from the typical within sector gains in any given year. It can be concluded that the focus on any particular switches does not alter the conclusion of limited differential pay of similar workers across sectors.

Figure 1.15: Transitions out of agriculture

(a) From agriculture to manufacturing




(b) From agriculture to services




Note: Raw mean wages before and after transitions out of agriculture are shown. Year 0 refers to the last year worked in the pre-transition sector and Year 1 refers to the first year in the post-transition sector. Year 1 refers to the first year in the post-transition sector. Coefficients $\gamma_j$ reflect wage changes from transitions after controlling for time fixed effects and changes in age and education.

## 1.10 Appendix. Within sector heterogeneity

In order to assess within sector heterogeneity, at the firm level, in explaining the agricultural gap, I follow the methodology described in Alvarez, Engbom and Moser (2016) to decompose wages into firm, individual and worker characteristic components, or

$$log(w_{ist}) = \phi_{J(i,t)} + \phi_i + F_t(education_{ist}, age_{ist}) + \varepsilon_{ist}$$

where $\phi_{J(i,t)}$ is a time varying firm-effect for firm $J$ where worker $i$ is employed at time $t$, and the other variables are defined as described in section 1.5. This estimation parallels the one conducted in equation (1) but it estimates average differences at the firm-level rather than at the sector level. As shown by Abowd et. al. (1999), worker and firm effects can only be separately identified within a set of firms and workers connected through the mobility of workers. Following a similar approach to Alvarez, Engbom and Moser (2016) and Card, Devicienti and Maida (2013), I find the largest set of firms connected through mobility of workers. I compute this set using a graph theoretical algorithm included in the MatlabBGL library. Because of the high mobility of workers in each of the sub-periods, these connected sets contain more than 95 percent of all workers and more than 70 percent of all firms.

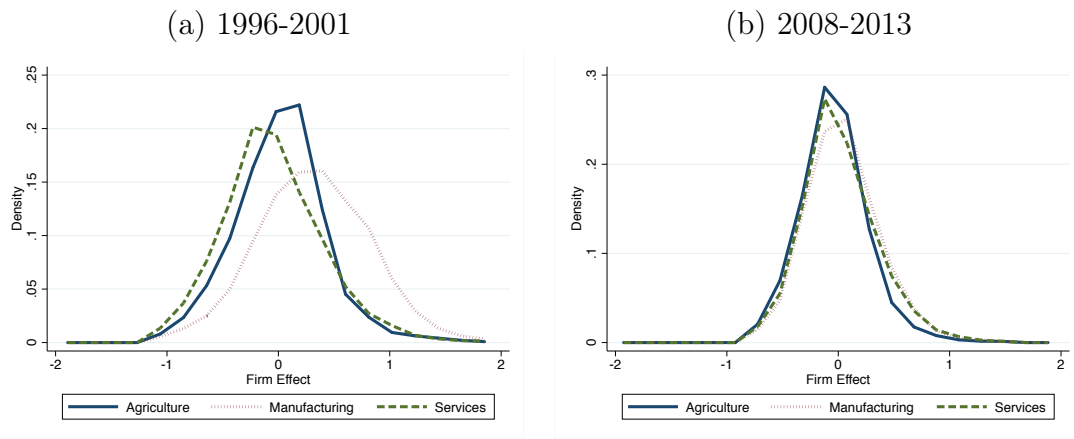
To estimate the model, I have to solve the normal equations. However, the model contains hundreds of millions of worker and firm effects in every period, which poses a computational challenge. I follow Card et. al. (2013) in approximating the system of normal equations without actually inverting the matrix using an iterative gradient method. By doing this, I obtain a vector of estimated firm effects for each period.

Figure 1.16 shows the distribution of firm-level premiums $\phi_{J(i,t)t}$ for the three sectors in the economy for the 1996-2001 and 2008-2013 periods. As mentioned before, agricultural firms pay lower average premiums relative to agriculture and services. However, there is great heterogeneity and overlap between the different sectors which cause the mean effect between sectors to be small. In fact, even in the early period when sector premiums used to be higher,

there is great overlap in the distribution of firm level premiums between agriculture and other sectors. For example, manufacturing firms did pay a premium relative to agriculture in the early periods, but there were also better agriculture firms that paid better premiums than manufacturing firms. By 2008-2013, the differences were greatly reduced but a significant wage gap remained.

Thus, consistent with the literature on the importance of firms in wage setting, firm premiums exist in Brazil, but mean differences in mean firm effects in relation to other sectors are much smaller than the overall wage gap. Since we are concerned with inter-sector differences, this observation motivates the study of sectors as uniform entities and a focus on worker-side mechanisms in explaining the between-sector wage gap magnitude and its decline.

Figure 1.16: Distribution of firm effects ($\phi_{J(i,t)}$) by sector



# Chapter 2

# Firms and the Decline of Earnings Inequality in Brazil

## 2.1 Introduction

Brazil has experienced a large reduction in earnings inequality since the mid-1990s. This came after decades of Brazil being infamously known as the most unequal country in Latin America, which itself ranked among the most unequal regions in the world. While the decline of earnings inequality in Brazil resembles other Latin American economies' experience during this period,[1] it stands in stark contrast to that of the U.S. and many developed countries, which saw inequality steadily increasing over the past two decades.[2] To investigate the sources of this decline, in this paper we decompose Brazil's earnings inequality evolution into changes in firm and worker characteristics on the one hand, and the returns to such characteristics on the other.

[1] See Lopez and Perry (2008) and Tsounta and Osueke (2014).

[2] Using administrative data, Kopczuk, Saez and Jae (2010) document in detail earnings inequality trends for the U.S., while chapter 8 of Atkinson and Bourguignon (2015) discusses inequality trends in middle- and high-income countries.

Guided by recent research, which suggests that firms are an important determinant of earnings dispersion in a number of high-income countries,[3] we decompose the sources of Brazil's inequality decline by exploiting a large administrative linked employer-employee dataset containing information on over one billion job spells between 1988 and 2012. By linking individual workers to their employers and tracking both over time, we are able to separately identify the contributions of firm- and worker-specific factors towards worker's pay (Abowd, Kramarz and Margolis, 1999). Applying this econometric framework repeatedly within overlapping subperiods, we decompose the overall inequality decline into changes between firms and between workers. We, we investigate the link between firm performance and the firm component of pay by linking our first stage estimation results to another confidential dataset containing detailed information on the characteristics of hundreds of thousands of Brazilian firms between 1996 and 2012.

We uncover three main results. First, firms played an important role in the decline in earnings inequality in Brazil over this period, explaining 45 percent of the fall in the variance of log earnings between 1996 and 2012. Compression in worker fixed effects explains an additional 24 percent of the decline, with the remaining part being attributed a decline in the covariance between worker and firm fixed effects and the residual. As worker heterogeneity is most important in levels, the compression in firm-specific pay components contributed more than proportionately towards Brazil's inequality decline.

Second, changes in the link between firm performance and pay accounts for a significant fraction of the compression in the firm component of workers' earnings. We first show that a substantial share of the variation in the firm component of pay is explained by differences in observable firm characteristics, with more productive firms paying more. Moreover, more than half of the decline in the firm component is accounted for by observable firm characteristics. All of this decline is driven by a weakening pass-through from productivity to pay, and none is due to firms becoming more similar in observable characteristics. A weaker

[3]See Abowd, Kramarz and Margolis (1999); Card, Heining and Kline (2013); Barth et al. (2014); Bloom et al. (2015), and Mueller and Ouimet (2015).

link between observable firm characteristics and worker pay thus explains 25 percent of the overall fall in the variance of log earnings over this period.

Third, a decline in the return to measures of ability such as experience and education explains a sizable share of the fall in the variance of the worker component of pay. In levels, age, education and occupation explain around 40 percent of the variance of the worker component of pay. However, we observe no compression in the underlying distributions of such characteristics over time. Instead, the inequality decline is driven by a rapid fall in the returns to observable measures of worker ability, especially the returns to education. We find that lower returns to education explain 13 percent and lower returns to age explain three percent of the overall fall in the variance of log earnings over this period.

This decomposition of the sources of Brazil's earnings inequality decline informs our understanding of various commonly proposed stories explaining the decline. On the worker side, our results do not support a widely-held belief that changes in educational attainment accounted for a significant share of Brazil's inequality evolution over the period.[4] While educational attainment increased over this period, we find that all else equal this resulted in a small net *increase* in inequality. We reach similar conclusions regarding the implications of changes in the age structure of the workforce. On the firm side, a reading of the existing literature would suggest that trade dynamics and the productivity evolution during this period could have been important drivers behind changes to the earnings distribution.[5] Yet, in line with U.S. trends, we find that also the Brazilian productivity distribution grew more dispersed over this period.

Our findings thus pose a challenge to candidate explanations for the decline in earnings inequality over this period. Our results suggest that a theory of the inequality decline needs to be consistent with the following three facts: (i) firm-level changes explain a significant share of initial inequality levels and an even larger share of its decline; (ii) a weaker pass-

[4]See for example Barros et al. (2010).

[5]See for example Dunne et al. (2004) and Faggio, Salvanes and Van Reenen (2010), who conclude that parts of the increase in earnings inequality in the U.S. can be explained by widening dispersion in the firm productivity distribution.

through from firm productivity to worker pay was a key driver behind the declining dispersion in the firm component of pay; and (iii) a lower return to worker ability explains a significant share of the decline in the worker component of pay. We conclude that changes in pay policies rather than in worker and firm fundamentals played an important role in the decline in inequality in Brazil during this period.

**Related literature.** With the current project we contribute to three broad strands of the literature. First, we provide a decomposition of earnings into worker and firm heterogeneity in a developing economy. With the increasing availability of large, administrative matched employer-employee datasets, a recent literature has started to examine the role of firms in wage determination. The first paper to make use of such large, linked employer-employee datasets to jointly study the role of worker unobservables and firms for pay is Abowd, Kramarz, and Margolis (1999, henceforth AKM), who study the role of firm and worker heterogeneity for wage inequality in France. They find an important role for firms in generating earnings inequality. Similar conclusions using the same methodology have been reached among others for the state of Washington in the U.S. (Abowd, Creecy and Kramarz, 2002), Denmark (Bagger, Sorensen and Vejlin, 2013), Austria (Gruetter and Lalive, 2009), and Germany (Card, Heining and Kline, 2013). The last paper is close to our methodology of applying the AKM framework in overlapping subperiods to study changes in wage determinants over time. They find that increasing dispersion in the firm-specific component of pay contributed significantly to rising earnings inequality in West Germany. Bloom et al. (2015) highlight firms as an important driver behind the increase in U.S. labor earnings inequality since 1980, but do not employ the AKM methodology to control for sorting of highly-paid workers into high-paying firms.

Second, we highlight the sources of changes in the worker and firm component of inequality by studying the role of changes in observable characteristics versus changes in the return to such characteristics. AKM study the link between firm observable characteristics and the

estimated firm effects, but do not focus on changes over time to this relationship. Menezes-Filho, Muendler and Ramey (2008) study the link between firm characteristics and wages in Brazil in the cross-section of linked data on worker earnings and firm characteristics in Brazil's manufacturing and mining sectors. Bagger, Jesper and Holloway (2014) investigate the role of labor misallocation in driving the positive correlation between labor productivity and wages at the firm using Danish data. Card, Cardoso and Kline (2015) study the degree of rent-sharing in Portugal with a particular emphasis on gender differences in profit participation and the allocation of workers across firms.

Third, we add to the literature on the sources of the decline in inequality in Brazil by providing a comprehensive decomposition of the decline. Many previous papers have studied the role of specific, isolated mechanisms in Brazil's inequality decline over the last two decades. For example, Ulyssea (2014) considers the role of worker flows between the informal and formal sectors. In related work, de Araujo (2014) studies the role of labor adjustment costs in propagating wage inequality in a frictional search framework. Dix-Carneiro and Kovak (2015) analyze the long-lasting impact of industry-specific tariff cuts in the presence of wage-equalizing migration. Barros et al. (2010) use Brazilian household data to study inequality trends since 1977 and decompose the decline in labor earnings inequality in Brazil since 1990. Given their data and method, those authors conclude that the inequality decline was in equal shares driven by education reform and labor market integration. Medeiros, Souza and de Castro (2014) use administrative tax return data to study the evolution of top income inequality in Brazil from 2006–2012, but they cannot distinguish between the role played by worker versus firm characteristics during that period. Using linked employer-employee data, Lopes De Melo (2013) conducts a static decomposition of earnings inequality levels in Brazil's formal sector into components due to firms and workers. Helpman et al. (2016) use the same worker-level data to show that a significant share of overall wage inequality is due to between-firm differences and that Brazil's trade liberalization starting in the late 1980s led to increasing between-firm earnings inequality.

**Outline.** The rest of the paper is structured as follows: Section 2.2 provides an overview of the main institutional changes and macroeconomic trends affecting Brazilian labor markets from 1988 to 2012. Section 2.3 summarizes the administrative datasets used in our empirical analysis and discusses sample selection and variable definitions. Section 2.4 provides descriptive statistics on trends in earnings inequality in Brazil during this time. Section 2.5 introduces the empirical framework we use to decompose the variance of log earnings into a worker and firm effect as well as the subsequent regressions we run to link these estimates to worker and firm fundamentals. Section 2.6 presents our main empirical results as well as checks on the validity of our empirical framework. Finally, Section 2.7 summarizes our key findings and concludes.

## 2.2 Institutions and macroeconomic trends in Brazil

During our period of study, Brazil resumed democratic elections (1989), ended a decade of hyperinflation (1994), and inaugurated two decades of sustained economic growth—between 1996 and 2012 real gross domestic product grew by 2.3 percent per year on average. In this section, we discuss some of the institutional changes that could have have affected inequality during this period, including labor regulation, trade liberalization and social policy.

Brazil had a highly regulated labor market before reforms started in the late 1980s. For instance, since 1965, a national Wage Adjustment Law mandated yearly wage increases for all workers in the economy and dismissal costs were high. After the transition to civil rule and the signing of a new constitution in 1988, flexibility in labor markets was further affected by firing penalties and an increased power of labor unions, which gather about a quarter of employed formal workers in Brazil.[6]The Wage Adjustment Law was finally abandoned in 1995, introducing a period of greater flexibility and less regulated wage-setting practices.Further legislation in 1997–1998 eased restrictions on temporary contracts and lowered dismissal bar-

[6]Bioto and Marcelino (2011) argue that there has been an uptake in labor strike activity in Brazil since the year 2000.

riers. Subsequently, formal employment increased by around five percent and unemployment fell from 10 percent in 2000 to around six percent in 2011 (World Bank, 2015). The overall labor participation rate has remained stable at 73–75 percent over this period.[7]

Hyperinflation also encouraged the adoption of automatic wage adjustment practices. From 1980 to 1989, yearly inflation averaged 355 percent, which was followed by a yearly average of 1,667 percent between 1990 and 1994 (World Bank, 2015). As a result, wage indexation to the minimum wage became the norm, with labor payments being adjusted first annually and then on a monthly basis proportionately to the previous period's realized inflation rate. In 1994, hyperinflation finally subsided with the introduction of the "Real Plan". This ambitious stabilization program introduced a gradual float of the local currency, tightened monetary and fiscal policy, and lowered inflation below two-digits.

In parallel to monetary stabilization, Brazil also undertook trade liberalization reforms during this period. Starting with initially high import tariffs that had substituted import bans from the previous decade, a series of trade liberalization bills in the late 1980s eliminated selected tariffs and eradicated quantitative import controls. When social democrat Fernando Henrique Cardoso became president in 1995, he strengthened this agenda with a reduction of tariff and non-tariff trade barriers to one tenth of their levels in 1987 (Pavcnik et al., 2004). The opening up to trade over the last 25 years has been frequently cited as a major contributor to the country's growth in total factor productivity (TFP)(Ferreira and Rossi, 2003; Ferreira, Leite and Wai-Poi, 2007; Moreira, 2004; Muendler, 2004; Córdova and Moreira, 2003). In addition, Helpman et al. (2016) argue that trade reforms contributed to the rise in income inequality seen in the late 1980s and early 1990s, and later to the start of the decline in wage dispersion in 1995.

Health, education and other social programs began expanding during the late 1990s, a trend that strengthened once the left-wing Workers' Party ascended to power in 2003. It doubled social expenditure as a fraction of GDP and, although it remains less than one

[7]Labor force as a percentage of total population aged 15–64, from OECD Employment and Labor Market Statistics.

percent, it is often portrayed as an important contributor to the reduction in household income inequality.[8] The reach of the public cash transfer program, Bolsa Família, increased to cover 11 million families in 2006, which comprised nearly 25 percent of the total population (Barros et al., 2010). Education spending increased reaching 5.5 percent of GDP in 2009 (compared to 3.5 percent in 2000 and 5.7 percent among G20). As we discuss in Section 2.4 this is reflected in a rapidly rising share of the labor force with a high school degree. Moreover, the quality of education relative to other countries, as measured by the international PISA scores, has also improved, with Brazil having the greatest increase in mathematics among 65 countries since 2003 (OECD, 2012).

The Worker's Party complemented social policies with minimum wage increases above the previous upward trend. Within their first year in office, they established a 20 percent increase in 2003 and continued to implement yearly increases averaging over 10 percent during the next 10 years. As a result, the minimum to median wage in Brazil increased from around 34 percent in 1996—similar to U.S. levels—to over 50 percent, which is close to the level in France. Engbom and Moser (2015) argue that this large increase in the minimum wage can explain a significant fraction of the reduction in earnings inequality in Brazil over the 1996–2012 period, while being consistent with the other facts we document in the current paper.

With this brief overview of recent developments in Brazil, we turn to a discussion of the data we use to decompose the decline in inequality experienced in Brazil over the past two decades.

## 2.3 Data

Our analysis uses two confidential administrative datasets from Brazil: the *Relação Anual de Informações (RAIS)* contains earnings and demographic characteristics of workers as

[8]Using household data, Barros et al. (2010) estimate that social programs accounted for about 20 percent of the decline in household income inequality.

reported by employers, and the *Pesquisa Industrial Anual Empresa (PIA)* contains detailed information on revenues and costs of large firms in Brazil's mining and manufacturing sectors. To make the reader familiar with these confidential data, we briefly discuss their collection, coverage, variable definitions, and sample selection.

### 2.3.1 Description of linked employer-employee data (RAIS)

**Collection and coverage.** The RAIS data contains linked employer-employee records that are constructed from a mandatory survey filled annually by all registered firms in Brazil and administered by the Brazilian Ministry of Labor and Employment *(Ministério do Trabalho e Emprego*, or *MTE*). Data collection was initiated in 1986 within a broad set of regions, reaching complete coverage of all employees at formal establishments of the Brazilian economy in 1994.[9] Fines are levied on late, incomplete, or inaccurate reports, and as a result many businesses hire a specialized accountant to help with the completion of the survey. In addition, MTE conducts frequent checks on establishments across the country to verify the accuracy of information reported in RAIS, particularly with regards to earnings, which are checked to adhere to the minimum wage legislation.[10]

The RAIS contains an anonymized, time-invariant person identifier for each worker, which allows us to follow individuals over time. It also contains anonymized time-invariant establishment and firm IDs that we use to link multiple workers to their employers and follow those over time. Although it would be possible to conduct part of our analysis at the establishment instead of firm level, this paper focuses on firms for three reasons. First, to the extent that there is substantial variation in pay across establishments within firms, our firm-level anal-

[9]Because registration with the central tax authorities is necessary and sufficient for a firm to be surveyed, the RAIS covers only workers in Brazil's formal sector. Complementing our analysis with data from the Brazilian household survey *Pesquisa Nacional por Amostra de Domicílios* (*PNAD*), we find that the formal sector employment share among male workers of age 18–64 grew from 64 to 74 percent between 1996 and 2012. Differential inequality trends between formal and informal sector workers are discussed at more length in Engbom and Moser (2015).

[10]In addition to being fined, non-compliant firms are added to a "Black List of Slave Work Employers," made available publicly under law Decree No. 540/2004. A recent version of the list dated March 2015 is available from Brazilian television news channel Repórter Brasil at `http://reporterbrasil.org.br/documentos/lista_06_03_2015.pdf`.

ysis provides a lower bound on the importance of workers' place of employment.[11] Second, we think that many of the factors that could give rise to employer-specific components of pay including corporate culture, company leadership, etc., act at the firm level. Additionally many regulations targeting pay policies differ as a function of firm-level employment, not establishment-level employment. Third, we will later use data on firm characteristics such as financial performance that are not available at the establishment-level.

**Variable definitions.** For each firm at which a worker was employed during the year, the RAIS contains information on the start and end date of the employment relationship, the amount the worker was paid and a broad set of worker and job characteristics. Reported earnings are gross and include regular salary payments, holiday bonuses, performance-based and commission bonuses, tips, and profit-sharing agreements. Although this is a broad measure of earnings, it does not contain other sources of income such as capital income or in-kind transfers. We divide total earnings from an employment relationship in a given year by the duration of the job spell.[12] This accounts to some extent for labor supply. As hours worked only exists for some years, we do not use this to construct a measure of per hour pay. Instead, to limit the impact of unmeasured labor supply differences, we focus on adult males.[13]

We define a consistent age variable by calculating the year of birth for any observation, and then setting an individual's year of birth as the modal implied value and finally reconstructing age in each year using this imputed year of birth.[14] Because age is only reported in bins prior to 2002, we code all subsequent years into the same age bins (18–24, 25–29, 30–39, 40–49, and 50–64 years old).

[11] As we will show later, however, the explanatory power of our model incorporating firm and person effects is high, leaving little variation to be explained by separate establishment level effects.

[12] That is, if an employment relationship is reported as active for seven months during the year, we divide total earnings reported for that employment relationship for that year by seven.

[13] In the years for which we have data on hours, we find relatively little variation in hours, with most adult males reporting 44 hours of work a week.

[14] We use age instead of experience throughout our analysis; results are similar using age plus six minus years of education as a measure of experience.

We define a consistent measure of years of schooling by first setting it to its modal value within a year in case of multiple job spells in a year and then ensuring that the years of schooling are non-decreasing across years. Subsequently, we define four education groups based on attained degree implied by the reported number of years of schooling and the education system in Brazil (primary school, middle school, high school, and college).

The data also contain information on detailed occupation classification of the job and detailed sector classification of the employer establishment. Both the industry and occupation classification systems underwent a significant change during the period we study. For occupations, we use the pre-2003 classification (*Classificação Brasileira de Ocupações*, or *CBO*) at the one-digit level. We also use two-digit sectoral classifications (*Classificação Nacional de Atividades Econômicas*, or *CNAE*) according to the pre-2003 period. We make occupations and sectors reported for 2003–2012 consistent with the older CBO and CNAE classifications by using conversion tables provided by IBGE. In order to achieve a high level of consistency between the old and the new classification schemes, we cannot go less coarse than one digit occupation and two digit sector, but we believe that for the purpose of this paper this restriction is not of major importance.

Our firm size measure is the number of full-time equivalent workers during the reference year. Importantly, we calculate this prior to making any sample restrictions so that it reflects to the greatest extent possible the total amount of labor used by the firm during the year. We calculate it as the total number of worker-months employed by the firm during the year divided by 12.

**Sample selection.** We exclude observations with either firm IDs or worker IDs reported as invalid as well as data points with missing earnings, dates of employment, educational attainment or age. Together, these cleaning procedures drop less than one percent of the original population, indicative of the high quality of the administrative dataset. Subsequently, to limit the computational complexity associated with estimating our model, we

restrict attention to one observation per worker-year. We impose this restriction by choosing the highest-paying among all longest employment spells in any given year. As the average number of jobs held during the year is 1.2 and there is not trend in this, we do not believe that loosening this restriction would meaningfully affect our results.

Finally, we restrict attention to adult male workers of age 18–64. We make this restriction as a trade-off between our results being comparable to a large part of the literature focusing on prime age males on the one hand and to obtain as complete as possible coverage of the changes in the Brazilian wage structure over the period on the other.[15] We have tried alternative sample restrictions, including focusing on only prime age males of age 18–49, as well as both male and female adult workers.

**Descriptive statistics.** Table 2.1 provides key summary statistics for the RAIS data for six subperiods of five years each with one year overlap between adjacent periods, namely 1988–1992, 1992–1996, 1996–2000, 2000–2004, 2004–2008, and 2008–2012. Since our analysis focuses on adult males and adult males working for large manufacturing and mining firms, we provide a brief comparison of these subpopulations to the overall population of formal sector employees. As we will be primarily concerned with the later four subperiods during which inequality declined markedly and for which we have firm level data, we focus our discussion on these periods.

Panel A shows statistics for the overall formal sector work force in Brazil and Panel B for the subpopulation of adult males. Adult males are consistently about 0.3–0.4 years older than the population average. They also have 0.78 years of schooling less than the overall sample in the 1996–2000 subperiod; this gradually drops to 0.65 years in the last subperiod. Finally, adult males earn about eight to nine log points more than the overall population, but the variance of log earnings is very similar to the overall population.

[15]The restriction to only male workers has the advantage of avoiding issues with changing patterns of female labor supply and labor market discrimination. In a separate ongoing research project, we investigate the degree to which firm-level average pay and profit sharing differs between male and female employees and how the gender pay gap has evolved over time.

Panel C presents statistics on the subpopulation of adult males working at large mining and manufacturing firms. Adult males in the PIA subpopulation are about 0.8 years younger than all adult males in the 1996–2000 subperiod, which gradually increases to 1.3 years younger in the last subperiod. They are similar to all adult males in terms of education. The PIA sample of adult males earned on average 27 log points more than all adult males in the 1996–2000 subperiod; this declined to only a 19 log point premium in the last subperiod. Finally, they display a two log point higher standard deviation of log earnings in the 1996–2000 period, which increases to four log points in the last subperiod.

### 2.3.2 Description of firm characteristics data (PIA)

**Collection and coverage.** The PIA data contain information on firm financial characteristics from 1996–2012. The dataset is constructed by the Brazilian National Statistical Institute (*Instituto Brasileiro de Geografia e Estatística*, or *IBGE*) based on annual firm surveys in the manufacturing and mining sector. This survey is mandatory for all firms with either more than 30 employees or above a revenue threshold as well as for an annual random sample of smaller firms.[16] As with RAIS, completion of the survey is mandatory and non-compliance is subject to a fine by national authorities. Each firm has a unique, anonymized identifier, which we use to link firm characteristics data from PIA data to worker-level outcomes in the RAIS data.

**Variable definitions.** The PIA dataset includes a breakdown of operational and non-operational revenues, costs, investment and capital sales, number of employees and payroll. All nominal values are converted to real values using the CPI index provided by the IBGE. Instead of the measure of firm size in the PIA, we prefer our measure of full-time-equivalent employees constructed from the RAIS as it accounts for workers only employed during part of the year. We define operational costs as the cost of raw materials, intermediate inputs,

[16]The revenue threshold for inclusion in the deterministic survey has grown over the years, standing at USD300,000 in 2012.

Table 2.1: RAIS summary statistics

| | | | Earnings | | Age | | Schooling | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| | # Worker-years | # Unique workers | Mean | St.d. | Mean | St.d. | Mean | St.d. |
| Panel A. All Formal Sector Workers (RAIS) | | | | | | | | |
| 1988–1992 | 165.5 | 41.9 | 1.10 | 0.86 | 31.91 | 11.47 | 7.65 | 4.45 |
| 1992–1996 | 162.1 | 43.4 | 1.18 | 0.86 | 33.20 | 11.32 | 8.08 | 4.41 |
| 1996–2000 | 174.6 | 47.0 | 1.19 | 0.84 | 33.68 | 11.27 | 8.60 | 4.27 |
| 2000–2004 | 202.7 | 52.7 | 1.00 | 0.80 | 34.02 | 11.33 | 9.49 | 4.05 |
| 2004–2008 | 254.2 | 62.7 | 0.81 | 0.74 | 34.26 | 11.48 | 10.25 | 3.78 |
| 2008–2012 | 326.5 | 76.2 | 0.71 | 0.71 | 34.55 | 11.66 | 10.78 | 3.52 |
| Panel B. Adult Male Workers | | | | | | | | |
| 1988–1992 | 86.5 | 25.5 | 1.24 | 0.87 | 33.26 | 10.82 | 7.04 | 4.30 |
| 1992–1996 | 87.3 | 26.4 | 1.29 | 0.87 | 33.88 | 10.79 | 7.37 | 4.27 |
| 1996–2000 | 92.7 | 28.8 | 1.27 | 0.85 | 33.97 | 10.78 | 7.82 | 4.16 |
| 2000–2004 | 105.3 | 32.5 | 1.07 | 0.80 | 34.14 | 10.92 | 8.70 | 4.00 |
| 2004–2008 | 126.9 | 37.3 | 0.88 | 0.75 | 34.44 | 11.11 | 9.50 | 3.79 |
| 2008–2012 | 154.2 | 43.9 | 0.80 | 0.72 | 34.91 | 11.34 | 10.13 | 3.59 |
| Panel C. Adult Male Workers at Large Manufacturing and Mining Firms (PIA) | | | | | | | | |
| 1996–2000 | 16.6 | 6.3 | 1.54 | 0.87 | 33.20 | 10.07 | 7.83 | 4.05 |
| 2000–2004 | 18.0 | 6.8 | 1.29 | 0.85 | 33.04 | 10.20 | 8.75 | 3.91 |
| 2004–2008 | 23.2 | 8.5 | 1.09 | 0.80 | 33.11 | 10.45 | 9.41 | 3.77 |
| 2008–2012 | 26.9 | 9.9 | 0.99 | 0.76 | 33.60 | 10.70 | 10.04 | 3.60 |

Notes: The number of worker-years and number of unique workers are reported in millions. Statistics on earnings are in log multiples of the current minimum wage, schooling is in years. Panel A includes all workers in the RAIS dataset. Panel B includes male workers that are between 18 and 64 years old. Panel C includes male workers age 18–64 working at large manufacturing and mining firms included in the PIA firm characteristics data. Means are computed by period. The standard deviation is calculated by first demeaning variables by year and then pooling the years within a subperiod.

electricity and other utilities, and net revenues as the gross sales value due to operational and non-operational firm activities net of any returns, cancellations, and corrected for changes in inventory.[17]. We finally construct value added as the difference between net revenues and intermediate inputs, and value added per worker as value added divided by full-time equivalent workers. This is our main measure of firm productivity. We have also constructed

[17]We have explored alternative revenue definitions such as only restricting attention to operational revenues or excluding certain types of non-operational revenues. Such robustness checks yield very similar results to what we report below.

alternative measures of firm productivity by cleaning value added per worker off industry-year effects and some measures of worker skill. In our main analysis, we focus on "raw" value added per worker and present results containing these alternative measures in the Appendix.

Our productivity measure differs from the commonly used total factor productivity (TFP) (Bartelsman, Haltiwanger and Scarpetta, 2009, 2013) since it does not control for capital intensity. A major reason for this is that we do not have data on capital, only on investment. To construct a measure of the capital stock, we would need to assume a depreciation rate to be able to impute capital using reported investment. We would also need to impute capital in 1996 since we do not have data prior to that, as well as for any firm that enters the PIA population. We have constructed such a measure of the capital stock using an assumed annual depreciation rate of five percent and using data on the aggregate capital stock at the subsector level.[18] However, the multiple imputations required to obtain capital as well as the fact that the investment data is incomplete for many firms lead us to prefer value added per worker as our measure of firm productivity.[19]

**Sample selection.** The PIA firm survey spans the universe of large firms (as defined above) in Brazil's manufacturing and mining sectors in addition to a random sample of smaller firms. Because parts of our analysis make use of the panel dimension on the firm side and to avoid issues with excessive sample attrition related to our later estimation procedure, we focus our analysis on the deterministic set of relatively large firms.

**Descriptive statistics.** Table 2.2 shows key summary statistics on firms during the four periods for which we have firm financial data: 1996–2000, 2000–2004, 2004–2008, and 2008–

[18]Each new firm starts with an initial capital equal to its current net investment plus a share of total capital in its subsector. The shares are given by taking the share of capital at a firm to be proportional to the share of total net revenues assuming a firm-level production function of the form $y = Ak^{\alpha}$ for $\alpha = 1/3$. Firms entering the PIA at a later year are initiated by applying the same method to get those firms' capital stock proportional to scaled firm revenues relative to the subsector total.

[19]In addition, several bargaining models of the labor market have in common that workers and capital owners split the surplus from production, and value added per worker is arguably the best measure of that surplus. Thus, to the extent that such models well describe Brazilian labor markets, value added per worker is an important metric.

2012. All results are weighted by the number of full-time equivalent workers employed by the firm. The number of firms in the PIA increased by 57 percent between the first and the last period. The average firm size increased by 32 percent and average real value added per worker grow by 27 percent. There is significant dispersion in both log firm size and log value added per worker across firms, with the standard deviation of the former being close to two and that of the latter exceeding one. Furthermore, there is no evidence of convergence in either measure. The standard deviation of firm size monotonically increases whereas the standard deviation of value added per worker first increases rapidly, then falls again in the last subperiod. To the extent that firm characteristics matter for employees' labor remuneration, these results suggest that the decline in earnings inequality in Brazil cannot be explained by declining dispersion in these characteristics over time.

Table 2.2: PIA summary statistics

| | | | Firm size | | Value added p.w. | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| | # Firm-years | # Unique firms | Mean | St.d. | Mean | St.d. |
| 1996–2000 | 142.4 | 51.1 | 6.34 | 1.80 | 11.15 | 1.13 |
| 2000–2004 | 168.7 | 59.9 | 6.26 | 1.85 | 11.19 | 1.32 |
| 2004–2008 | 202.4 | 73.0 | 6.49 | 1.96 | 11.22 | 1.34 |
| 2008–2012 | 230.2 | 80.2 | 6.62 | 2.05 | 11.30 | 1.31 |

Note: The number of firm-years and number of unique firms are reported in thousands Firm size is the log number of full-time equivalent employees. Value added per worker is the log of real value added per worker. Means and standard deviations are weighted by the number of full-time employees. The standard deviation is calculated by first demeaning variables by year and then pooling the years within a subperiod.

## 2.4 Inequality trends in Brazil from 1988–2012

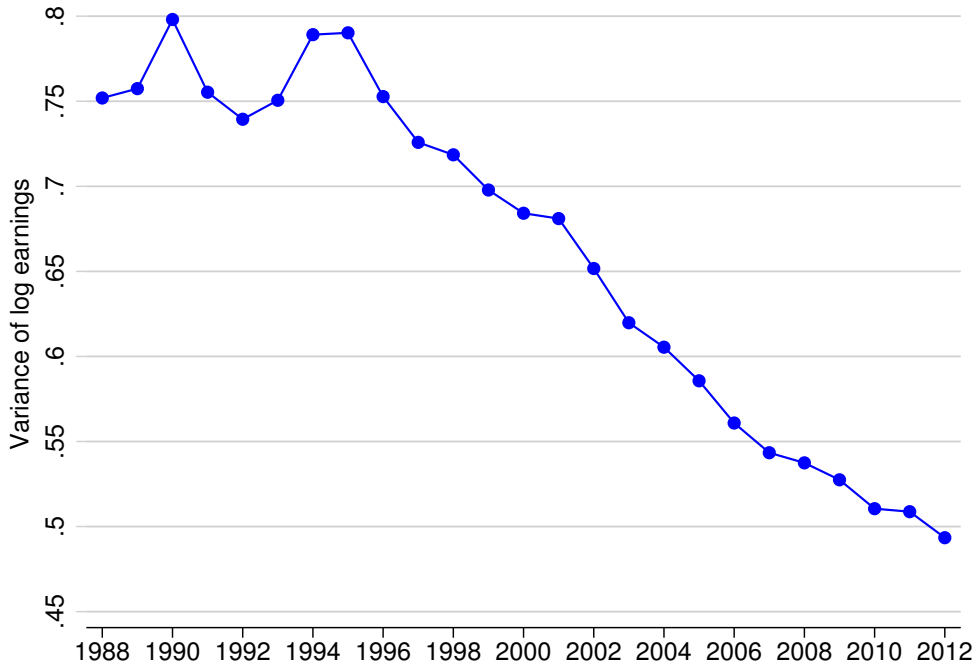
In this section, we first document Brazil's rapid decline in earnings inequality. We then demonstrate that the decline in inequality in Brazil occurred throughout a large part of the earnings distribution. Subsequently, we present results from a series of Mincer regressions, which provides a first look at possible factors behind the decline. Finally, we provide some

suggestive evidence of firms being an important source of inequality as well as a factor behind the decline in inequality in Brazil.

### 2.4.1 The evolution of earnings inequality

Starting from high initial levels, Brazil experienced a rapid and steady declined in earnings inequality from 1995 onwards. The inequality decline followed years of stable or slightly increasing inequality between the late 1980s and the mid-1990s. While the inequality decline is apparent in many measures, Figure 2.1 plots the variance of log earnings, which is a common inequality measure in the labor literature. Between 1996 and 2012, the variance of log earnings in Brazil declined by 26 log points or 35 percent, from 0.75 to 0.49 over the period. To put this decline into context, the U.S. saw an increase in the same measure of 22 log points from 1968–2006.

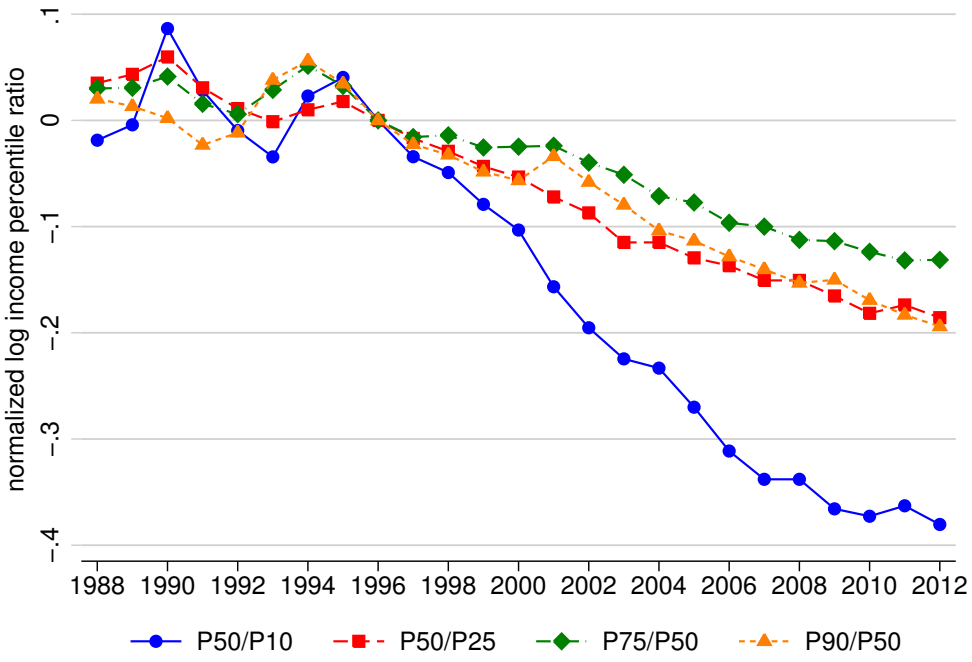
Figure 2.1: Variance of log earnings in Brazil, 1988–2012



### 2.4.2 Compression in different parts of the earnings distribution

Figure 2.2 plots the log percentile ratios of earnings, normalized to zero in 1996. Two striking facts emerge from the graph: First, there was widespread compression in the distribution of earnings—inequality declined past the 75th percentile. Second, the amount of compression gradually declines as one moves further up in the distribution. For instance, whereas the log 90–50 percentile ratio falls by 20 log points, the log 50–10 ratio falls by a remarkable 35 log points. Similarly, compression in the log 50–25 percentile ratio exceeds compression in the log 75–50 ratio.[20]

Figure 2.2: Log percentile ratios of the earnings distribution in Brazil (1996 = 0)

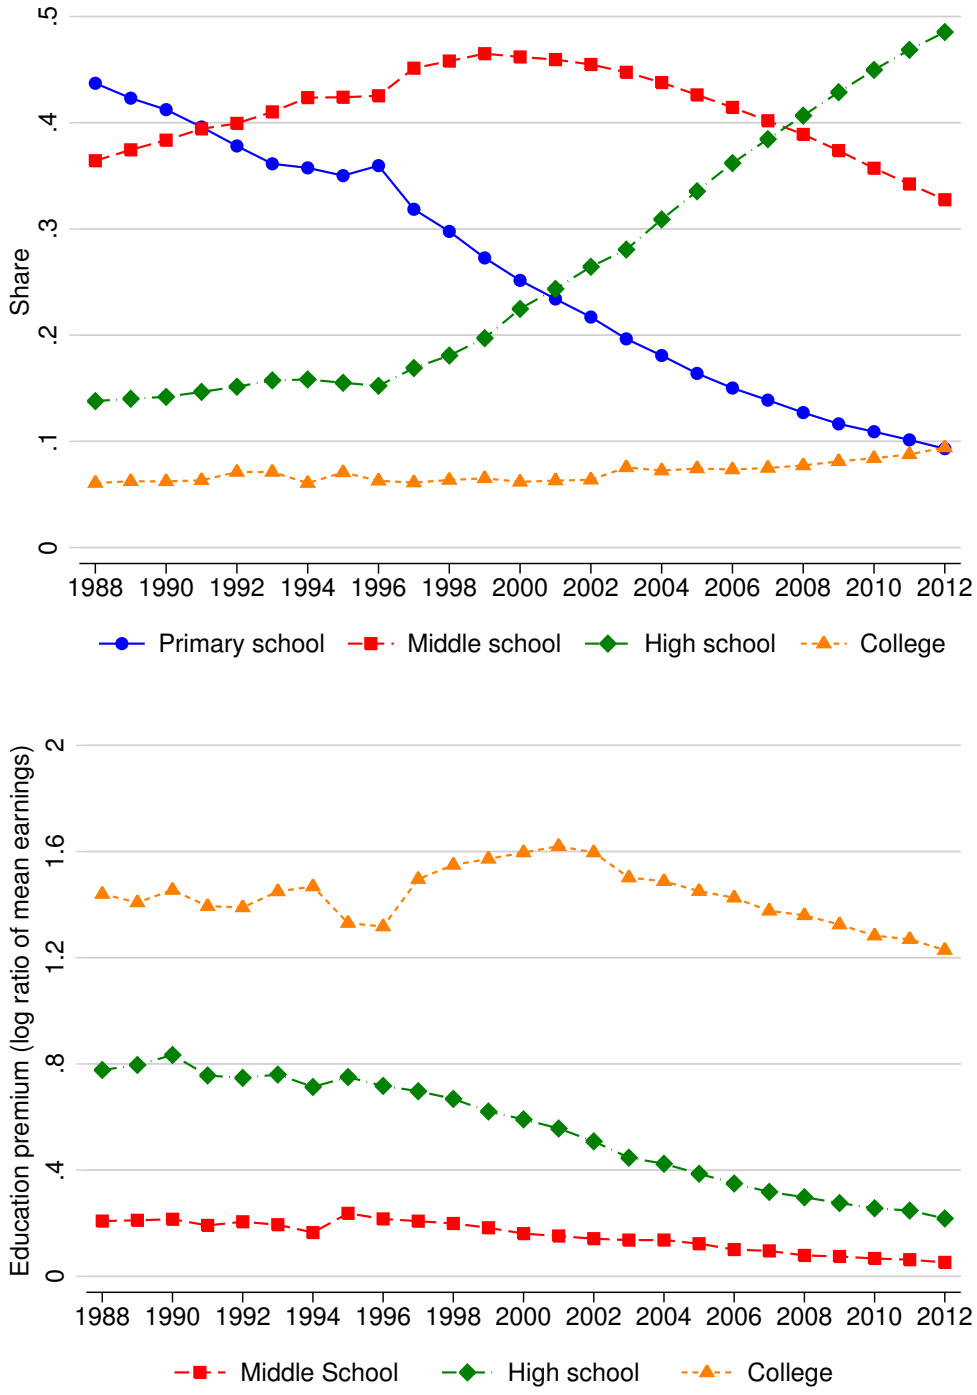



[20]For ease of interpretation and to show that there was significant real wage growth at all percentiles of the income distribution over the period for which such data exist, Appendix 2.8.1 plots the real percentile earnings levels evolution from 1996–2012, normalized to zero in 1996.

### 2.4.3 The (un)importance of worker observables

One candidate explanation for the decline in earnings inequality is increasing educational attainment. As can be seen in the left panel of Figure 2.3, the fraction of the Brazilian formal sector workforce with a high school degree rose rapidly during this time, while the fraction with primary school fell sharply (the fraction with a middle school degree and a college degree remained relatively flat). There were also important changes to the premia associated with a higher degree, as can be seen in the right pane of Figure 2.3 . In particular, the premia associated with a middle school and high school degree relative to the lowest education group fell rapidly over the past 20 years.

Figure 2.3: Educational attainment (top) and education premia (bottom)




To provide a first look at whether changes in observable worker characteristics were an important driver of earnings inequality in Brazil over this period, we run a series of Mincer

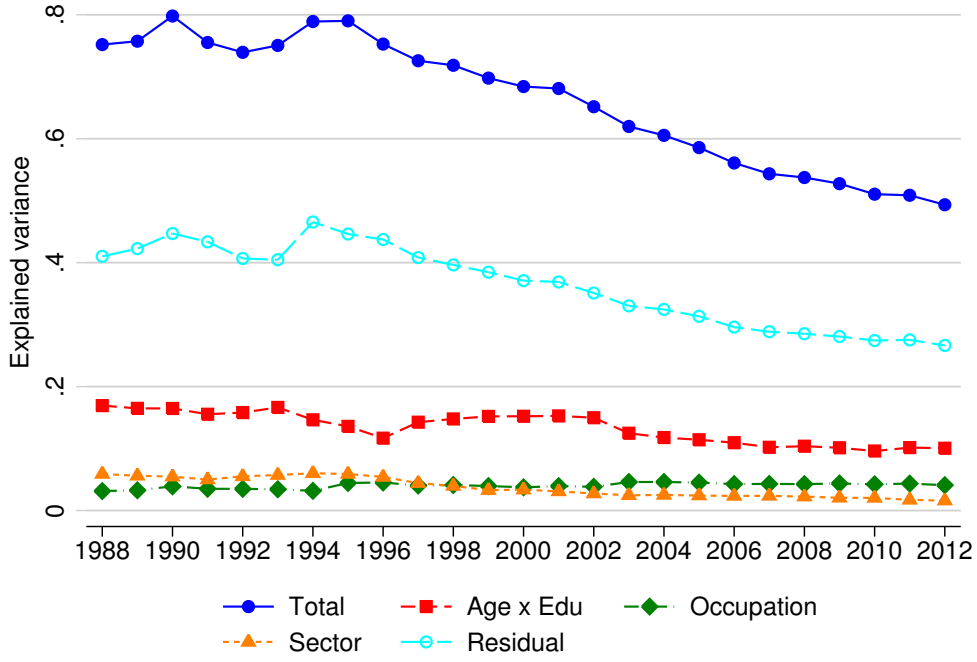
regressions. In particular, we regress log earnings of individual $i$ in year $t$, denoted $y_{it}$, on five age group dummies interacted with nine education group dummies, two digit occupation dummies and two digit sector dummies,

$$\log(y_{it}) = age_{it} \times edu_{it} + occ_{it} + sec_{it} + \varepsilon_{it}$$

Note that all explanatory variables are allowed to vary freely by year. Based on this regression, we calculate the predicted value due to each component and report the variance of these predicted values.

Figure 2.4 plots the results. In levels, worker observables jointly explain about 45 percent of the overall variance in log earnings. This does not change much over time, and hence worker observables explain close to 45 percent of also the fall in inequality. Decomposing this, age and education account for roughly 20 percent of the variance of log earnings and 27 percent of the decline. Inequality between occupations increases relative to overall inequality from four to eight percent of total variance (and in fact it also grows slightly in absolute terms). The fraction of total inequality explained by differences in means across sectors falls from eight to three percent—this accounts for 12 percent of the overall decline in variance. Finally, covariances between the explanatory variables increase slightly in importance from 11 percent to 14 percent of total variance and explain four percent of the overall fall in inequality. We conclude that even when controlling for detailed worker characteristics, more than half of the level of inequality as well as its decline is residual in nature.

Figure 2.4: Variance decomposition from Mincer regressions




### 2.4.4 The evolution of earnings inequality between and within firms

As a first step towards understanding the role of firms for earnings inequality, we investigate the variance of earnings within and between firms. To this end, we define between-firm inequality as the variance of the average log earnings at the firm across firms (weighted by firm size) and within-firm inequality as the variance of the difference between workers' log earnings and the average log earnings at their firm. Based on these definitions, one could imagine two hypothetical polar extremes. First, average earnings could be identical across firms so that overall earnings inequality is completely due to variance in earnings within firms. In this case, a firm is just a microcosm of the overall economy. Second, all workers could earn the same wage within the firm so that inequality arises entirely due to differences in earnings across firms. In reality, the question is which channel is quantitatively most important.

Formally, we follow Bloom et al. (2015) in letting $y_{ijt}$ denote earnings of worker $i$ employed by firm $j$ in year $t$, then:

$$y_{ijt} = \underbrace{\overline{y_t}}_{\text{economy average}} + \underbrace{\left(\overline{y}_t^j - \overline{y_t}\right)}_{\text{employer deviaion}} + \underbrace{\left(y_{ijt} - \overline{y}_t^j\right)}_{\text{worker deviation}}$$

Re-arranging and taking variances on both sides we get

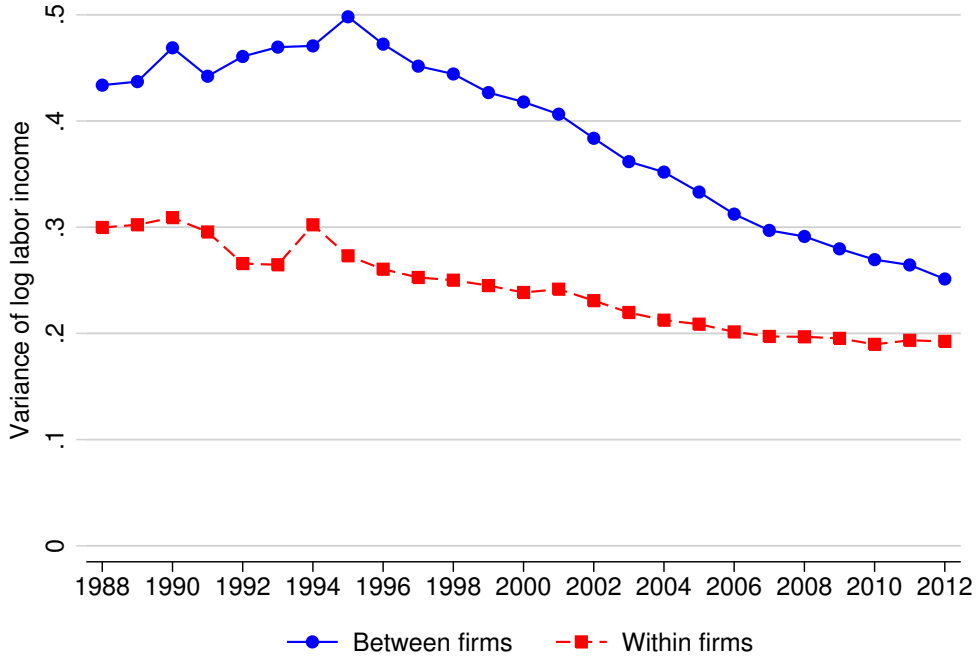
$$Var\left(y_{ijt} - \overline{y_t}\right) = Var\left(\overline{y}_t^j - \overline{y_t}\right) + Var\left(y_{ijt} - \overline{y}_t^j\right) + \underbrace{2Cov\left(\overline{y}_t^j - \overline{y_t}, y_{ijt} - \overline{y}_t^j\right)}_{=0}$$

where the last term is zero by construction. Simplifying, we have our decomposition of the overall variance of earnings into between- and within-firms inequality:

$$\underbrace{Var\left(y_{ijt}\right)}_{\text{overall}} = \underbrace{Var\left(\overline{y}_t^j\right)}_{\text{between firms}} + \underbrace{\overline{Var\left(y_{ijt} \middle| i \in j\right)}}_{\text{within firms}}$$

Figure 2.5 plots this decomposition over time in Brazil. We note two insights: Firstly, there is significant variability in earnings within firms, but an even greater amount of earnings inequality across firms. Secondly, although both measures of inequality fell during this time, the decline was particularly pronounced across firms: inequality across firms declined by 25 log points or 45 percent between 1988 and 2012, whereas within-firm inequality dropped by 10 log points or 33 percent.

Figure 2.5: Variance of log earnings within and between firms




Another way of illustrating the importance of firm is to compare earnings growth of different workers with average earnings growth at their employers. Making use of the link between workers and firms in the matched employer-employee data (RAIS), one can ask how much of the earnings growth accruing to a certain income group was mediated by rises in average pay at firms employing workers from that income group.[21]

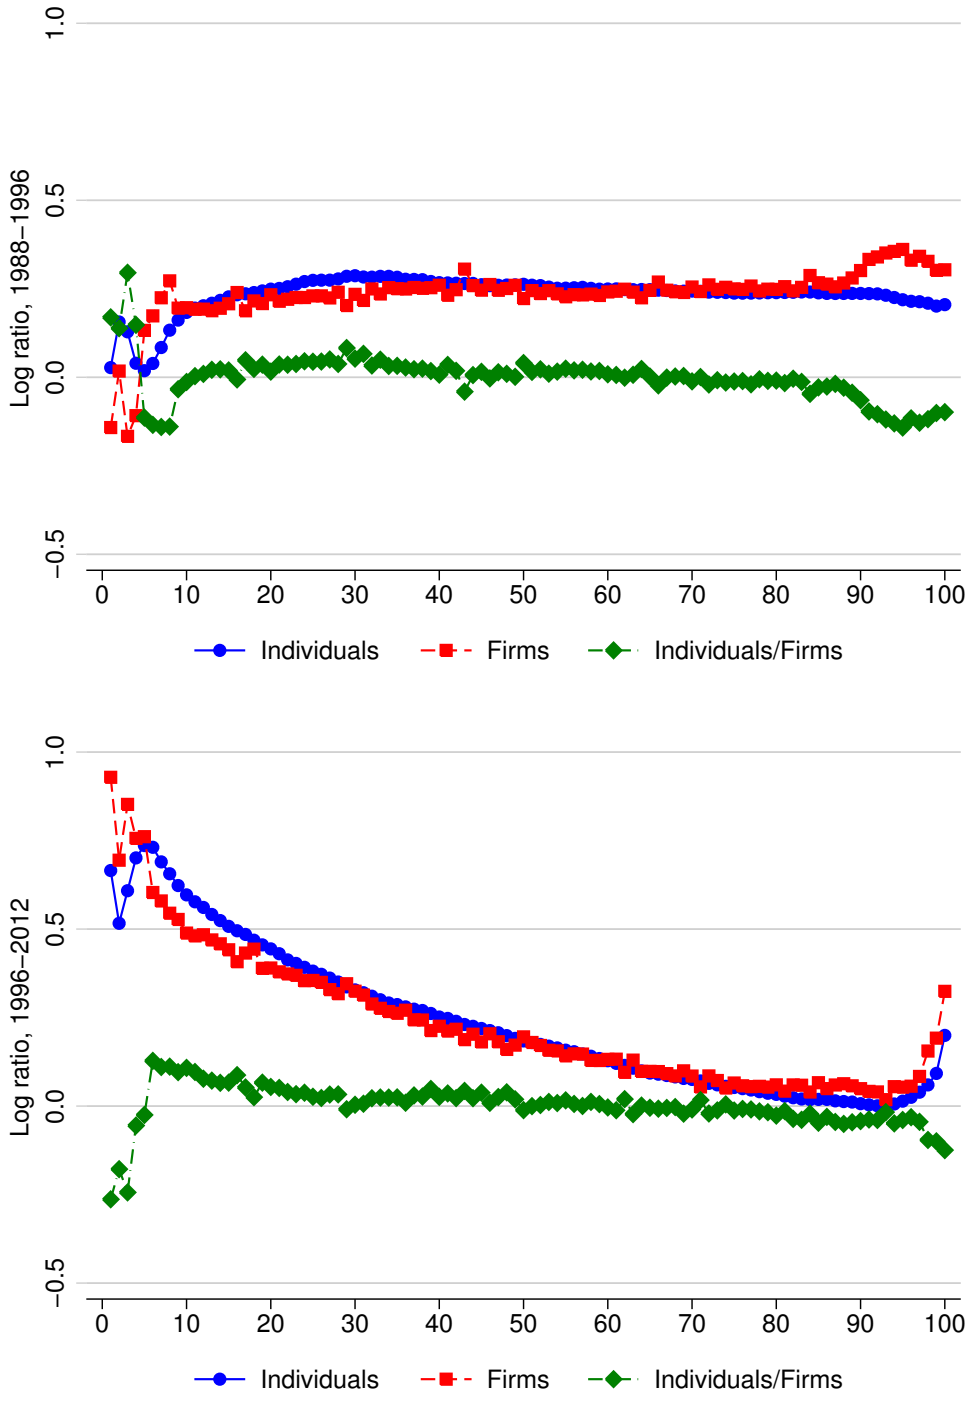
The results of this exercise are shown in Figure 2.6, first for the period 1988–1996, when earnings inequality remained roughly constant, and then for 1996–2012, when earnings inequality declined rapidly. Average earnings growth (solid blue line with circles) was relatively evenly distributed throughout the earnings distribution between 1988 and 1996, and firm average pay (solid red line with squares) grew equally in line with the growth rate of wages. The period from 1996–2012, on the other hand, was marked by a rapid catch-up of the lowest earnings groups, which in turn is almost entirely explained by the growth of firm average

[21]Similar calculations are presented for the U.S. in Barth et al. (2014) and Bloom et al. (2015).

earnings among those groups. Throughout both periods, there were no significant changes in within-firm earnings inequality (solid green line with diamonds).

Although informative, however, these type of decompositions of raw earnings cannot necessarily be interpreted as firms differing fundamentally in the way they compensate their workers. The reason is that some firms could hire workers who always get paid more regardless of where they work (maybe because they are more productive, have a higher bargaining power, etc). In this case differences in pay across firms would arise as a result of recruitment policies and not pay policies. The next section formalizes our approach to identifying the importance of firm pay policies for earnings inequality.

Figure 2.6: Individual labor earnings growth between and within firms, 1988–1996 (top) and 1996–2012 (bottom)



## 2.5 Empirical framework

For a long time, economists have recognized that worker observables fail to explain a large fraction of the variance of earnings. As we showed above, this is true also for Brazil—even with detailed occupation and sector controls, Mincer regressions explain less than half of the overall variation in earnings in Brazil. Furthermore, a recent literature has argued that there are important differences across firms in terms of pay (Abowd, Kramarz and Margolis, 1999), and we presented some evidence above suggesting that firms might be an important determinant of earnings also in Brazil. Motivated by these insights, we estimate two-way fixed effects econometric models controlling for both unobserved worker and firm heterogeneity. To be able to speak to changes over time in the components of inequality, we estimate our model separately in six sub-periods covering 1988–1992, 1992–1996, 1996–2000, 2000–2004, 2004–2008, and 2008–2012, respectively. Subsequently, we correlate the estimated firm and worker effects with observed characteristics of firms and workers in order to investigate what led to changes in the firm and worker component of pay over time.

### 2.5.1 Introducing the two-way fixed effects model

In order to identify the two-way fixed effects framework pioneered by Abowd, Kramarz and Margolis (1999), one needs to observe a panel of workers with the ability to link multiple workers to the same firm. Our data satisfy these requirements. Within each subperiod, we observe a large number $I$ of workers for up to five years while working at $J$ firms for a total of $N$ worker-years. Let $J(i,t)$ give the employer of worker $i$ in year $t$. We assume that the log earnings of individual $i$ in year $t$ within a period time window $p$, denoted $y_{it}$, can be written as the sum of a worker effect, $\alpha_i^p$, a firm effect, $\alpha_{J(i,t)}^p$, a time dummy, $\gamma_t$, and an error term, $\varepsilon_{it}$. Although our current paper does not provide a microfoundation to this reduced form empirical model, in subsequent work we show how it can be rationalized in

a frictional labor market with firm productive heterogeneity and worker ability differences (Engbom and Moser, 2015).

Our specification does not control in this first stage for observable worker and firm characteristics. Instead, we correlate the estimated fixed effects with observable characteristics of workers and firms in a second stage of our analysis. We prefer this specification to avoid identifying these effects off changes within workers and firms during the limited time frame of each subperiod. With regards to age, we additionally notice that unrestricted age controls would be perfectly collinear with person effects and the time dummies.[22] Similarly, we also correlate worker observables, including age, education, and occupation, with the estimated worker effect in a second stage. We argue that any error due to growing earnings with age is likely to be of second order importance within our five year subperiods. We thus estimate

$$\log y_{it} = \alpha_i^p + \alpha_{J(i,t)}^p + \gamma_t + \varepsilon_{it}$$

for $t \in p = \{t_1, \ldots, t_5\}$ and where $\alpha_i^p$ denotes the individual fixed effect of worker $i$ in period $p$, $\alpha_{J(i,t)}^p$ denotes the firm effect of the employer of worker $i$ at year $t$, $\gamma_t$ is a year dummy, and $\varepsilon_{it}$ is an error term that satisfies the strict exogeneity condition $\mathbb{E}\left[\varepsilon_{it} \middle| i, t, J\left(i,t\right)\right] = 0$.[23]

As shown by Abowd, Kramarz and Margolis (1999), worker and firm effects can only be separately identified within a set of firms and workers connected through the mobility of workers. Table 2.3 presents summary statistics on the largest set of workers in each subperiod—this covers 97–98 percent of all workers in each subperiod. Given that it covers such a large fraction of all adult males, it is not surprising that this subpopulation looks very similar to the overall population in all observable dimensions. Thus the restriction to the largest connected set imposed in the rest of our analysis appears to be inoccuous.

As identification critically derives from workers switching between firms, Table 2.4 presents statistics on the fraction of switchers in each subperiod. The degree of labor

[22] Although restrictions could be imposed to address this collinearity problem, we note that for instance the popular restriction advocated by Deaton (1997) requires many years to be well identified.

[23] We later discuss in greater detail the empirical validity of the assumption on the error term.

Table 2.3: Summary statistics on workers in largest connected set (as a fraction of adult males)

| | | | Earnings | | Age | | Schooling | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| | # Worker-years | # Unique workers | Mean | St.d. | Mean | St.d. | Mean | St.d. |
| 1988–1992 | 85.4 (99%) | 25.1 (98%) | 1.3 | 0.9 | 33.3 | 10.8 | 7.0 | 4.3 |
| 1992–1996 | 85.6 (98%) | 25.8 (98%) | 1.3 | 0.9 | 33.9 | 10.8 | 7.4 | 4.3 |
| 1996–2000 | 90.2 (97%) | 27.9 (97%) | 1.3 | 0.8 | 34.0 | 10.8 | 7.8 | 4.2 |
| 2000–2004 | 102.1 (97%) | 31.4 (97%) | 1.1 | 0.8 | 34.2 | 10.9 | 8.7 | 4.0 |
| 2004–2008 | 123.7 (97%) | 36.2 (97%) | 0.9 | 0.8 | 34.4 | 11.1 | 9.5 | 3.8 |
| 2008–2012 | 151.0 (98%) | 42.8 (98%) | 0.8 | 0.7 | 34.9 | 11.3 | 10.1 | 3.6 |

Notes: We report in parentheses the proportion of the reported statistics relative to the group of adult males listed in Table 2.1. Earnings are in log multiples of the minimum wage, schooling is years of education.

mobility is high in Brazil with more than 30 percent of the population switching firm at some point in the subperiod. The average number of firms worked at during the five years in each subperiod is about 1.5. There is no strong trend in either statistic.

Table 2.4: Frequency of switches, by period

| | (1) | (2) | (3) |
|---|---|---|---|
| | # Unique workers | Average # of jobs | % switchers |
| 1988–1992 | 25.1 | 1.55 | 0.36 |
| 1992–1996 | 25.8 | 1.47 | 0.32 |
| 1996–2000 | 27.9 | 1.44 | 0.31 |
| 2000–2004 | 31.4 | 1.45 | 0.31 |
| 2004–2008 | 36.2 | 1.53 | 0.36 |
| 2008–2012 | 42.8 | 1.64 | 0.40 |

Note: Number of unique workers in millions. A switcher is defined as a worker who is associated with two or more employers during the period.

The assumption we impose on the error term is often referred to in the literature as that of requiring *exogenous mobility*. As explained by AKM, this rules out dependency of the error term on the worker effect, the firm effect or the time controls. In particular, a worker is not allowed to switch between firms based on the unobserved error term, because if say matches with a particularly poor match effect are more likely to break up, the residual of remaining matches does not have mean zero. Moreover, this assumption rules out assortative matching

of the type found in Roy models, since these models emphasize the complementarity in matches between workers and firms, whereas the AKM framework imposes log-additivity between the two.

We investigate whether this assumption is violated in two ways. First, we follow Card, Heining and Kline (2013) in dividing estimated firm effects into quartiles and investigate whether the gain in the firm component of those switching between for instance the first and fourth quartile is similar to the loss of those making the reverse switch. To the extent that the labor market is better characterized by assortative matching as in Roy models, we would expect these to be very different. Second, we examine the distribution of error terms across worker and firm effects quantiles to check for systematic variation, which could be an indication that our log additive model is misspecified.

Based on our estimated equation, we decompose the variance of log earnings within any subperiod into the variance of the worker component, the firm component, the year trend and the residual, as well as the covariance between the worker and the firm component, the worker and year component, and the firm and year component:

$$\begin{aligned} Var\left(\log y_{it}\right) &= Var\left(\alpha_i\right) + Var\left(\alpha_{J(i,t)}\right) + Var\left(\gamma_t\right) + 2Cov\left(\alpha_i, \alpha_{J(i,t)}\right) \\ &\quad + 2Cov\left(\alpha_i, \gamma_t\right) + 2Cov\left(\alpha_{J(i,t)}, \gamma_t\right) + Var\left(\varepsilon_{it}\right) \end{aligned} \tag{2.1}$$

We note that sampling error in the estimated person and firm effects will cause us to overestimate the variance of worker and firm effects and induce a negative bias in the covariance between worker and firm effects. We do not attempt to correct for this. Instead, we assume that this error is constant over time, so that even if the level of our estimated variances is slightly overstated, the changes we document over time are still valid. [24]

[24]To better estimate firm effects, Bonhomme, Lamadon and Manresa (2015) suggest restricting attention to firms whose fixed effect is "well-identified" due to a high number of switchers. In practice, this procedure boils down to restricting attention to workers at firms with at least 10 switchers during the estimation period. With this restriction, we find a slightly more pronounced role for worker effects in explaining both the initial levels and the decline of earnings inequality between 1996–2000 and 2008–2012.

### 2.5.2 Determinants of the estimated firm effects

In the second stage of our empirical investigation, we study how the estimated firm effects relate to observable measures of firm performance available in the PIA survey. In particular, we are interested in understanding what firm characteristics are related to pay, and whether any changes in the distribution of firm effects over time can be explained by underlying changes in firm characteristics or the way the labor market translates those into pay. Since the PIA only covers the set of large manufacturing and mining firms,[25] we are forced to restrict attention to only these firms and workers when linking firm effects to firm characteristics. We implement this by first estimating the AKM model for the universe of firms and then subsequently restricting attention to only large firms.

Consider a given subperiod and let $a_j$ be the estimated firm component of pay, $\text{VA}_j$ average log value added per worker during the subperiod and $\text{S}_j$ a set of 26 two digit subsector controls. For each subperiod, we regress

$$a_j = \gamma_0 + \gamma_1 \text{VA}_j + \text{S}_j + \varepsilon_j$$

Notice that all regressions are run with sub-period averages of all variables. Additionally, we have considered versions including a range of other firm characteristics, but as these are only marginally important we do not show them here. Based on the above regression, we compute the variance due to value added per worker as

$$Var\left(\hat{a}_j\right) = \left(\hat{\gamma}_1\right)^2 Var\left(\text{VA}_j\right)$$

In a similar vein, we can compute the variance of firm effects due to sub-sector differences.

In order to isolate the importance of a compression in firm fundamentals versus a compression in the pass-through from such fundamentals to pay, we consider two counterfactuals.

[25] As described in Section 2.3, we restrict attention to the deterministic stratum of PIA containing only large firms. We drop small firms contained in the random stratum to ensure that firms stay in the sample for multiple years for our estimation procedure below.

First, we keep the pass-through from value added per worker to pay, $\hat{\gamma}_1$ constant at the estimated level in 1996–2000 and let the variance of value added per worker change as in the data. Second, we keep the variance of value added per worker at its 1996–2000 level and letting the estimated coefficient $\hat{\gamma}_1$ change as in the data. Similarly, we decompose the change in the variance due to sub-sectors into changes in sub-sector pay premia versus a change in the sectoral composition of the economy. A comparison of the two counterfactuals allows us to address whether a change in the variance of firm pay is explained by changes in underlying firm characteristics or due to a change in the degree of pass-through from firm value added to worker pay.

### 2.5.3 Determinants of the estimated worker effects

In the second stage of our analysis, we also investigate what factors influence the worker component of pay. To this end, we regress the estimated worker effects, $a_i$, on five age bin dummies, four education dummies and ten occupation dummies:

$$a_i = \text{age}_i + \text{edu}_i + \text{occ}_i + \varepsilon_i$$

We report results both with and without occupation controls. We have also examined versions of this regression with age and education interacted as well as including sector controls, but neither of these alternatives changes the estimated results significantly. Based on the above regression estimates, we predict the variance due to each of age, education and occupation.

Subsequently, we decompose the evolution of the variance of the age, education and occupation components into that due to a change in the underlying distribution of such characteristics versus a change in the return to them. That is, we first keep the estimated return to the characteristic of interest constant and change only the underlying distribution to match its evolution in the data. This shows how important changes in the distribution

of worker characteristics were for the overall decline. Subsequently, we instead change the return to the characteristic of interest as in the data, while holding the underlying distribution constant. This allows us to evaluate how important changes in the return to age, education and occupation were for the overall decline.

## 2.6 Results

In this section, we first present the results from our two-way fixed effects model decomposing earnings inequality into a firm and a worker component. Subsequently, we investigate the sources of the firm and worker component of pay. Finally we provide additional results evaluating the assumptions imposed by our econometric model.

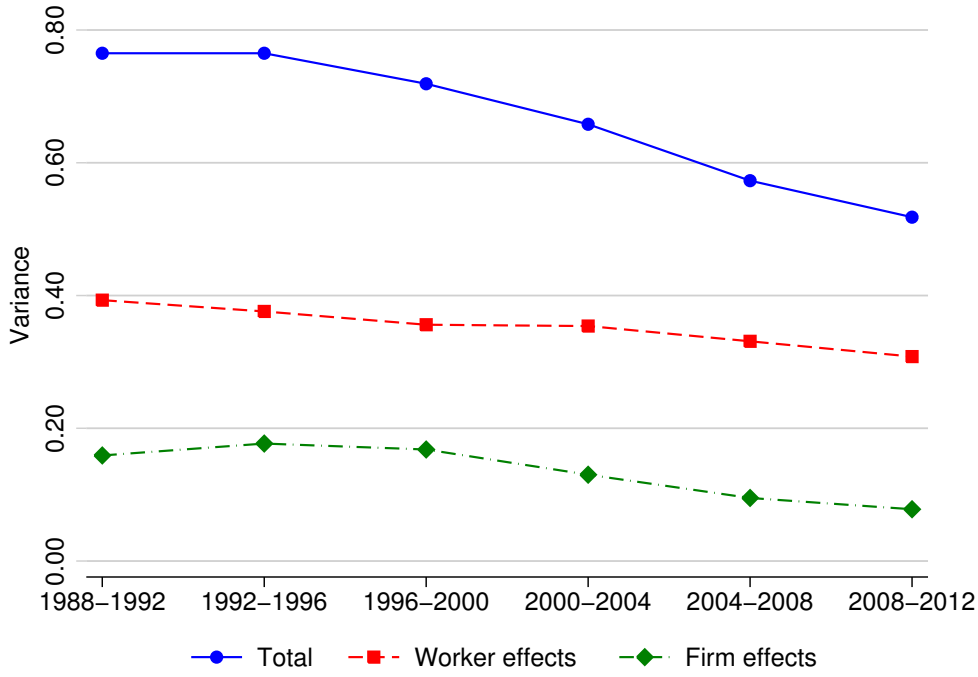
### 2.6.1 AKM decomposition

Table 2.5 presents the variance decomposition based on the estimation results of the AKM model in equation (2.1) for each of the six five-year subperiods between 1988 and 2012. To illustrate the relative importance of the various components of this decomposition over time, Figure 2.7 plots the variance of raw earnings (solid blue line with circles), the variance of estimated worker effects (dashed red line with squares), and the employee-weighted variance of firm effects (dash-dotted green line with diamonds) in each subperiod.

Two important results emerge from our analysis. First, although firm heterogeneity is a non-negligible source of earnings inequality *levels*, worker heterogeneity is the single most important determinant. In the 1996–2000 subperiod, the variance of worker fixed effects makes up 50 percent of the total variance of log earnings. This share increases monotonically to 60 percent by the last subperiod. The variance of firm effects, on the other hand, makes up 23 percent of the variance of log earnings in the 1996–2000 subperiod, decreasing to 15 percent by the last subperiod.

Second, in terms of explaining *time trends*, we observe a more than proportionate fall in the variance of firm effects. Between 1996–2000 and 2008–2012, the variance of firm effects falls from 17 to eight log points whereas the variance of person effects falls from 36 to 31 log points. Additionally, the declining variance of each component is reflected in a lower covariance between the two, with the correlation between worker and firm effects staying fairly constant at around 0.3 throughout the period.[26] Given the large role played by firms in the decline, understanding the drivers of more equal pay across employers over time is an important question which we address in the following section.

Figure 2.7: Variance decomposition from AKM model with firm and worker fixed effects




When studying the link between firm effects and firm performance, we limit attention to the manufacturing and mining sector, for which we have data on firm performance and characteristics.[27] Table 2.6 compares AKM estimates for this subpopulation with the over-

[26] Given that sorting patterns—measured by the correlation between worker and firm effects—remains fairly constant over time, the true reduction in earnings inequality due to firms would be even greater if the covariance term were distributed proportionately between workers and firms.

[27] As noted earlier, we impose the restriction to the PIA subpopulation after estimating the AKM model on the entire population of adult males.

Table 2.5: AKM variance decomposition results

| | (1) 1988–1992 | (2) 1992–1996 | (3) 1996–2000 | (4) 2000–2004 | (5) 2004–2008 | (6) 2008–2012 | (7) Change 1988 to 2012 | (8) Change 1996 to 2012 |
|---|---|---|---|---|---|---|---|---|
| Var. of log earnings | 0.77 (100%) | 0.77 (100%) | 0.72 (100%) | 0.66 (100%) | 0.57 (100%) | 0.52 (100%) | -0.25 (100%) | -0.20 (100%) |
| Var. of worker effects | 0.39 (51%) | 0.38 (49.2%) | 0.36 (50%) | 0.35 (54%) | 0.33 (58%) | 0.31 (60%) | -0.09 (34%) | -0.05 (24%) |
| Var. of firm effects | 0.16 (21%) | 0.18 (23.1%) | 0.17 (23%) | 0.13 (20%) | 0.10 (17%) | 0.08 (15%) | -0.08 (33%) | -0.09 (45%) |
| Var. of year effects | 0.02 (2%) | 0.01 (2%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | -0.02 (6%) | 0.00 (0%) |
| 2×Cov. worker and firm | 0.14 (18%) | 0.14 (18%) | 0.14 (20%) | 0.12 (18%) | 0.10 (18%) | 0.10 (20%) | -0.04 (14%) | -0.04 (22%) |
| 2×Cov. worker and year | -0.01 (-2%) | -0.02 (-2%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.01 (-2%) | 0.00 (0%) |
| 2×Cov. firm and year | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0.%) | 0.00 (0%) | 0.00 (0%) | 0.00 (-1%) | 0.00 (-1%) |
| Var. of residual | 0.08 (10%) | 0.07 (10%) | 0.06 (8%) | 0.05 (7%) | 0.04 (7%) | 0.04 (7%) | -0.04 (17%) | -0.02 (10%) |
| # worker years | 85.4 | 85.6 | 90.2 | 102.0 | 123.7 | 151.0 | | |
| # firm years | 0.98 | 1.04 | 1.23 | 1.44 | 1.75 | 2.22 | | |
| $R^2$ | 0.90 | 0.91 | 0.92 | 0.93 | 0.93 | 0.93 | | |

Note: Variance decomposition is $Var\left(y_{it}\right) = Var\left(a_i\right) + Var\left(a_j\right) + Var\left(\hat{\gamma}_t\right) + Var\left(e_{it}\right) + 2Cov\left(a_i, \hat{\gamma}_t\right) + 2Cov\left(a_j, \hat{\gamma}_t\right) + 2Cov\left(a_j, \hat{\gamma}_t\right)$. Cells contain variance explained by each decomposition element. The share of the total variance explained by each decomposition element is given in parentheses.

all population. The overall variance of log earnings is four log points higher in the PIA subpopulation in the 1996–2000 subperiod and falls by 19 log points over the subsequent three subperiods (versus 20 log points in the overall population). The variance of worker effects is three log points higher in the 1996–2000 subperiod and four log points greater in the 2008–2012 subperiod. The variance of firm effects is two log points less in 1996–2000 and falls by eight instead of nine log points. We conclude from this that trends in inequality are similar in the PIA subpopulation as in the overall population.

### 2.6.2 The link between firm effects and firm characteristics

Given that a reduction in dispersion of firm effects has been an important element in the decline, a natural question is whether differences in pay at the firm level are related to observable characteristics of the firm and whether changes in such observable characteristics explain the changes we observe in the firm-specific component of pay over time. Using measures of firm performance from the PIA data, Table 2.7 reports results from regressing estimated firm effects on log value added per worker both controlling for and not controlling for two-digit subsectors.[28]

Several features are worth highlighting. First, greater value added per worker is associated with a higher firm component of pay. Second, the amount of dispersion in firm effects explained by firm performance is notable, with an $R^2$ above 0.6 (around 0.5 without subsector controls). Third, the pass-through from value added per worker to the firm component of pay declines substantially over time. In 1996–2000, a one log point increase in value added per worker was associated with a 0.24 log points increase in estimated firm effects (0.26 without subsector controls); in 2008–2012, the same increase in value added per worker was only associated with a 0.13 log point higher firm effect (0.15 without subsector controls).

[28]We find that other firm characteristics only marginally improve the explanatory power of the model. Moreover, the variance of the predicted values due to other firm characteristics is essentially zero, and hence we do not report them here.

Table 2.6: AKM decomposition results for workers at large manufacturing&mining firms (as fraction of largest connected set)

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | 1996–2000 | 2000–2004 | 2004–2008 | 2008–2012 | Cumulative change from 1996 to 2012 |
| Variance of log earnings (% of pop. estimate) | 0.76 (100%) | 0.74 (100%) | 0.65 (100%) | 0.58 (100%) | -0.19 (100%) |
| Variance of worker effects (% of pop. estimate) | 0.39 (51%) | 0.39 (53%) | 0.38 (58%) | 0.35 (61%) | -0.04 (20%) |
| Variance of firm effects (% of pop. estimate) | 0.15 (19%) | 0.12 (16%) | 0.08 (13%) | 0.07 (11%) | -0.08 (43%) |
| Variance of year effects (% of pop. estimate) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) |
| 2×Cov. worker and firm effects (% of pop. estimate) | 0.18 (23%) | 0.18 (24%) | 0.15 (23%) | 0.12 (21%) | -0.06 (30%) |
| 2×Cov. worker and year effects (% of pop. estimate) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) |
| 2×Cov. firm and year effects (% of pop. estimate) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (0%) | 0.00 (-1%) |
| Variance of residual (% of pop. estimate) | 0.05 (7%) | 0.05 (6%) | 0.04 (6%) | 0.04 (6%) | -0.02 (10%) |
| # Worker years (% of pop. estimate) | 16.60 (18%) | 17.90 (18%) | 22.80 (18%) | 26.30 (17%) | |
| # Firm years (% of pop. estimate) | 0.11 (3%) | 0.13 (3%) | 0.16 (3%) | 0.18 (2%) | |
| $R^2$ (% of pop. estimate) | 0.93 (101%) | 0.94 (101%) | 0.94 (101%) | 0.94 (101%) | |

Note: Variance decomposition of AKM model estimated using manufacturing firms covered by PIA. The ratio between estimates using manufacturing firms relative to AKM estimates using all sectors is given in parentheses. Worker-years, unique workers, firm-years, and unique firms are in millions.

Table 2.7: Regression of estimated firm effects on firm productivity

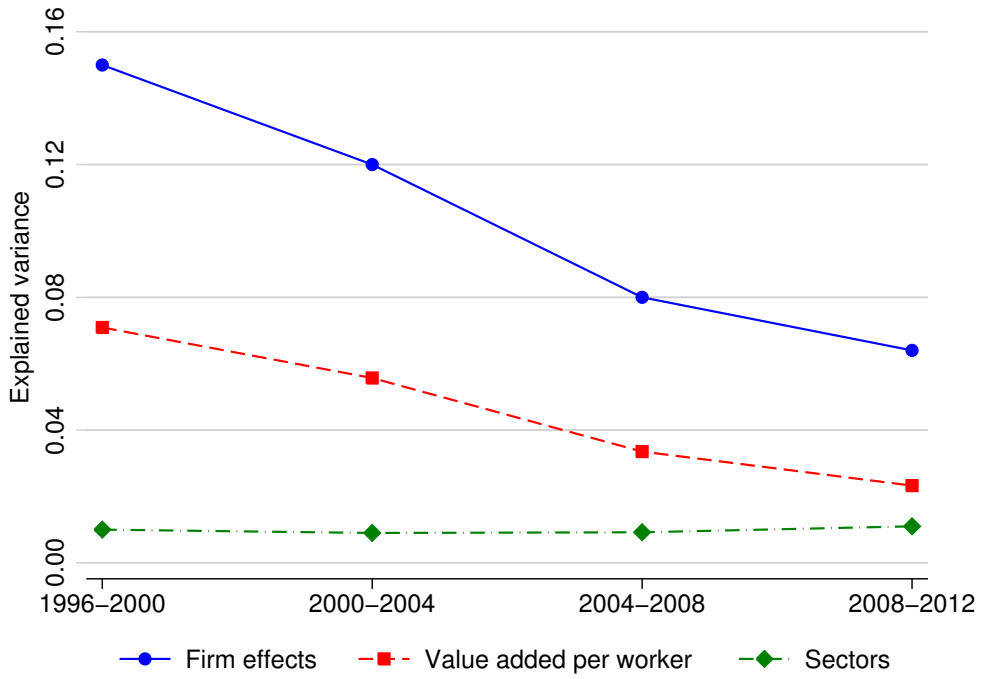
| | 1996–2000 | | 2000–2004 | | 2004–2008 | | 2008–2012 | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Value added p.w. | 0.258 | 0.236 | 0.201 | 0.182 | 0.163 | 0.143 | 0.147 | 0.126 |
| Sector controls | No | Yes | No | Yes | No | Yes | No | Yes |
| Worker-years | 16.6 | 16.6 | 17.9 | 17.9 | 22.8 | 22.8 | 26.3 | 26.3 |
| $R^2$ | 0.583 | 0.655 | 0.562 | 0.650 | 0.537 | 0.639 | 0.483 | 0.624 |

Note: Dependent variable is the estimated firm effect $a_j$. Independent variable is log value added per worker. Number of worker-years in millions.

Figure 2.8 evaluates the implication of the estimation results in Table 2.7 for the compression in firm pay in Brazil. It plots the overall variance of firm effects (solid blue line with circles), the predicted variance only due to value added per worker (dashed red line with squares), and the predicted variance only due to sectors (dash-dotted green line with diamonds) across periods. Value added per worker explains almost half of the initial variance in firm effects, while sectors initially explain seven percent. Between 1996–2000 and 2008–2012, the variance explained by value added per worker declines by five log points or 25 percent of the overall decline in the variance of log earnings in Brazil over this period. The variance due to sectoral variation, on the other hand, is small and remains constant in levels.[29]

[29]We separately investigated the role of sectors in the overall economy, not just the manufacturing and mining sectors, and find a significant role for variation in firm-level pay within sectors; results are available upon request.

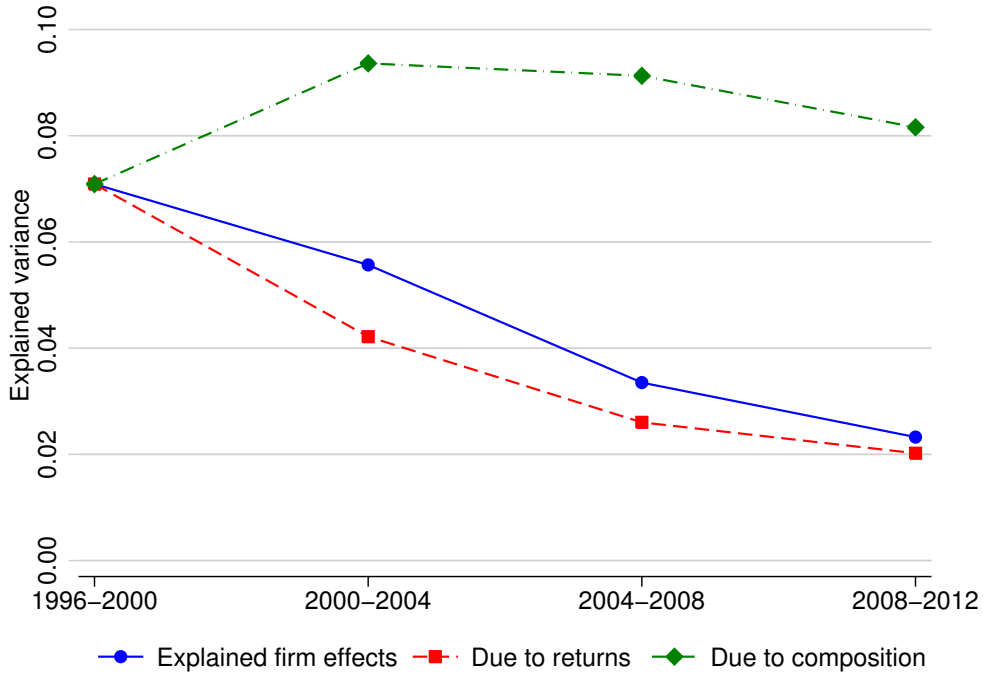
Figure 2.8: Regression of firm effects on productivity and sector controls




To quantify the importance of changes in the firm productivity distribution versus the pass-through from productivity to pay, we consider the two counterfactual exercises outlined in Section 2.5. Thus, we first hold the pass-through from value added per worker to the firm component of pay constant at its estimated 1996–2000 value and allow only the variance of value added per worker to change as in the data. Second, we hold the variance of value added per worker fixed and change only the pass-through to match the data.

Figure 2.9 plots the result from this exercise. The predicted variance of firm effects due to value added per worker (solid blue line with circles) declines from seven to two log points, the predicted variance from value added per worker holding pass-through constant (dashed red line with squares) increases from seven to eight log points, and the predicted variance holding the dispersion in value added per worker constant (dash-dotted green line with diamonds) falls slightly more than the total predicted variance. We conclude that, ceteris paribus, a declining pass-through from firm performance to pay contributed significantly to reduced earnings inequality in Brazil during this period.

Figure 2.9: Variance of predicted firm effects, variance predicted by fixed pass-through and variance predicted by fixed productivity distribution




### 2.6.3 The link between worker effects and worker characteristics

We finally turn to our results regarding the link between observable worker characteristics and the worker component of pay. Table 2.8 shows the result from a regression of the estimated worker effects on age and education with and without occupation controls. We note the following: First, more able workers, as measured by labor market experience or education level, are paid more. For instance, the oldest age group earns on average 50 log point more than the youngest age group conditional on place of work. Secondly, such worker characteristics explain roughly 40 percent of the variance in worker effects, leaving substantial room for unobserved worker heterogeneity. Third, both the age gradient and the return to education declines over time, with a particularly pronounced fall in the latter. For instance, a college degree was associated with a 116 log point higher worker component of pay in 1996–2000 but only a 78 log point higher worker effect in 2008–2012.

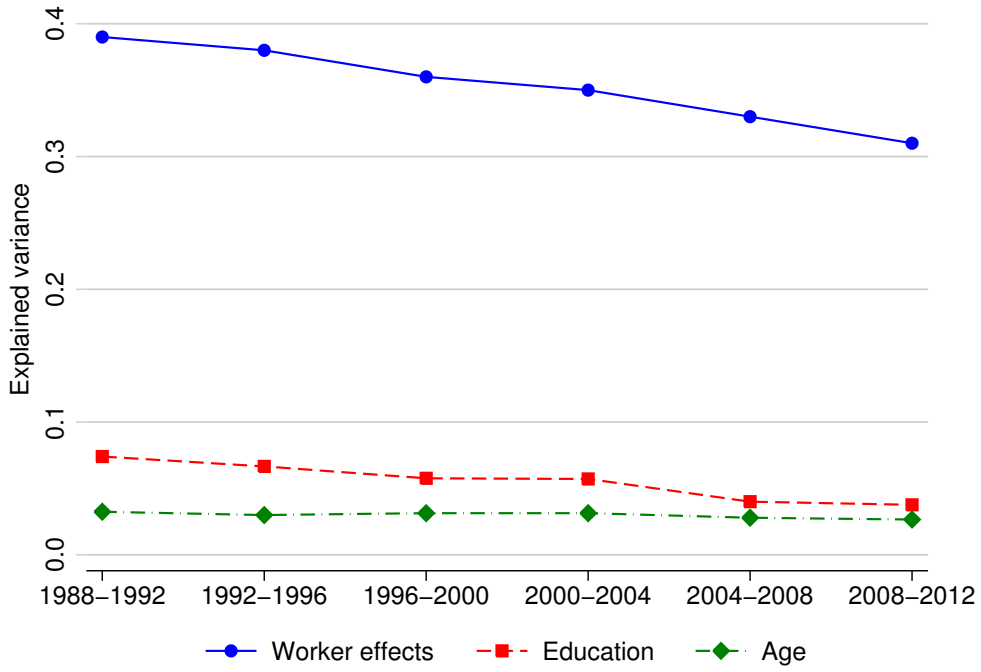
Table 2.8: Regression of estimated worker effects on worker characteristics

| | 1988–1992 | | 1992–1996 | | 1996–2000 | | 2000–2004 | | 2004–2008 | | 2008–2012 | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) | (12) |
| *Age groups* | | | | | | | | | | | | |
| 25–29 | 0.20 | 0.18 | 0.18 | 0.18 | 0.19 | 0.18 | 0.18 | 0.18 | 0.16 | 0.15 | 0.16 | 0.15 |
| 30–39 | 0.39 | 0.35 | 0.35 | 0.32 | 0.35 | 0.32 | 0.34 | 0.32 | 0.31 | 0.29 | 0.30 | 0.28 |
| 40–49 | 0.52 | 0.47 | 0.50 | 0.46 | 0.51 | 0.47 | 0.49 | 0.45 | 0.45 | 0.41 | 0.42 | 0.39 |
| 50–64 | 0.48 | 0.44 | 0.47 | 0.44 | 0.50 | 0.46 | 0.52 | 0.48 | 0.52 | 0.47 | 0.51 | 0.47 |
| *Education groups* | | | | | | | | | | | | |
| Middle school | 0.21 | 0.16 | 0.19 | 0.15 | 0.16 | 0.12 | 0.14 | 0.10 | 0.13 | 0.08 | 0.11 | 0.07 |
| High school | 0.61 | 0.51 | 0.56 | 0.46 | 0.45 | 0.37 | 0.37 | 0.29 | 0.32 | 0.21 | 0.27 | 0.19 |
| College or more | 1.21 | 1.10 | 1.19 | 1.04 | 1.16 | 1.00 | 1.19 | 1.00 | 1.15 | 0.82 | 1.10 | 0.78 |
| Occupation controls | No | Yes | No | Yes | No | Yes | No | Yes | No | Yes | No | Yes |
| # Workers | 85.4 | 85.4 | 85.6 | 85.6 | 90.2 | 90.2 | 102.0 | 102.0 | 123.7 | 123.7 | 151.0 | 151.0 |
| $R^2$ | 0.35 | 0.40 | 0.34 | 0.37 | 0.34 | 0.38 | 0.35 | 0.39 | 0.36 | 0.41 | 0.37 | 0.44 |

Note: Dependent variable is the estimated worker effect $a_i$. Number of workers in millions. Education estimates are relative to "less than middle school (<7 years)" category. Age estimates are relative to "age 18–24" category.

To quantify the contribution of the estimation results in Table 2.8 towards the declining dispersion in worker heterogeneity in Brazil, Figure 2.10 plots the overall variance of worker effects (solid blue line with circles), the predicted variance only due to age structure (dash-dotted green line with diamonds), and the predicted variance only due to variation in educational attainment (dashed red line with squares). We do not plot the variance due to occupation as this remains small and steady at around one, two log points over the entire period (in particular, it increases slightly from 0.019 in 1988–1992 to 0.023 in 2008–2012).[30] We find that while education and age explain only around 20–25 percent of total worker heterogeneity, they account for almost half of the decline in the variance of worker effects.

Figure 2.10: Regression of worker effects on age and education




We further decompose the decline in the dispersion due to age and education into one due to changes in returns and one due to changes in the composition of the workforce.[31] Figure 2.11 plots in the top (bottom) panel in solid blue with circles the overall compression

[30] Also not shown, the sum of covariances increases modestly from 0.032 in 1988–1992 to 0.040 in 2008–2012.

[31] Since the variance due to occupation is small and does not change much over time we do not present a decomposition of this.

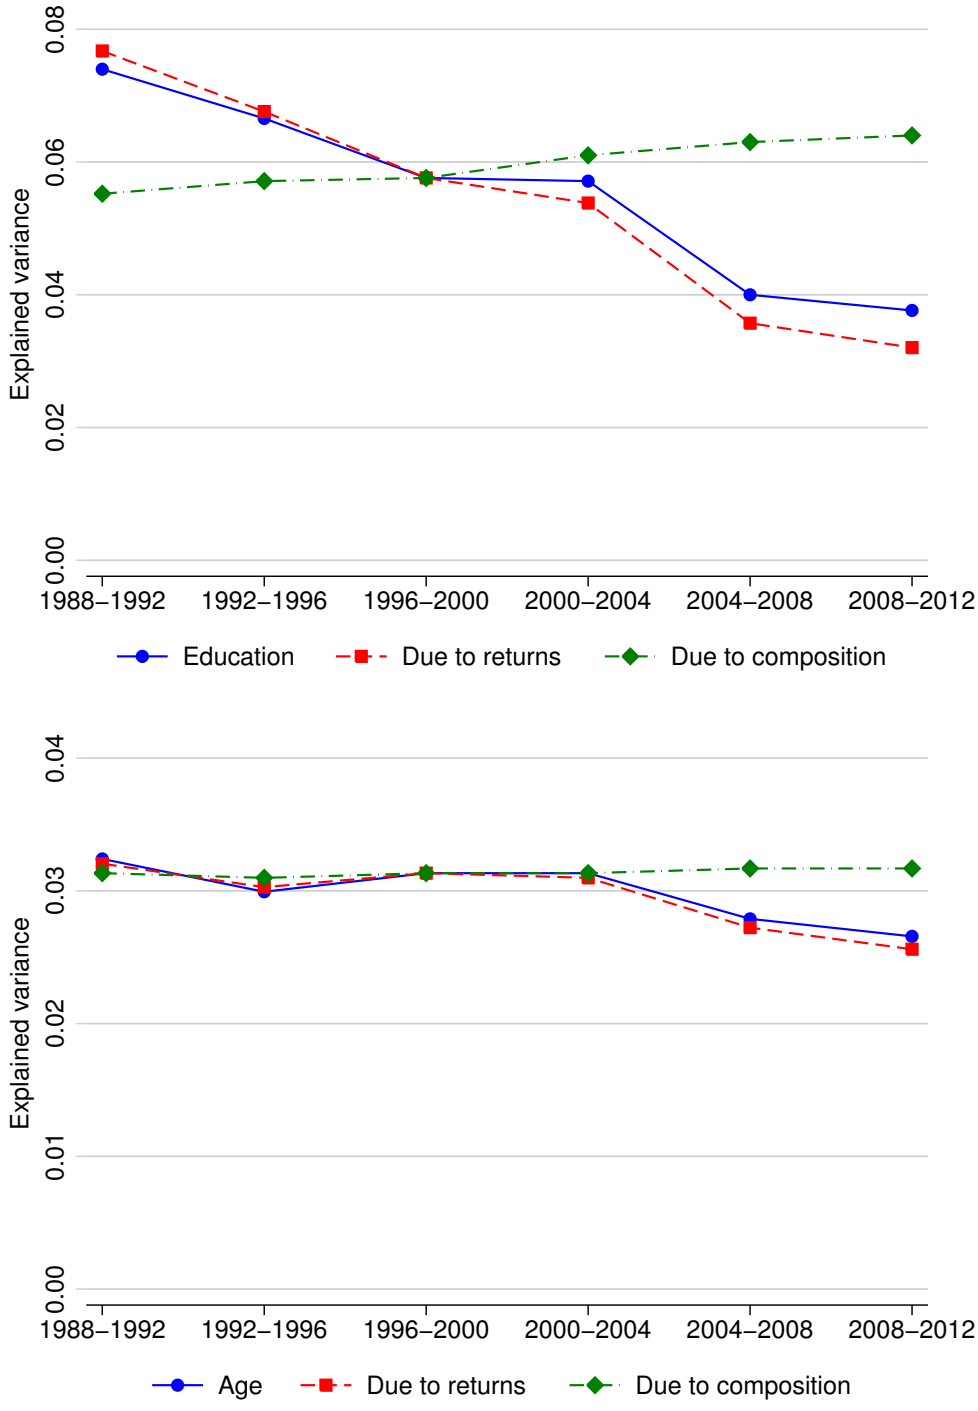
explained by education (age), in dash-dotted green with diamonds the compression resulting from holding the returns to education (age) fixed, and in dashed red with squares that when holding the underlying education (age) distribution fixed.

Holding returns fixed, changes over this period in the age structure as well as the educational attainment of the workforce resulted in a small net increase in inequality. Given the large increase in educational attainment documented earlier, it might come as a surprise that this did not induce a larger change in the underlying distribution of educational attainment. This is the result of two offsetting forces: on the one hand, the share of the lowest education group shrinks, but on the other more workers obtain high-school or college degrees, which earn more than the average. Altogether, these two effects have small positive net effects on inequality.

At the same time the rapid decline in the return to education—and to a lesser extent age—resulted in significant earnings compression. In fact as can be seen by comparing the solid blue lines with circles to the dashed red line with squares, the compression in returns resulted in an even larger fall in inequality than the overall fall. Holding everything else constant, the fall in the return to education (age) explains 13 (three) percent of the overall decline in the variance of log earnings between 1996–2000 and 2008–2012.[32]

[32]Of course, the compression in the return to education could have partly resulted from changes in educational attainment. The counterfactual of holding the distribution constant while changing returns should be interpreted with this qualification in mind.

Figure 2.11: Decomposition of variance of worker effects predicted by age (top) and education (bottom)



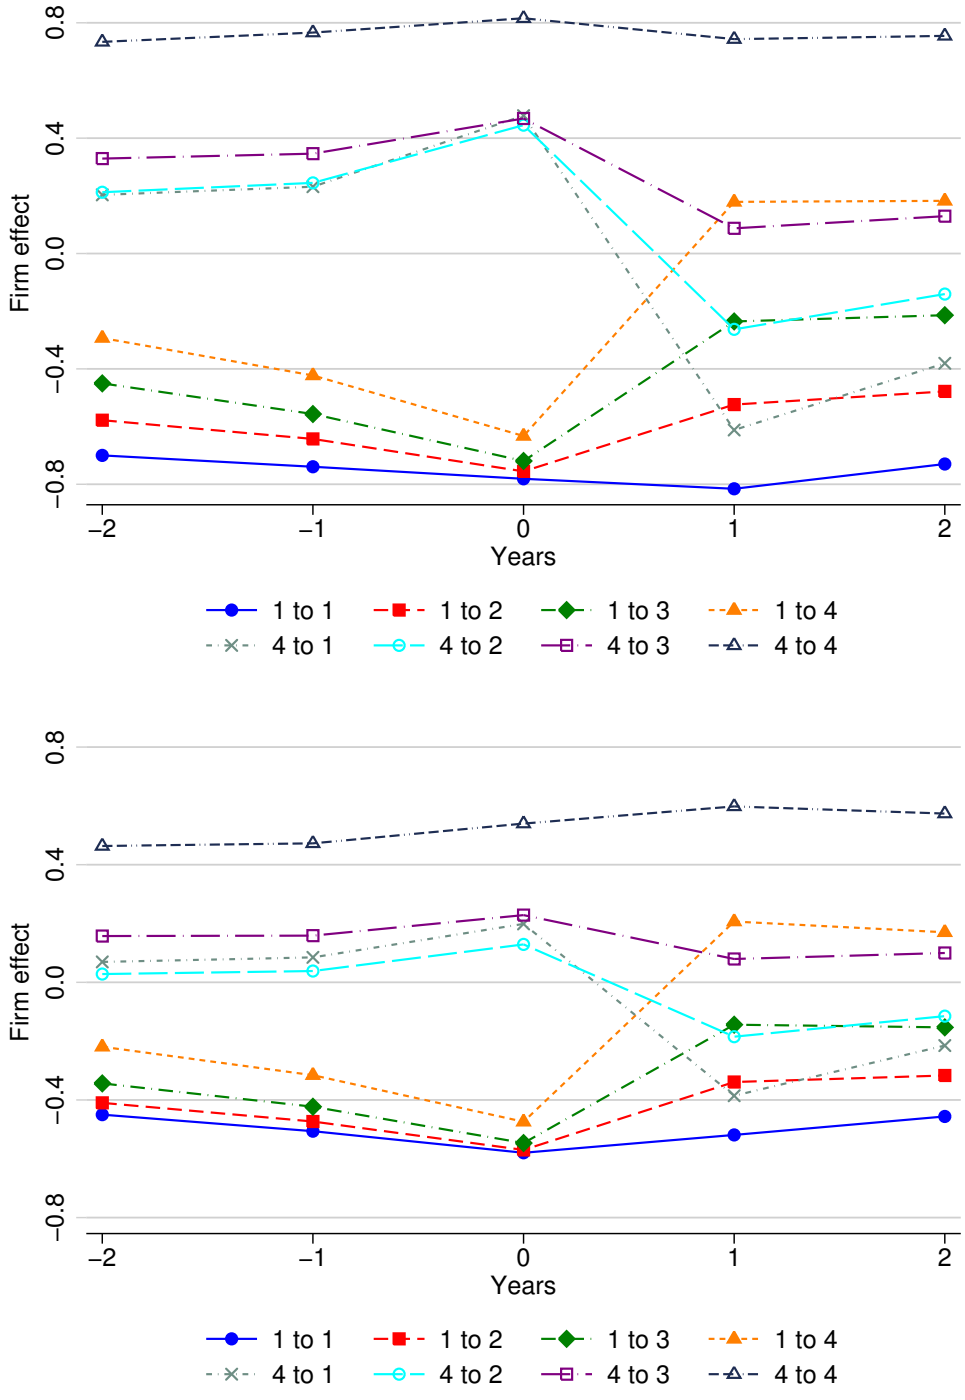
### 2.6.4 Empirical support for the AKM model

A large part of our analysis has been based on estimation results from repeated applications of the AKM framework. Although the consistently high $R^2$ coefficient of our model suggests that the model provides a good fit to the data, our estimates can be biased if the residual is correlated with either the firm or worker component of earnings. To investigate this possibility further, we replicate using our Brazilian data two diagnostics exercises proposed by Card, Heining and Kline (2013) for the case of Germany.

Figure 2.12 shows the average firm effect of workers who switch firms up to two years prior to the switch and two years after the switch for the first and last period of our sample. Switchers are classified by the firm effect quartile of the pre and post transition firms. Consistent with the AKM specification, workers that switch from the lowest quartile experience gains in firm effect and workers that switch from the highest quartile experience losses. Additionally, the gains of those switching up are similar to the losses of those making the reverse switch.[33] This suggests that our additive model is consistent with the pattern observed among workers who transition between firms.

[33]For expositional ease we only show switches out of the first and fourth quartile, but other quartiles display similar pattern and are available on request.

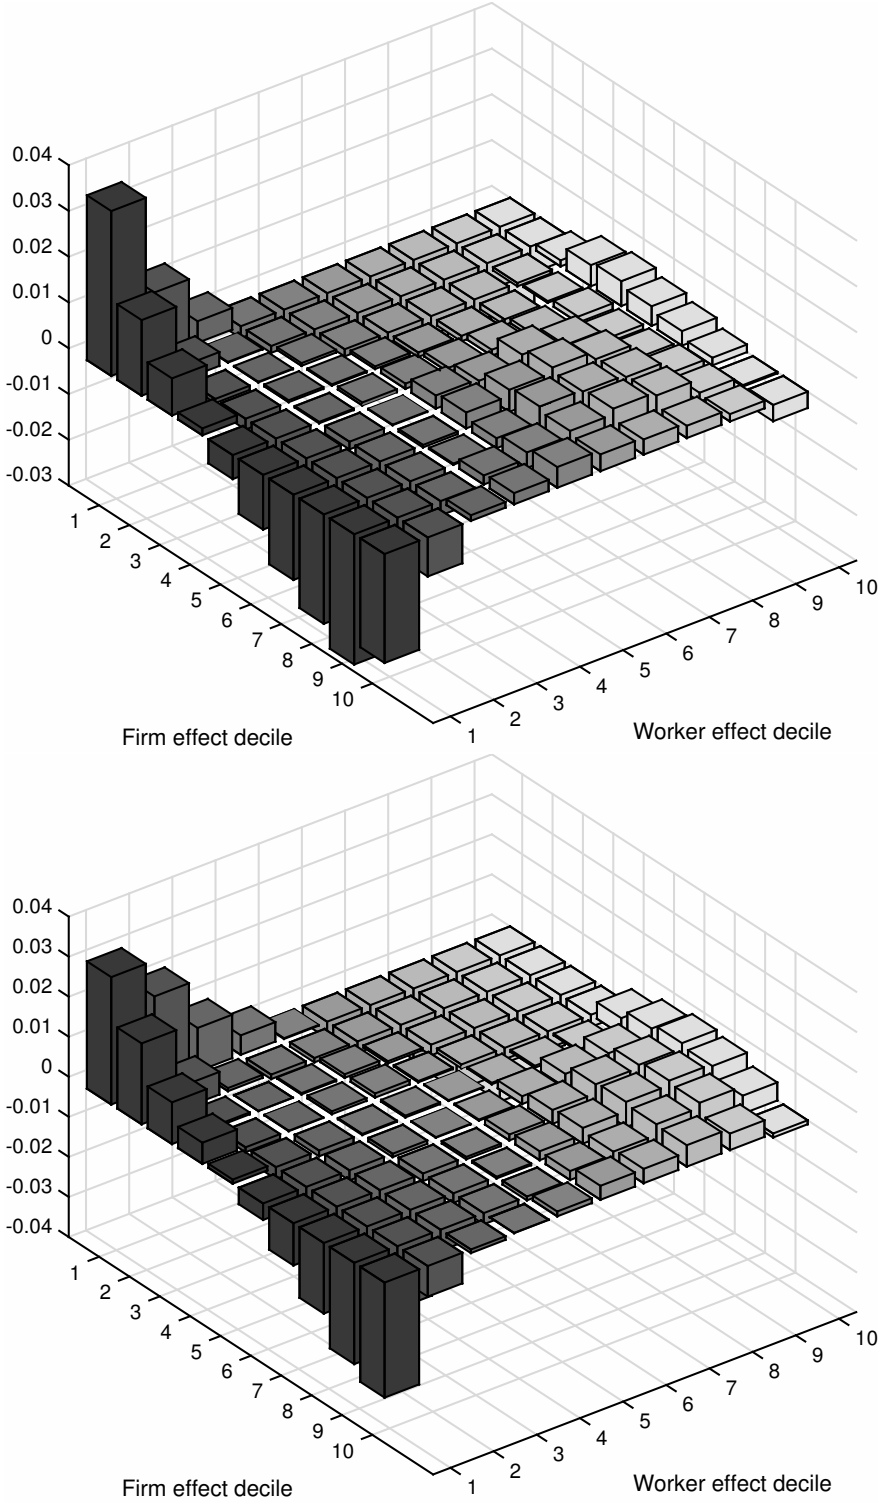
Figure 2.12: Average changes in firm effects of workers switching employers, 1996–2000 (top) and 2008–2012 (bottom)



Second, we find little evidence in the data for match effects that are systematically correlated with either person of match effects. Figure 2.13 shows the average estimated residual by decile of worker and firm effect. There is some evidence of misspecification for the lowest decile of workers in the sense that they display a systematically positive residual while working at the lowest paying firms. However, the magnitude of the error is modest and beyond the lowest decile of worker effects, errors do not exhibit any systematic relationship with firm and worker effects.[34] This reassures us that the log additive assumption is an appropriate characterization of Brazilian wage setting policies.

[34]The systematic variation among the lowest decile of worker effects is due to those individuals having a strongly positive error term when working at the lowest-paying firms, and strongly negative error term when working at the highest-paying firms. In other words, for lowest worker effects individuals the pay gradient across firms is very flat. We argue in Engbom and Moser (2015) that this pattern is consistent with a frictional labor market with a binding minimum wage.

Figure 2.13: AKM residual by firm and worker fixed effect deciles, 1988–1992 (top) and 2008–2012 (bottom)



## 2.7 Conclusion

In this paper, we document and dissect a large decline in earnings inequality in Brazil between 1996 and 2012. To understand the sources of this decline, we estimate a high-dimensional fixed effects model of earnings controlling for both worker and firm heterogeneity. Our results suggest that firms contributed more than proportionately to the decline. While more productive firms pay more, compression in firm productivity was not a factor behind declining inequality. Instead, we find that a declining pass-through from firm productivity to pay played an important role behind the decline. On the worker side, higher ability workers as measured by educational attainment or experience are paid more, yet also such measures of worker ability show no compression over time. Rather, a significant share of the explained decline in worker heterogeneity is due to rapidly falling returns to measures of worker ability. We conclude that changes in pay policies, rather than changes in firm and worker fundamentals, played a significant role in Brazil's inequality decline.

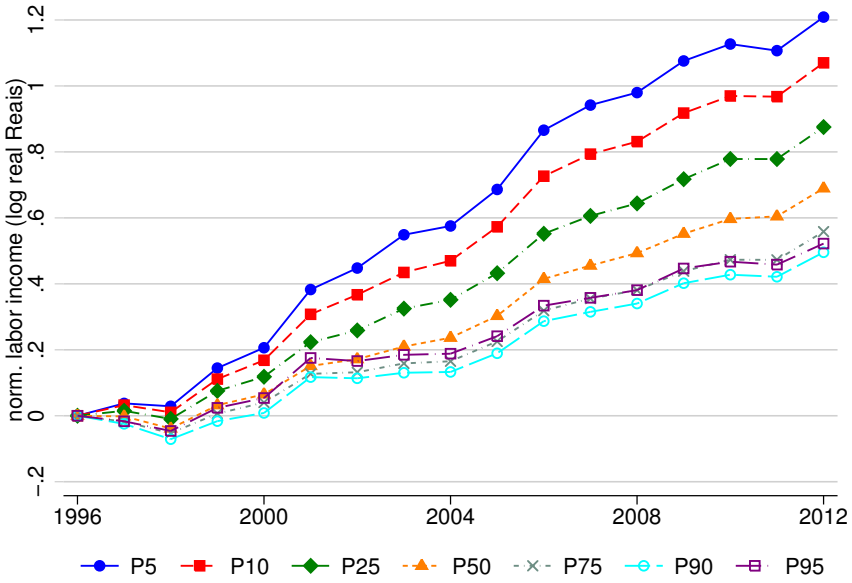
Our paper suggests a set of facts that a potential theory of the inequality decline in Brazil should be consistent with. Such a theory must generate pay differences between firms for identical workers, and they have to be strongly positively correlated with firm productivity. Moreover, any promising explanation needs to generate a compression in firm pay differences over time, not through compression in firm productivity but through a weaker link between productivity and pay. On the worker side, it should generate compression primarily through declining returns to ability, and not through a compression in underlying ability.

We think that promising candidates behind the decline in inequality in Brazil are changes in the nature of wage setting. In follow-up work Engbom and Moser (2015), we argue that the rapid rise in the minimum wage in Brazil during this period can explain a significant fraction of the decline in inequality, while being consistent with the stylized facts presented in this paper. Other institutional reforms, including changes in union bargaining structure and labor compensation laws, represent interesting avenues for future research.

## 2.8 Additional figures

### 2.8.1 Evolution of real earnings levels

Figure 2.14: Earnings levels evolution



### 2.8.2 Cross-section and time series comparison of alternative productivity measures

Figure 2.15: Cross-sectional comparison of alternative productivity measures

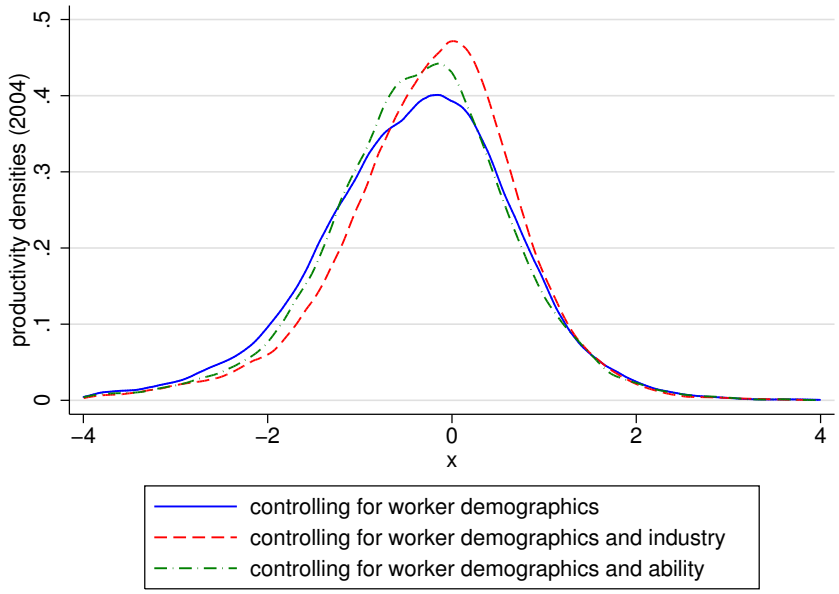
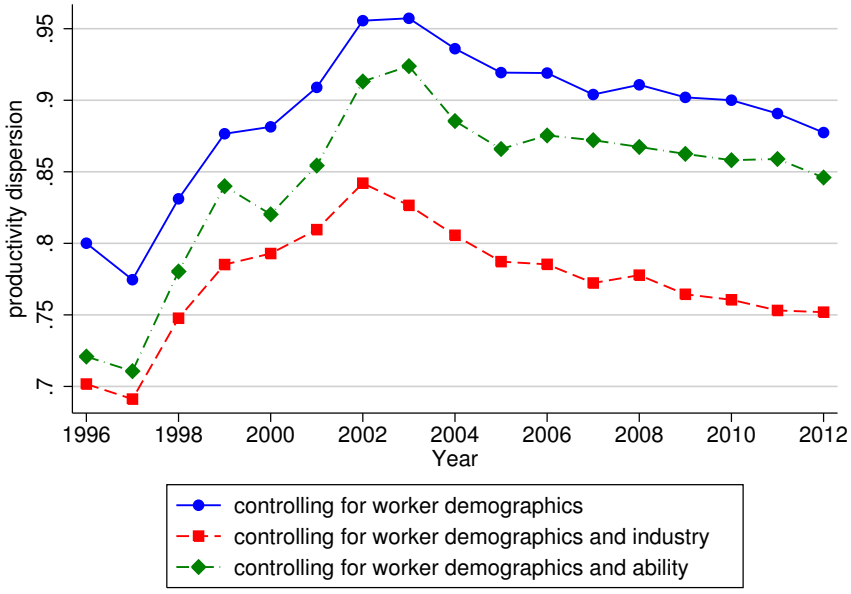
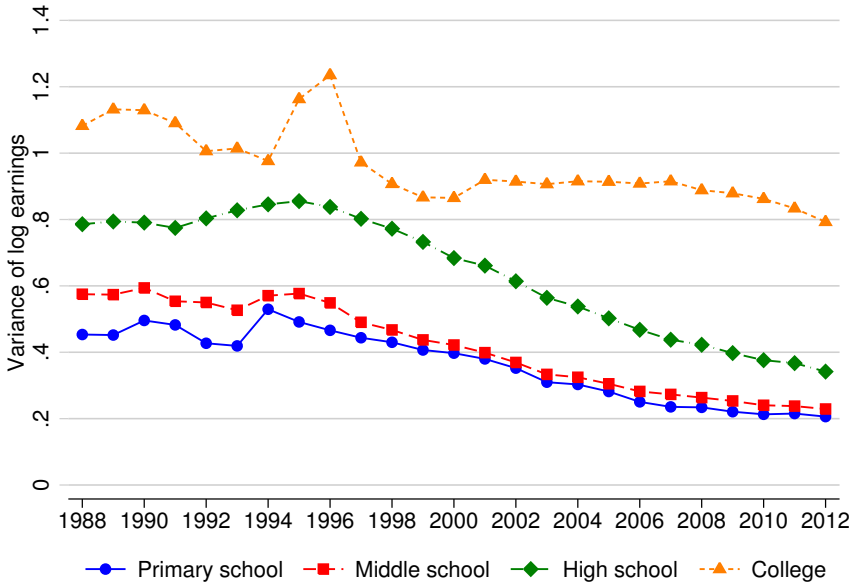



Figure 2.16: Evolution of dispersion of alternative productivity measures



### 2.8.3 Earnings inequality by education groups

Figure 2.17: Within education group inequality



# Chapter 3

# Unproductive Land Savings

## 3.1 Introduction

Most of the cross-sectional variation in incomes across countries is due to TFP differences.[1] Although the body of literature explaining this variation is both large and eclectic, an empirical fact has gained renewed attention in recent years: while the GDP per worker of the richest 5% of countries is 34 times that of the bottom 5%, this same ratio is a staggering 78 in the agricultural sector.[2] More perplexing yet is the fact that, while developing countries are far less productive in agriculture, they also devote a larger share of their resources to this sector.[3] This is not due to a lack of access to better technology. Within one country, the productivity gap between agriculture and other sectors varies by a factor of 10, suggesting that resources within a country could be transferred to a more productive sector but are not.[4] Even within the agricultural sector there is great variation among the productivity of individual farmers, with some large productive farms producing more efficiently by investing intermediate goods in their land. From a growth accounting perspective, there are too

[1] Klenow and Rodriguez-Claire (1997) and Hall and Jones (1999) argue that physical and human capital can account for less than 50% of total variation.

[2] The importance of agricultural productivity differences was first emphasized by Kuznets (1971), and more recently confirmed in a larger sample by Caselli (2005) and Restuccia, Yang, and Zhu (2008).

[3] Cai (2012) claims that the labor force share dedicated to agriculture varies from 85% in some poor countries to under 2% in some of the rich countries.

[4] Gollin, Parente and Rogerson (2004).

many farmers in poor countries wasting their land and labor on unproductive agriculture. Economists then measure this misallocation as aggregate productivity differences. But why do poor farmers not sell their land and labor to more productive sectors? Why do they stay small and unproductive? Why do poor farmers stay farmers? This paper argues that the lack of development in the financial sector may play a role in accounting for agricultural productivity differences, for small farmers may stay farmers not only due to lack of opportunity or education, but also due to lack of access to an effective savings technology.

Agricultural land is not only a means of production, but it is also a saving instrument. In an environment with institutional limitations, where informational or agency problems impede the issuing of other types of securities, farmers have a limited number of options to save. Most of the 2 billion people living under $2 a day do not even have access to commercial banking, and in most countries investment options are limited to low-interest accounts when available, some forms of physical capital, and land.[5] As a result, farmers, even if unable to produce high yields on their land or unwilling to make necessary investments to increase productivity, may nonetheless choose to hold land in order to save for the future. This paper presents a model where financial underdevelopment alone – and deficiencies in the availability of savings instruments in particular – leads to land misallocation and underinvestment, which in turn gives rise to agricultural productivity differences across countries.

The rest of the paper is organized as follows. Section 2 discusses how the mechanism above is related to recent literature on agricultural investments and productivity differences across countries, as well as productivity gaps between sectors. Section 3 presents a model that posits financial underdevelopment as a source of misallocation, underinvestment, and low agricultural productivity. Section 4 computes a numerical example to illustrate the model's mechanics. Section 5 presents cross-country empirical evidence consistent with the theory. Section 6 concludes and suggests avenues for future research.

[5]Demigurc-Kunt and Klapper (2012).

## 3.2 Literature review

This paper is related to a recent strand of literature that focuses on productivity gaps between the agricultural sector and the rest of the economy. Gollin, Parente and Rogerson (2004) describe the variation of these gaps across countries, and highlight the role of unmeasured home production as part of the reason why these are observed. Herrendorf and Schoellman (2011) analyze data from US states and conclude that the gap can mostly be explained by differences in human capital. In a similar line of thought, Lagakos and Waugh (2012) emphasize the role of subsistence levels in producing a friction in worker self-selection across sectors. In this story, poor workers that are unproductive in agricultural work may nonetheless self-select into the agricultural sector if they are poor enough, because they need to satisfy minimum consumption needs. Nevertheless, although human capital, measurability problems, and other explanations certainly do play a role in generating productivity gaps, it is recognized that this cannot be the whole story. Gollin, Lagakos and Waugh (2014), for example, find that even after adjusting for human capital differences between sectors, lower costs of living in agricultural areas[6], and hours worked, there are still large residual gaps in agriculture.

On the other hand, and closer to the reasoning of the model to be presented, Restuccia, Yang and Zhu (2008) and Restuccia and Adamopoulos (2013) highlight the role of intermediate goods in increasing farm productivity. They find that differences in import barriers can have a large impact on the use of productive intermediate goods and thus agricultural productivity.[7] Moreover, Donovan (2011, 2013) explains how, even without price distortions to intermediate goods, differences in productivity-increasing investments can arise from income levels alone. He observes that minimum subsistence levels can promote risk-averse behavior in poor agents, which lowers their propensity to invest in intermediate goods. This last mechanism will be key to the model below.

[6] As in Ravallion, Chen, and Sangraula (2009).

[7] Recent evidence from Anderson and Nelgen (2013) suggest that barriers to trade have been substantially reduced in recent years, adding importance to explaining productivity differences beyond these factors.

Furthermore, a related literature has emphasized the importance of land distribution in explaining the variation in agricultural productivity. One striking fact motivates this: the average farm size of the richest 20% of countries is 34 times that of the poorest 20%.[8] If large farms are more productive, this alone can have significant effects on overall productivity measures. Vollrath (2010) finds that an unequal distribution of land, when controlling for average farm size, can have a negative impact on productivity; moreover, Vollrath (2009) finds that land misallocation can account for up to 80% of variation in agricultural TFP. In a similar line of thought, Adamopoulos and Restuccia (2013) emphasize the role of land policies in explaining land misallocation. They find that decreasing returns to scale alone can only explain about 25% of the variation in farm sizes, and that size-ceilings on land ownership have a great effect on productivity. However, the persistence of land misallocation may not be entirely explained by institutional or historical failures, for it may persist even when the institutional environment is improved. Udry (2010), for example, argues that improvements in property right enforcement in Africa have not resulted in significant yield increases. Therefore, a theory that addresses why land misallocation is observed beyond regulatory barriers seems to be important in studying agricultural productivity differences.

The importance of financial underdevelopment in explaining this land misallocation is motivated by literature on savings in rural areas. Rosenzweig and Wolpin (1993), for example, argue that farmers in India use bullocks as a means of precautionary savings as a result of the poor options available to them. More recently, evidence from Anagol, Etang, and Karlan (2013) suggests that the lack of savings options may be quite severe in some countries. They show that in rural India farmers decide to invest in cows and buffaloes, even when these provide large negative returns.[9] Similarly, Dupas et. al. (2014) present evidence on the existence of negative interest rate savings accounts in Kenya. The following model

[8] Restuccia (2011)

[9] They find negative returns of -6% and -13% for cows and buffalos even when labor is valued at zero. When labor is valued at market wages, these are -64% and -39%.

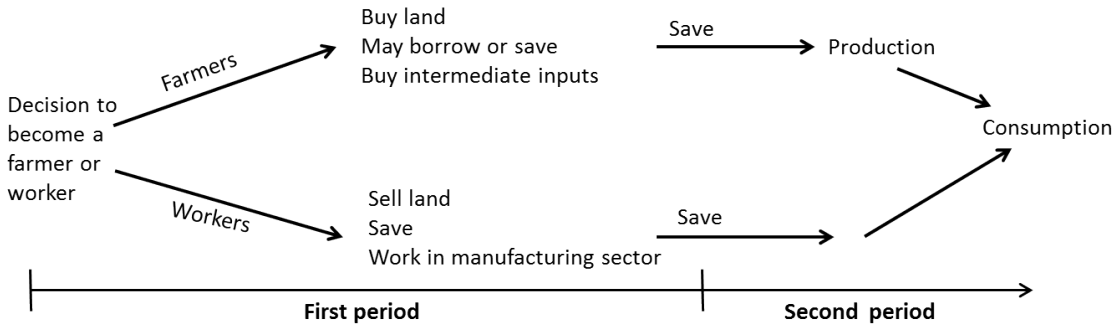
suggests that, in absence of trust-worthy storage technologies, land's role as a savings vehicle becomes more prominent, which hinders its role as a means of production.

## 3.3 Theoretical Model

This section presents a simple theoretical framework to study how deficiencies in the financial sector may lead to the misallocation of land, low investment in intermediate inputs, and ultimately low agricultural productivity. In this framework, land serves a dual function. On one hand, it is a savings mechanism that provides an alternative to ordinary savings accounts or holding cash. On the other, land is a resource needed for production. The interaction of these two functions induces land misallocation, since there are agents who hold land without investing much in it, hence lowering overall productivity. Furthermore, by providing an alternative explanation to differences in the distribution of land, the model departs from heterogeneous agent models that assume variation in labor productivity, skill, or idiosyncratic parameters in the preferences. In contrast, this model shows that an environment with equally capable agents who have differentiated access to reliable financial services – and savings assets in particular – can generate similar results.

We consider a small open economy where goods are traded internationally. Figure 3.1 shows the timing of the model. In the first period, agents can become either farmers or workers in the manufacturing sector. In this period, they also trade land among themselves and buy intermediate inputs from international markets. In the second period, production and consumption take place. In this economy, land can only be bought to be used in agricultural production so that land rental markets are shut down. In other words, due to either informational frictions or transactions costs, farmers are unable to rent their land to large agricultural producers; thus, each farmer must produce on the land they buy. The full environment is now described in detail.

Figure 3.1: Timing of the Model




### 3.3.1 Preferences

There is a continuum of agents with unit mass, where each agent has identical preferences over agricultural and manufacturing goods. Following the literature on structural change and the dual economy, agents have a subsistence requirement of $\overline{a}$ agricultural goods; hence, preferences are non-homothetic and the share of expenditure on agricultural goods will be decreasing with wealth.[10] In particular, preferences over consumption goods are given by

$$U(c) = \alpha \log(c_a^i - \overline{a}) + (1 - \alpha) \log(c_m^i)$$

where $c_a^i$ and $c_m^i$ stand for the consumption of agricultural and manufacturing goods of farmers ($i = f$) and manufacturing workers ($i = n$).

### 3.3.2 Endowments

Each agent is endowed with the same units of land, $\widehat{l}$, and $\widehat{m}$ units of manufacturing goods, where $\widehat{m}$ is drawn from a distribution with density $\phi_o(\widehat{m})$, and support $[\underline{m}, \overline{m}]$.[11] In addition, agents are endowed with one indivisible unit of labor that they can spend either in farming

[10] Examples of the use of these preferences are found in Laitner (2000), Kongsamut, Rebelo, and Xie (2001), Gollin, Parente, and Rogerson (2007), Donovan (2011, 2012), Cai (2012), among others.

[11] $\underline{m} > 0$

$(F = 1)$, or as workers in the manufacturing sector $(F = 0)$ which pays an exogenous wage of $w$. [12]

### 3.3.3 Technology

#### 3.3.3.1 Agricultural production

An agent who decides to become a farmer has access to a production technology that takes the form

$$f(x, l, z) = A[\theta(xz)^{\rho} + (1 - \theta)(l)^{\rho}]^{\frac{\gamma}{\rho}}$$

where $x$ are intermediate inputs (e.g. fertilizer), $l$ is the land used in production, $\theta$ captures the relative importance of intermediate goods, and $\rho$ determines the elasticity of substitution.[13] Moreover, $\gamma$ sets the extent to which there are decreasing returns to scale, and $A$ is the total productivity factor in the economy. In the model, one unit of labor from the agent is required to have access to the technology and produce in a farm. Since we are interested in how land owners make decisions about optimal farm size (in terms of land) and intensity in input use, we abstract from the optimal choice of agricultural workers within a farm.[14]

Crucially,$z$ is an idiosyncratic shock[15] with density $h_z(z)$, support $[\underline{z}, \overline{z}]$, and $E(z) > 0$ that affects the productivity of intermediate inputs for an individual farmer.[16] That is, intermediate inputs increase the expected yield, but they also exacerbate risk. This feature is

[12]Since we are interested in land misallocation dynamics in the agricultural sector, we abstract from explicitly modeling a manufacturing sector.

[13]Following Adamopoulos and Restuccia (2013), which also studies agricultural production across countries, the CES production technology is chosen. This is convenient, since farms with different intermediate input intensities are subject to different ouput risk.

[14]We also restrict the choice of farmers to being either full-time farmers or full-time manufacturing workers. This is at odds with nutritional wage effects described in Mirrlees (1975), Stiglitz (1976), Bliss and Stern (1978), among others. In this story, having a plot of land allows the worker to be employable in another sector by allowing him to satisfy the minimum subsistence requirement. In our model, however, this would not alter the results and would only magnify the effect of unproductive land savings.

[15]Since this is a small open economy, the analysis remains unchanged if we consider a combination of aggregate and idiosyncratic shocks.

[16]In the presence of a minimum subsistence requirement and log utility, $\underline{z}$ must be high enough to induce some agents to become farmers.

motivated by empirical evidence on intermediate input use by farmers, particularly fertilizer. First, fertilizer does not appear to decrease the variance of yields (Taxler et al. (1995)) and may actually increase it (Just and Pope (1979)), which is consistent with the production function chosen. Second, recent evidence suggests that input choices are sensitive to risk in the presence of incomplete markets. Karlan et al. (2012), for example, study the factors that affect the use of fertilizer and other inputs in Ghana and conclude that exposure to uninsured risk is a key determinant for low intermediate input use. In a randomized trial, they offer insurance to farmers and find that when outcome risk is reduced, farmers spend more on fertilizer, chemicals, and other inputs. Similarly, Zerfu and Larson (2010) analyze a panel data set from Ethiopia and find a negative effect of incomplete insurance markets and uncertainty on outcomes – due to a lack of education or weakness in the institutions in charge of information diffusion– on fertilizer use. Altogether, this evidence highlights the reactiveness of intermediate inputs to risk, and motivate the use of a production function that associates volatility in output to the choice of $x$.

#### 3.3.3.2 Saving and borrowing

In the first period, everyone can borrow ($b$) or save ($s$). Since the model's objective is to analyze the connection between financial underdevelopment and agricultural productivity, and not the origins of the former, this paper takes a reduced form approach to model financial markets. In particular, agents have access to different storage technologies depending on how wealthy they are. An agent with an endowment of $\widehat{m}$ has access to a technology that returns $r_s(\widehat{m})$ without any exposure to risk in the second period. The analysis in this paper is restricted to the case where $r_s(\widehat{m})$ is a weakly increasing function, which reflects the differentiated returns available to different parts of the wealth distribution. In particular, in an environment where banking has little presence and institutional failures prevent the proper functioning of equity markets, richer individuals may have access to higher returns on their savings than poorer ones.

Agents can also borrow resources to invest in land and intermediate inputs in the first period at a rate $r_b$, where $r_b > r_s(\widehat{m})$ for all possible endowments. Hence, though agents have different storage technologies, they have access to the same borrowing rate.[17]

### 3.3.4 Competitive equilibrium

A competitive equilibrium is now defined for a small open economy where the prices of manufacturing (numeraire), agricultural ($p_a$), and intermediate goods ($p_x$) are set internationally. Since agents only differ in terms of their endowment of manufacturing goods $\widehat{m}$, the equilibrium allocations are defined for each agent as a function of their endowment. Thus, given endowments, storage and borrowing technologies $r_s(\widehat{m}), r_b$, a manufacturing sector wage $w$, and international prices $p_a, p_x$, a *competitive equilibrium* is defined by a price of land $p_l$, and a set of allocations $F(\widehat{m})$, $c_a^n(\widehat{m})$, $c_m^n(\widehat{m})$, $c_a^f(\widehat{m}, z)$, $c_m^f(\widehat{m}, z)$, $s(\widehat{m})$, $l(\widehat{m})$, $x(\widehat{m})$, $b(\widehat{m})$ such that:

1. Given prices, $F(\widehat{m})$, $c_a^n(\widehat{m})$, $c_m^n(\widehat{m})$, $c_a^f(\widehat{m}, z)$, $c_m^f(\widehat{m}, z)$, $s(\widehat{m})$, $l(\widehat{m})$, $x(\widehat{m})$, $b(\widehat{m})$ solve the agent's problem:

$$\begin{aligned}
& \max_{x,l,s,F\varepsilon\{0,1\}} F * E\left(\max_{c_a^f, c_a^f} \alpha \log(c_a^f - \overline{a}) + (1-\alpha)\log(c_a^f)\right) \\
& +(1-F)\left(\max_{c_a^n, c_m^n} \alpha \log(c_a^n - \overline{a}) + (1-\alpha)\log(c_m^n)\right) \\
\text{s.t. } s + p_x x + p_l l &= \widehat{m} + p_l \widehat{l} + b \\
c_m^f + p_a c_a^f &= p_a f(x, l, z) + r_s(\widehat{m}) s - b r_b \\
c_m^n + p_a c_a^n &= r_s(\widehat{m})(w + \widehat{m} + p_l \widehat{l}) \\
b &\geq 0 \\
s &\geq 0
\end{aligned}$$

[17] The analysis is similar if we allow borrowing rates to differ across agents with different endowments, and this would only reinforce the results. Since we want to focus on the effect of storage return differences, we restrict ourselves to the case where $r_b(\widehat{m}) = r_b$.

2. Allocations $l(\widehat{m})$ fulfill the land market clearing condition:

$$\int \phi_o(\widehat{m}) l(\widehat{m}) d\widehat{m} = \int \phi_o(\widehat{m}) \widehat{l} d\widehat{m} \equiv L$$

### 3.3.5 The agent's problem

To solve for the agent's optimal allocation and working choices, we proceed by backward induction. First, we describe how agents allocate consumption given the resources available after storage returns and agricultural production are realized, and payments on existing debt made. Second, we analyze the optimal choice of inputs by a farmer. Third, we describe how agents make the decision to become farmers or workers.

#### 3.3.5.1 Consumption choice and risk aversion

Consider the problem in the second period, after production is realized but before consumption decisions are taken. Let $y$ be the sum of agricultural output, wages and storage returns available to the farmer in the second period after payments on his debt are made. In other words, $y$ reflects the resources available for consumption in the second period. For agents with enough ex-post resources $y$ to satisfy the subsistence requirement ($y > p_a \overline{a}$),[18] this leads to the optimal allocation rule,

$$\begin{aligned} c_a^i &= \overline{a} + \frac{\alpha}{p_a}(y - p_a \overline{a}) \\ c_m^i &= (1-\alpha)(y - p_a \overline{a}) \end{aligned}$$

Simply put, agents spend all of their resources on agricultural goods until they meet the subsistence requirement; then, they distribute the rest among the two goods according to the

[18] $q > p_a \overline{a}$ is assumed, so that everyone in this economy has enough resources to survive.

weight parameter $\alpha$. In what follows, it will be useful to define the indirect utility function $U(y)$:

$$\begin{aligned} U(y) &= \phi - \alpha \log(p_a) + \log(y - p_a \overline{a}) \\ \text{where } \phi &\equiv \alpha \log(\alpha) + (1-\alpha)\log(1-\alpha) \end{aligned}$$

Note that because of the presence of a subsistence requirement in agriculture goods, agents exhibit decreasing relative risk aversion (RRA) with respect to realized wealth.

$$RRA(y) = -y\frac{U''(y)}{U'(y)} = \frac{y}{y - p_a\overline{a}}$$

Without a subsistence requirement ($\overline{a} = 0$), we would have a Cobb-Douglas utility function and relative risk aversion would be constant. Nonetheless, with a subsistence requirement ($\overline{a} > 0$), the relative risk aversion is decreasing, so that wealthier individuals will be more tolerant to risk. To see this more clearly, the effect of risk on the choice of farmers in the first period can be decomposed by taking the Taylor expansion of the expected indirect utility expression:

$$E(U(y)) \approx \log(E(y - p_a\overline{a})) - \frac{Var(y)}{2(E(y - p_a\overline{a}))^2}$$

Thus, preferences exhibit variance aversion as reflected in the negative sign of the second term of the expansion. More importantly, the quadratic term in the denominator indicates that this aversion decreases when there is more wealth.

The relationship between wealth and agents' attitudes towards risk is crucial to the mechanics of the model for two reasons. First, because richer agents are more tolerant to risk, they will bear the increased risk of intermediate input use, which will ultimately produce a positive correlation between farm size and productivity. Second, risk-less manufacturing

work will be more attractive to the poorest workers, since they are the most averse to risk. It is important to point out that decreasing RRA is not only theoretically useful, but it is also a feature that has been observed empirically by the micro literature. Ogaki and Zhang (2001), for example, present evidence from farmers in India and Pakistan that confirms this behavioral pattern.

#### 3.3.5.2 The farmer's problem

We now focus on agents who decide to become farmers. Once they are farmers, their problem is reduced to a choice of optimal saving, borrowing, land and intermediate inputs to maximize expected utility in the second period. This can be simplified by using the indirect utility expression derived above, and writing second period utility in terms of ex-post resources $y^f$. Following this strategy, and combining two of the budget constraints, we get the following formulation of the problem:

$$\begin{aligned} & \max_{x,l} E(\log(y^f(x,l,z,\widehat{m},p_l) - p_a\overline{a})) \\ \text{s.t. } y^f(x,l,z,\widehat{m},p_l) &= p_a f(x,l,z) + r_s(\widehat{m})(\widehat{m} + p_l\widehat{l} + b - p_x x - p_l l) - br_b \\ b &\geq 0 \\ \widehat{m} + p_l\widehat{l} + b - p_x x - p_l l &\geq 0 \end{aligned}$$

That is, a farmer must choose intermediate inputs and land optimally, taking benefits, risks and costs into account. Taking first order conditions, we can start elucidating the trade-offs between borrowing, saving, land and intermediate input use faced by the farmer.

$$x \quad : \quad E(\frac{p_a[\theta(xz)^\rho + (1-\theta)(l)^\rho]^{\frac{\gamma}{\rho}-1}\gamma\theta z^p x^{\rho-1}}{y^f(x,l,z,\widehat{m},p_l) - p_a\overline{a}}) = \\ r_s(\widehat{m})p_x E(\frac{1}{y^f(x,l,z,\widehat{m},p_l) - p_a\overline{a}}) + p_x\lambda_s \tag{3.1}$$

$$l \quad : \quad E(\frac{p_a[\theta(xz)^\rho + (1-\theta)(l)^\rho]^{\frac{\gamma}{\rho}-1}\gamma(1-\theta)l^{\rho-1}}{y^f(x,l,z,\widehat{m},p_l) - p_a\overline{a}}) = \\ r_s(\widehat{m})p_l E(\frac{1}{y^f(x,l,z,\widehat{m},p_l) - p_a\overline{a}}) + p_l\lambda_s \tag{3.2}$$

$$b : (r_s(\widehat{m}) - r_b)E(\frac{1}{y^f(x,l,z,\widehat{m},p_l) - p_a\overline{a}}) = -\lambda_b - \lambda_s \tag{3.3}$$

A key variable of interest in our analysis is the productivity of farms owned by agents at different wealth levels. This will be determined by the intensity of intermediate input use per unit of land, so it is instructive to analyze what factors influence this choice. By combining equations (1) and (2), the following expression for the optimal use of intermediate goods per unit of land arises.

$$\frac{x}{l} = \left(\frac{\theta}{(1-\theta)}\frac{p_l}{p_x}\right)^{\frac{1}{1-\rho}} * \left(\frac{E(g(x,l,z,\widehat{m},p_l))}{E(g(x,l,z,\widehat{m},p_l)z^\rho)}\right)^{\frac{1}{1-\rho}(1-\frac{\gamma}{\rho})}. \tag{3.4}$$

$$\text{where } g(x,l,z,\widehat{m},p_l) = \frac{[\theta(xz)^\rho + (1-\theta)(l)^\rho]^{\frac{\gamma}{\rho}-1}}{y^f(x,l,z,\widehat{m},p_l) - p_a\overline{a}}.$$

Here, the variation in intermediate input use per unit of land is decomposed into two components. The first reflects an effect of relative prices that is consistent with the literature. Restuccia, Yang and Zhu (2008), for example, interpret $p_x$ as price distortions introduced often by policy. In this view, policies in several developing countries make intermediate goods

more expensive relative to the price of land and agricultural goods, mitigating their use. The second is a new term that reflects risk aversion induced by the minimum subsistence requirement present in the preferences. Since land is a risky investment and the use of intermediate goods increases risk, farmers take this into account when selecting the optimal intermediate input to land ratio. Because of the decreasing RRA pattern discussed before, the intensity of intermediate input use will depend on the level of wealth, with wealthier farmers willing to bear the risk of using more intermediate inputs per unit of land.

A second variable of interest is the optimal farm size (i.e. land demand) chosen by farmers at different levels of wealth. How much land is used by each farmer will depend on whether the farmer is rich enough to self-finance production or needs to borrow resources in order to do so. Thus, we now derive land demand for agents with different borrowing and saving choices.

The first thing to notice is that agents either save or borrow, but they do not do both.[19] Intuitively, nobody willing to borrow at a high interest rate would be willing to invest those resources at a lower rate. Therefore, depending on their endowment, farmers can adopt three possible behaviors. Richer farmers ($\lambda_s = 0$, $\lambda_b > 0$ ) take part of their endowment and put it in the saving technology available to them; moreover, they produce at a rate such that the marginal utility of an extra unit invested in land is the same as the return they could get in their savings technology. Their demand for land is thus directly affected by the rate of return on savings, as it can be seen by combining equations (2) and (3).

$$l^r = \left(\gamma(1-\theta)\frac{p_a}{p_l}\right)^{\frac{1}{1-\rho}} \left(\frac{1}{r_s(\widehat{m})}\right)^{\frac{1}{1-\rho}} \left(\frac{E(g(x,l,z,\widehat{m},p_l)))}{E(\frac{1}{y(x,l,z,\widehat{m},p_l)-p_a\overline{a}})}\right)^{\frac{1}{1-\rho}} \tag{3.5}$$

Similarly, the poorest farmers ($\lambda_s > 0$, $\lambda_b = 0$) need to borrow to finance intermediate goods and agricultural production. These agents make choices so that the marginal utility

[19]To see this, first consider equation 3. Since it must be true that agents consume above the minimum subsistence requirement $y(x,l,z,\widehat{m},p_l) > p_a\overline{a}$ for all agents and $r_s(\widehat{m}) - r_b < 0$, then it cannot be the case that the Lagrangian multipliers are both zero. Hence, agents must be either borrowers or savers, but they cannot be both.

from an extra unit invested in land equals the utility sacrificed from reducing future consumption because of borrowing. This intuition is reflected in the expression for land demand below, which is similar to the expression derived for the rich, but takes the opportunity cost of borrowing – not of saving– as an input.

$$l^p = \left(\gamma(1-\theta)\frac{p_a}{p_l}\right)^{\frac{1}{1-\rho}} \left(\frac{1}{r_b}\right)^{\frac{1}{1-\rho}} \left(\frac{E(g(x,l,z,\widehat{m},p_l)))}{E(\frac{1}{y(x,l,z,\widehat{m},p_l)-p_a\overline{a}})}\right)^{\frac{1}{1-\rho}} \tag{3.6}$$

Finally, agents in the middle of these two groups ($\lambda_s > 0, \lambda_b > 0$) will find borrowing rates too high to take a loan and saving rates too low to compete with the returns they can get by farming. As a result, they produce, but only do so by relying exclusively on self-finance. Their land demand $l^m$ is therefore in between the two groups described before.

$$\begin{aligned} l^p &< \left(\gamma(1-\theta)\frac{p_a}{p_l}\right)^{\frac{1}{1-\rho}} \left(\frac{1}{r_b}\right)^{\frac{1}{1-\rho}} \left(\frac{E(g(x,l,z,\widehat{m},p_l)))}{E(\frac{1}{y(x,l,z,\widehat{m},p_l)-p_a\overline{a}})r_b - \lambda_b}\right)^{\frac{1}{1-\rho}} = \\ l^m &= \left(\gamma(1-\theta)\frac{p_a}{p_l}\right)^{\frac{1}{1-\rho}} \left(\frac{1}{r_s(\widehat{m})}\right)^{\frac{1}{1-\rho}} \left(\frac{E(g(x,l,z,\widehat{m},p_l)))}{E(\frac{1}{y(x,l,z,\widehat{m},p_l)-p_a\overline{a}}) + \frac{\lambda_s}{r_s(\widehat{m})}}\right)^{\frac{1}{1-\rho}} \\ &< l^r \end{aligned}$$

The analysis above highlights how farmers balance trade-offs between risk free storage, borrowing and land investments. While poor farmers' decisions on optimal farm size will rely on the borrowing rate, this is not the case for middle and richer farmers. Moreover, similar to the use of intermediate goods per unit of land, the optimal size decision for each type of farmer can be decomposed into three components: first, the effect of the relative price of land with respect to the price of output; second, the opportunity cost from borrowing or saving; and third, the risk term from land investments. As before, this last term will depend on initial wealth, which is key to our analysis. In the model, wealthier individuals are willing to take more risk; therefore, they demand more land and invest more intermediate inputs in their farms. A corollary of this is that a piece of land is more valuable in the hands of a rich

farmer than in the hands of a poor farmer, and richer farmers choose bigger farm sizes. This has two implications. On one hand, larger farms owned by the rich will be more productive. On the other, since farming is less valuable to the poor, they will be the first to leave to the manufacturing sector. We analyze this choice next.

#### 3.3.5.3 Occupational choice

The agent's decision to become a farmer or a worker in manufacturing depends on how the utility from working in manufacturing compares to the expected utility from farming, given the optimal borrowing, saving, and input use decisions derived above. The occupational choice of the farmer is therefore described by the function $F(\widehat{m})$, where $F\varepsilon\{0,1\}$ is the amount of labor dedicated to farming, and

$$F(\widehat{m}) = \left\{ \begin{array}{ll} 1 & E(\log(y^f(x(\widehat{m}), l(\widehat{m}), z, \widehat{m}, p_l) - p_a\overline{a})) > \log(y^n(\widehat{m}, p_l) - p_a\overline{a}) \\ 0 & \text{otherwise} \end{array} \right\}$$

As argued before, farms from wealthier agents are more productive and – because of decreasing RRA in wealth – farming is more valuable to the rich. Thus, the decision to join farming will be determined by a cutoff $\overline{c}$, where $F(\widehat{m}) = 1$ for all $\widehat{m} \geq \overline{c}$ and zero otherwise. The key here is that the number of farmers in the economy - and therefore average farm size- is a function of the difference between expected utilities in the two sectors. In turn, this difference is a result of the mechanics of the agricultural sector, the outside option offered by wages in the manufacturing sector, and the level of financial development as reflected in the return of the storage technology available. It is important to highlight the role of savings returns in this setting. Given a fixed wage, prices, and parameters of the production function, a decrease in the return of storage (savings) decreases the benefits from working in the manufacturing sector (since $y^n(\widehat{m}, p_l) = r_s(\widehat{m})(w + \widehat{m} + p_l\widehat{l})$), making the farming sector relatively more attractive. Hence, a decrease in returns to storage will increase the number of (small) farmers, and reduce average farm size.

## 3.4 A numerical example

To illustrate the effects of the storage technology on land distribution and agricultural productivity, a numerical example is computed. The goal of this section is not to provide definitive quantitative predictions, but to illustrate the qualitative features described above. As a result, and given the abstract nature of the model, we do not undertake a formal calibration exercise but choose sensible parameters that illustrate the model's mechanics.

For the production function and preferences, we use parameters from Adamopoulos and Restuccia (2013), who calibrate the production function using US census data and use this function to analyze international differences in agricultural productivity. In particular, they use $\theta = .89$, $\gamma = .54$, and $\rho = .24$ to match agricultural land and capital shares as well as agricultural capital to land ratios. Since $x$ is interpreted as agricultural capital in their model, this parametrization exaggerates the role of intermediate inputs in the setting of this paper but retains the qualitative properties associated with decreasing returns to scale and the complementarity of land and intermediate inputs in production. Similarly, parameters $\overline{a} = .035$ and $\alpha = 0.01$ – which are calibrated to target the agricultural employment share in a different environment– are also used to illustrate qualitative properties. In this analysis, we have abstracted from international price and TFP effects and have therefore used $p_a = 1$, $p_x = 1$, and $A = 1$. Finally, for simplicity, the shock $z$ can take two values, $\{\overline{z}-\delta_z, \overline{z}+\delta_z\}$, with equal probabilities, where $\overline{z} = 1$ and $\delta_z$ is chosen ad hoc so that the risk of investing in land is low enough to encourage some agents to become farmers, but high enough to encourage some agents to become workers in the manufacturing sector given the other parameters of the model.[20]

[20]In the model, if there is no risk, every agent would chose to become a farmer since this would offer higher returns. Similarly, if there is too much risk, nobody would become a farmer. The value of $\delta_z = .25$ is chosen to illustrate the case where there are both manufacturing workers and farmers in the economy given the other parameters.

In the next subsection, aggregate productivity measures of competitive equilibria are computed for economies at different levels of financial development. Later, the mechanisms through which aggregate productivity differences arise are analyzed.

### 3.4.1 Allocations and aggregate productivity in competitive equilibrium

Consider an economy where manufacturing endowments are distributed uniformly, but the richer half of the population has access to a better storage technology than the poorer half. Concretely, the median of manufacturing endowments, $\widetilde{m}$, is such that $r_s(\widehat{m}) = r_s^r$ for all $\widehat{m} \geq \widetilde{m}$ and $r_s(\widehat{m}) = r_s^p = r_s^r - \delta_r$ for all $\widehat{m} < \widetilde{m}$. Intuitively, $\delta_r$ reflects the level of underdevelopment in the savings options for the poor. Since we do not have information on the rates of savings and borrowing available for different parts of the wealth distribution, $r_s^r = 6\%$ and $r_r^b = 10\%$ are chosen ad hoc at values that seem reasonable in magnitude given the data.[21] Figure 3.2 presents aggregate measures of the equilibria as we increase the gap $\delta_r$ between the storage returns available to the rich and poor. Thus, the savings rate available to the poor falls as we go from left to right along the graph towards a less developed financial system.

Several features are worth noting. First, the fall in the storage rate induces a fall in average farm size. This is because the fall in storage rates decreases the benefits of selling the land endowment and working in manufacturing. Thus, otherwise indifferent workers become farmers in this environment, which increases the total number of farmers and reduces land use per farmer. Second, the price of land increases as a result of the increased demand for land. This result stems from the fact that in this economy land is only used as an input for agricultural production. Financial underdevelopment leads to more farmers, and this in turn increases prices. However, in the presence of another sector that takes land as an input and is less developed in financially underdeveloped economies, this result could be

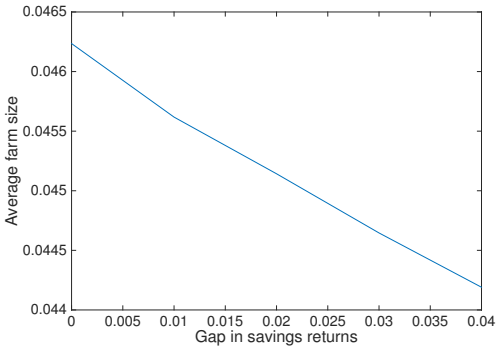
[21] In the data, the average gap between saving and borrowing rates is 4.4%.

eliminated without affecting the other predictions of the model. Since we are only interested in the dynamics of farm land distribution and agricultural productivity in this paper, we have abstracted from modeling other sectors in the economy.

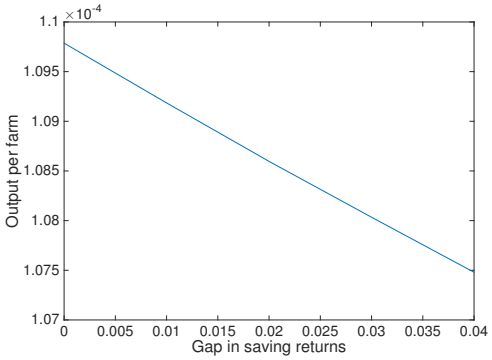
Third, and most importantly, financially underdeveloped economies are less productive at agriculture. This happens because more poor agents become unproductive farmers and demand more land when savings options for the poor worsen. As a consequence of having more poor farmers, aggregate output per farm and value added are reduced. This reasoning leads the main result of this section: that deficiencies in accessibility to high-quality savings technologies may lead to an unproductive land distribution. In other words, when the poor are not fully integrated into the financial system and have no access to good investment alternatives, they use land as a means of saving and lower agricultural productivity. We now proceed to study the mechanics of farmers' decisions that give rise to this equilibrium result.

Figure 3.2: Equilibrium outcomes by storage technology dispersion

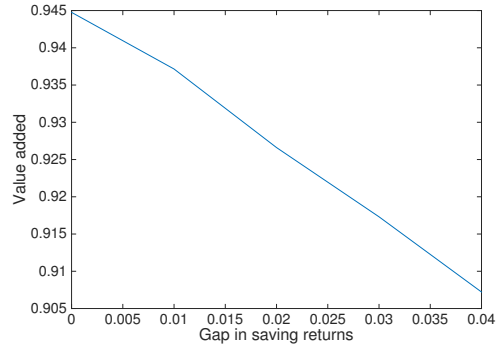
(a) Average farm size



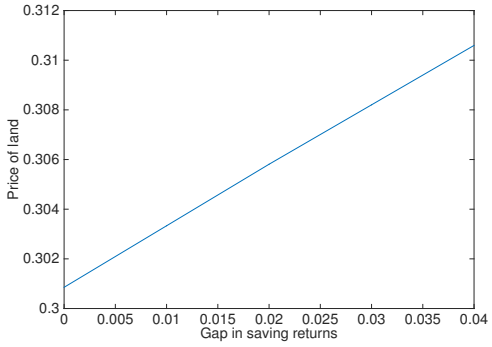

(b) Production per farm




(c) Value added




(d) Price of land




Note: Equilibria for different spreads of rates on savings. Endowments of manufacturing goods are uniform. Parameters used are $\theta = .89, \gamma = .54, \rho = .24, A = 1, \bar{a} = 0.035, \alpha = 0.01, L = 0.035, p_a = 1, p_x = 1, \bar{z} = 1, \delta_z = 0.25, w = ap_a, r_b = 1.1, r_s^r = 1.06$.

.

### 3.4.2 Behavior of individual agents

The competitive equilibrium results above are caused by the differentiated behavior of agents at different levels of wealth. To illustrate this process, we fix the price of land in this section, and analyze the decisions of individual agents. Figure 3.3 shows the optimal choice of land, intermediate inputs, and savings at different wealth endowments $I$ (where $I = \widehat{m} + \widehat{l}p_l$) given

fixed rates on borrowing and saving ($r_b, r_s(\widehat{m}) = r_s$). A land use of zero in the figure implies that the agent decides to be a worker in the manufacturing sector.

We can see *four different types* of behavior depending on the agent's wealth endowment: manufacturing work, farming with borrowing, self-financed farming, and farming with saving. At the lowest end of the wealth spectrum, an extremely poor agent is more averse to the risk of land investments and has less resources to invest in intermediate inputs for his land. This agent finds the benefits from selling his land, earning a manufacturing wage, and storing these earnings to be greater than the benefits from risky farming ($U(y^n(\widehat{m}, p_l) > E(U(y^f(x(\widehat{m}), l(\widehat{m}), z, \widehat{m}, p_l)))$. In addition, the decision to join the manufacturing workforce is affected by the return on his savings technology. This is because the storage technology affects the value of selling his land and going to work in the manufacturing sector. An agent that refuses to farm has $y^n(\widehat{m}, p_l) = r_s * (\widehat{m} + p_l * \widehat{l})$ resources available for consumption in the second period; therefore, a fall in the storage technology return available to him will make farming relatively more attractive. As we move from a developed to an underdeveloped financial system for the poor, there is a leftward shift in the first cutoff of Figure 3.3, for a previously indifferent agent will now find selling his land less worthy. As a result, more agents become farmers when lower returns on savings are available to the poor. It is this mechanism what drives the falling average farm size seen in Figure 3.2.

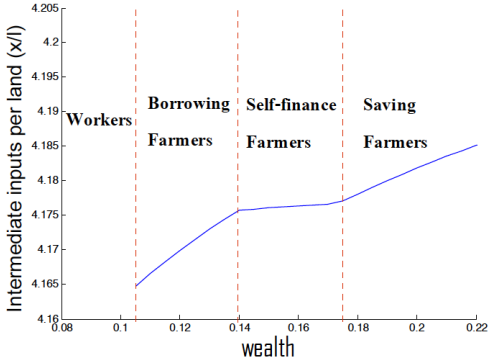
Moving upwards along the wealth spectrum, a slightly richer agent is able to tolerate the risk of investing in agriculture, but he needs to borrow in order to produce. Moreover, as we increase wealth, agents stop borrowing and the middle case of self-financed farming appears. An agent with a wealth level within this range will find it too expensive to borrow and can achieve more returns by investing resources in their own land instead of putting them into savings. As a result, he neither borrows nor saves but finances his production with his own endowment. Finally, the richest agents are wealthy enough to save part of their wealth in the storage technology, invest more in intermediate inputs, and are thus the ones who hold the most productive farms.

Figure 3.3: Farmer behavior by wealth level

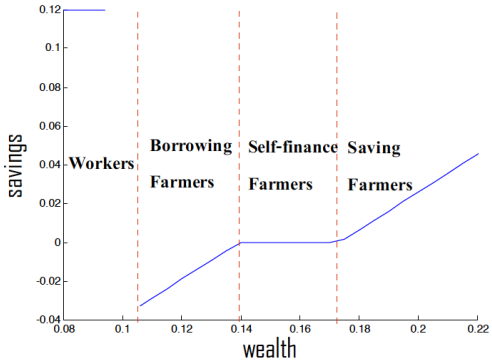
(a) Land demand




(b) Intermediate input usage per unit of land




(c) Savings per unit of land

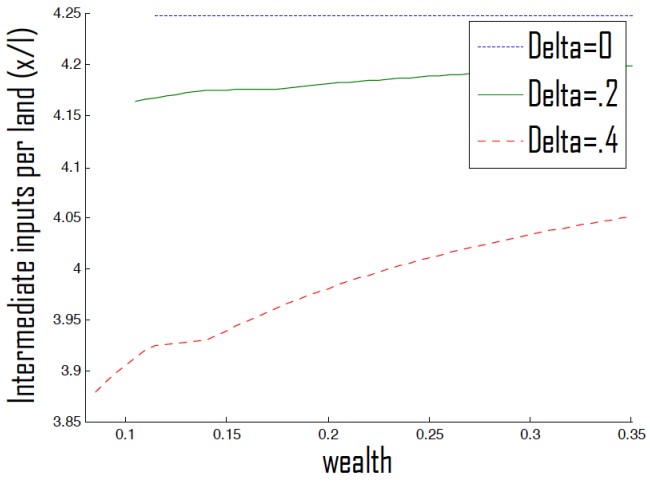



Note: Land demand (l), intermediate input use per unit of land (x/l), and savings (s) are presented. Parameters: $\theta = .89, \gamma = .54, \rho = .24, A = 1, \bar{a} = 0.035, \alpha = 0.01, L = 1,$ $p_a = 1$,$p_x = 1, \bar{z} = 1, \delta_z = 0.25$,$w = \bar{a}p_a, r_b = 1.1, r_s = 1$.

In addition, Figure 3.3 shows land and intermediate input use are also increasing in wealth within each farmer type. This reflects the decreasing RRA induced by the minimum subsistence requirement and can be inferred from the last term in equations (4), (5), and (6). In order to clarify the mechanics of how risk affects farmers' investment decisions at different levels of wealth, Figure 3.4 presents the optimal intermediate input to land ratio $(x/l)$ for farmers with different wealth levels, holding the price of land fixed. Each line represents the farmers' optimal decisions for different mean preserving spreads of the

distribution $h_z(z)$ (different values of $\delta_z$). Without risk ($\delta_z = 0$), the last term of equation 4 disappears, and intermediate input use is independent of wealth as depicted by the flat line in figure 3.4. When risk increases, this term becomes relevant and reduces the demand for intermediate inputs as seen in in the lowering of the curves. Moreover, because the minimum subsistence requirement induces decreasing RRA with respect to wealth, the differences in intermediate input use are much greater as we approach the minimum subsistence levels towards the left of the figure. This is precisely the key driver of productivity differences between small and large farms in the model within each farmer type. It is because of this that, within each type described above, intermediate input use –and productivity per farm– are increasing in wealth. This wealth-dependent effect of risk therefore explains the upward slopping curves within each group in Figure 3.3. It also reveals why as we go from a developed to an undeveloped financial system, as described in Figure 3.2, the falling average farm size decreases productivity. Poor farms are more sensitive to risk and therefore less productive; as a result, low returns on savings produce more poor, small, unproductive farms.

Figure 3.4: Intermediate input usage by risk level




Note: Delta refers to the spread of the support of z. Parameters used are $\theta = .89$, $\gamma = .54, \rho = .24, A = 1, \bar{a} = 0.035, \alpha = 0.01, L = 1, p_a = 1, p_x = 1, z = 1, w = \bar{a}p_a$, $r_b = 1.1, r_s = 1$.

## 3.5 Empirical evidence

This section presents suggestive cross-country evidence on three claims of the model above: 1) that financial underdevelopment is associated with the distribution of land; 2) that a larger average farm size increases the use of productivity-increasing intermediate investments, and 3) that high investments in intermediate goods are strongly correlated with higher aggregate productivity.

### 3.5.1 Data

Measures of aggregate agricultural production and farm size distribution compiled by Vollrath (2007) for 92 countries are used. The land data, originally gathered by Deninger and Squire (1998), comes from the decennial agricultural censuses of the Food and Agriculture Organization of the UN (FAO) which date from 1953 to 1993. Though more than one observation is available for 54 countries, the analysis below will only take the most recent data point for each country. The measure we use for land distribution is the average farm size, where a farm is an economic unit of agricultural production under single management, without regards to legal ownership. Output and labor force variables are from the FAOSTAT, as well as a measure of fertilizer use which serves as a proxy for total intermediate investments.

Following the literature, Peterson's land quality index (PLQI) constructed by Peterson (1987) is used to control for differences in the quality of land. This consists of the predicted value of agricultural land per acre divided by the average value of land per acre across all countries. The log of the value of land per acre is itself a weighted average of the percentage of non-irrigated crop land, the percentage of irrigated land, and the log of the long-run average annual precipitation.[22]

Financial development data comes from the most recent Database on Financial Development and Structure from the World Bank compiled by Beck et. al. and updated in 2013.

[22]These weights are derived using US agricultural census data. Therefore, the measure of land quality is the value of production from US farmland with the same underlying characteristics.

The relevant measure used is the interest rates spread between loans and deposits. We use this both because it is a widely used measure of financial development in the literature, and because it provides a proxy for the returns available to commercial agricultural savers in comparison to the returns available to richer agents in other sectors of the local economy and abroad.

Finally, for the instrumental variable strategy part of the estimation, legal origins data from LaPorta et al. (1999) are used. This data classifies legal systems as coming from British, French, Socialist, Scandinavian, or German traditions, and, despite problems with the validity of the exclusion restriction, it has been used in the literature as a source of exogenous variation of financial development.

### 3.5.2 Variation in land distribution, financial development and agricultural output

There is great variation in land distribution, productivity and financial development across countries. While the median farm size in Sub-Saharan Africa is around 2.3 Ha, it is 20 Ha in OECD countries, and the variance is significant within each of the regions. In addition, as mentioned in the introduction, productivity in agriculture varies much more than overall productivity measures across countries, which is consistent with the variance in agricultural production per land and worker (measured as output at international prices after input expenditures are deducted) and the use of fertilizer summarized in Table 3.1.[23] Furthermore, interest rate spreads vary significantly across the world, with better financial systems present in OECD countries and the worst in Latin America and Africa. The following sections exploit this variation and use this measure as a proxy for the savings returns available to the poor relative to the rich across countries.

[23]Land quality plays a significant role in this variation as countries with the lowest production have significant desert regions classified as agricultural land. However, significant differences remain after controlling for land quality.

Table 3.1: Summary Statistics

| Variable | Mean | Std. Dev. | Min. | Max. |
|---|---|---|---|---|
| Agric. output per Ha | 580.442 | 739.115 | 6.004 | 4847 |
| Agric. output per worker | 5682.13 | 9367.435 | 34.044 | 38345 |
| Average farm size (Ha) | 70.247 | 378.771 | 0.8 | 3601 |
| Fertilizer per Ha | 58.694 | 82.788 | 0* | 357 |
| Agric. population per Ha | 0.364 | 0.53 | 0.001 | 3.026 |
| Interest rate spread | 4.418 | 2.67 | 0.706 | 13.28 |
| Land quality (PLQI) | 102.7 | 45.555 | 27 | 249 |
| Agric. land irrigated (%) | 0.088 | 0.151 | 0** | 0.743 |
| Agric. land (thousands Ha) | 29910.963 | 76182 | 12 | 464481 |
| Income Gini | 40.716 | 9.461 | 23.94 | 60.84 |

* The lowest fertilizer use per Ha is Mali with 0.0004.

** The lowest percent of irrigated land is Botswana with 0.00004

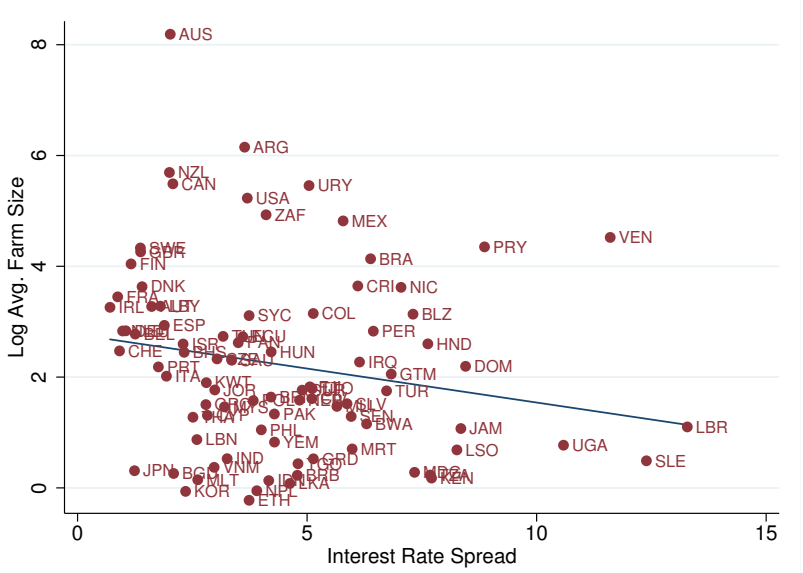
Note: Agricultural data comes from the most recent agricultural census available.

Interest rate spreads are taken as the long run average of observations for a given country.

### 3.5.3 Evidence on the impact of financial development on land misallocation

Figure 3.5 plots the level of financial development, as measured by the long run average of the interest rate spread in years before the agricultural census, to the log of average farm size. As expected, the predictive power of the financial measure is quite low, with population density playing a major role in countries like Australia and Argentina. Table 3.2 controls for this and, following the literature, uses the PLQI to control for land quality. A control for income inequality is also included. Once this is done, some patterns emerge. As suggested by the model, countries that are financially underdeveloped exhibit lower average farm sizes. Moreover, the magnitude of the results is significant, since a basis point increase in the interest spread is associated with a 10-20% decrease in average farm size.

Figure 3.5: Avg. farm size by interest rate spread




Interpreting the coefficient in the regression above as a causal effect, however, is problematic because of the possibility of omitted variable bias. It may be the case that there are unobserved covariates that are correlated with financial development and with inequity in the distribution of land. As an attempt to correct this, we follow some of the literature on the impact of financial development and use legal origins as an instrument. It is worth noting, however, that legal origins may affect property rights and other institutions through a variety of channels beyond financial structure.[24] Though we report the results of this specification below, this should be interpreted as suggestive, but non-conclusive, evidence of how financial development might affect farm size distribution. The results suggest that, though the coefficient loses significance, the negative relationship remains and is consistent with the connection between land allocation and financial development predicted by the model.

[24]See La Porta et al. (2008) for an overview of the multiple channels through which legal origins affect property rights, financial development and other economic outcomes, as well as potential problems in using them as instruments.

Table 3.2: Log avg. farm size (Ha)

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Interest rate spread | -0.123 | -0.102 | -0.108 | -0.204** | -0.321 |
| | (0.066) | (0.051) | (0.055) | (0.066) | (0.291) |
| Log labor per Ha | | -0.675*** | -0.697*** | -0.635*** | -0.582** |
| | | (0.089) | (0.121) | (0.125) | (0.185) |
| PLQI | | | 0.000540 | -0.000221 | -0.00134 |
| | | | (0.004) | (0.005) | (0.006) |
| Log irrigated land | | | | -0.639 | -1.022 |
| | | | | (1.357) | (1.712) |
| Income Gini | | | | 0.0410* | 0.0586 |
| | | | | (0.019) | (0.045) |
| Constant | 2.767*** | 1.376*** | 1.330 | 0.431 | 0.507 |
| | (0.340) | (0.322) | (0.744) | (0.966) | (1.071) |
| $N$ | 88 | 86 | 78 | 73 | 72 |
| $R^2$ | 0.039 | 0.432 | 0.455 | 0.497 | 0.474 |
| $F$ | 3.460 | 31.60 | 20.63 | 13.24 | 10.91 |
| Instruments | | | | L. Origins | L. Origins |

Standard errors in parentheses.

*$p < 0.05$, **$p < 0.01$, ***$p < 0.001$

### 3.5.4 Evidence on the impact of land distribution on intermediate investments

Figure 3.6 presents how average farm size by country is related to the use of intermediate inputs. A positive relationship is observed, though country income appears to explain much of the cross sectional variation. Table 3.3 presents regressions that attempt to control for both land quality and income. Income is included in the regression to control for the productivity of other sectors, as we are interested in the effect of land distribution on intermediate input use. It can be seen that income is not significant in this regression and average farm size is still positively associated with fertilizer use per hectare as predicted in the model. Hence, the positive correlation generated by the model between average farm size and intermediate input use is consistent with the pattern observed across countries.

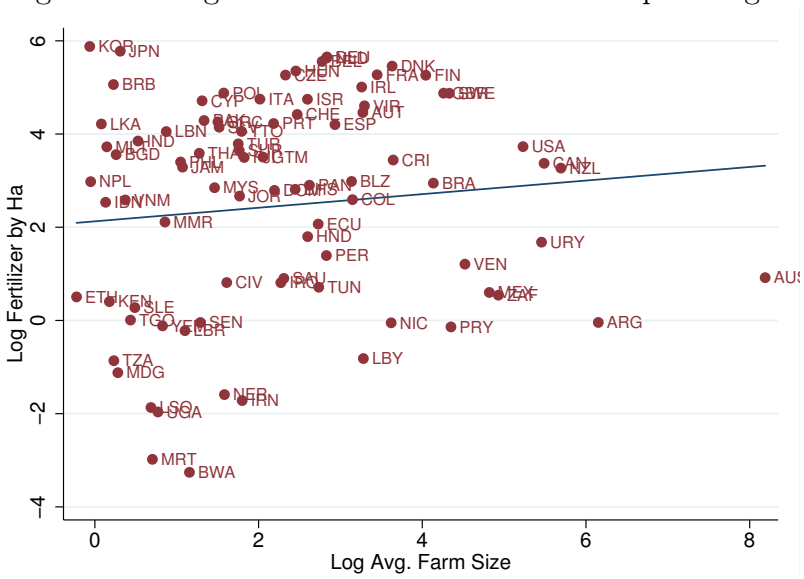
Table 3.3: Log fertilizer per Ha

| | (1) | (2) | (3) |
|---|---|---|---|
| Log avg. farm size | 0.146 | 0.463** | 0.411** |
| | (0.163) | (0.159) | (0.145) |
| PLQI | | 0.0312*** | 0.0245*** |
| | | (0.006) | (0.006) |
| Log irrigated land | | | 5.588** |
| | | | (1.973) |
| Log GDP | | | 5.05e-13 |
| | | | (0.000) |
| Constant | 2.127*** | -1.784* | -1.370 |
| | (0.456) | (0.854) | (0.836) |
| $N$ | 87 | 80 | 71 |
| $R^2$ | 0.009 | 0.271 | 0.399 |

Standard errors in parentheses.

$^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Figure 3.6: Avg. farm size and intermediate input usage

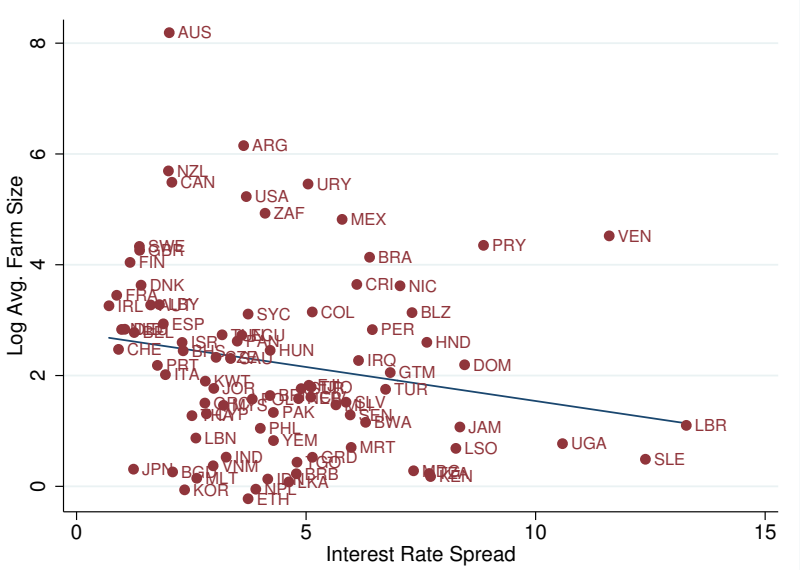



### 3.5.5 Evidence on the impact of intermediate investments on aggregate productivity and the productivity gap

Though the distribution of land and financial deficiencies may affect agricultural productivity through other channels, intermediate use seems to be the most prominent one. In fact, by

taking one measure of intermediate input use – namely fertilizer – we can explain a significant share of the variance in agricultural productivity across countries ($R^2 = .76$). Table 3.4 shows a strong correlation between fertilizer use and agricultural production per unit of land, which has been highlighted in the literature before. Moreover, when average farm size is included in this regression, the sign of the coefficient is negative and insignificant, suggesting that the main vehicle through which land distribution affects productivity is by influencing the use of intermediate input use, as modeled in the theoretical section.

Figure 3.7: Productivity and intermediate input usage




## 3.6 Conclusion

When the poor lack access to a reliable savings asset, they have an incentive to hold their land – even when not willing or able to make it productive. Poor savings options produce poor farmers – probably in large numbers – and may partially explain why we observe small average farm sizes in the developing world even when property rights are improved and old colonial structures removed. Furthermore, the link between financial development and the

Table 3.4: Log agricultural production per Ha

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Log fertilizer per Ha | 0.508*** | 0.542*** | 0.504*** | 0.513*** |
| | (0.031) | (0.033) | (0.033) | (0.036) |
| Total agric. land | | -0.00000276** | -0.00000240** | -0.00000211* |
| | | (0.000) | (0.000) | (0.000) |
| PLQI | | | 0.00489** | 0.00438* |
| | | | (0.002) | (0.002) |
| Log avg. farm size | | | | -0.0303 |
| | | | | (0.053) |
| Constant | 4.304*** | 4.294*** | 3.891*** | 3.979*** |
| | (0.110) | (0.125) | (0.183) | (0.240) |
| $N$ | 87 | 78 | 73 | 73 |
| $R^2$ | 0.756 | 0.793 | 0.840 | 0.841 |

Standard errors in parentheses

$^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

distribution of land adds a piece to the puzzle that is explaining the dramatic variation in agricultural productivity across the world.

This paper departs from the literature on how to model heterogeneity across farmers, which regularly takes the form of variation in labor productivity, skill, or some idiosyncratic parameter in the preferences. This feature is often used to explain the connection between land distribution and aggregate productivity. In contrast, the model presented suggests that even when individuals have identical skills and are unconstrained to borrow, there can be dispersion in farm sizes with a direct impact on agricultural productivity. In particular, heterogeneity in access to financial services may camouflage in the data as differences in skill or preferences. Given micro evidence from developing countries that shows farmers are willing to invest in negative-return assets to store wealth, a financial story may be an important factor. The implications of this change in interpretation are relevant, since the latter one suggests that better access to savings services can improve productivity and accelerate the transition of labor out of unproductive agriculture.

The cross-country evidence shown here is consistent with the model's predictions; nonetheless, it remains far from being conclusive. Micro evidence on the saving and invest-

ing patterns of farmers at different wealth levels would be useful in better understanding the connection between financial access and agricultural investments. At the macro level, time variation in a country's land distribution and its relationship to financial development can also help, particularly in places that suffered a transformation in financial accessibility for the poor. Lastly, the question remains on the relative importance of this mechanism in explaining overall variation in agricultural productivity. To answer this, a calibration exercise that considers different mechanisms in the literature is needed. This is left for future work.

# Bibliography

Abowd, John M, Francis Kramarz and David N Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67(2):251–333.

Abowd, John M., Robert H. Creecy and Francis Kramarz. 2002. "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data.".

Akerman, Anders, Elhanan Helpman, Oleg Itskhoki, Marc-Andreas Muendler and Stephen Redding. 2013. "Sources of Wage Inequality." *American Economic Review* 103(3):214–219.

Alvarez, Jorge, Niklas Engbom and Chris Moser. 2016. "Firms and the Great Decline in Labor Earnings Inequality in Brazil.".

Anagol, Santosh, Alvin Etang and Dean Karlan. 2013. "Continued Existence of Cows Disproves Central Tenets of Capitalism?" *National Bureau of Economic Research* (w19437).

Anderson, K. and S. Nelgen. 2013. *Updated Global Estimates of Distortions to Agricultural Incentives, 1955 to 2011.* World Bank.

Atkinson, Anthony B. and François Bourguignon, eds. 2015. *Handbook of Income Distribution.* Vol. 2 ed. Elsevier.

Autor, David H., Lawrence F. Katz and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90(2):300–323.

Bagger, Jesper, Bent Jesper and Royal Holloway. 2014. "Productivity and Wage Dispersion: Heterogeneity or Misallocation.".

Bagger, Jesper, Fabien Postel-Vinay and Jean-marc Robin. 2014. "Tenure , Experience , Human Capital and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *American Economic Review* pp. 1–54.

Bagger, Jesper, Kenneth L Sorensen and Rune Vejlin. 2013. "Wage sorting trends." *Economics Letters* 118(1):63–67.

Barros, Ricardo, Mrela de Carvalho, Samuel Franco and Rosane Mendonça. 2010. "Markets, the State and the Dynamics of Inequality: The Case of Brazil." *UNDP Discussion Paper* .

Bartelsman, Eric, John Haltiwanger and Stefano Scarpetta. 2009. Measuring and Analyzing Cross-country Differences in Firm Dynamics. In *NBER Chapters*. National Bureau of Economic Research, Inc pp. 15–76.

Bartelsman, Eric, John Haltiwanger and Stefano Scarpetta. 2013. "Cross-Country Differences in Productivity: The Role of Allocation and Selection." *American Economic Review* 103(1):305–334.

Barth, Erling, Alex Bryson, James C Davis and Richard Freeman. 2014. "It's Where You Work: Increases in Earnings Dispersion Across Establishments and Individuals in The US." *NBER working paper* (w20447).

Beck, Cihak, Demirguc-Kunt, Freyen and Levine. N.d. "Financial Development and Structure Dataset. Updated April 2013. World Bank.".

Bernard, Andrew B. and J. Bradford Jensen. 1995. "Exporters, Jobs, and Wages in U.S. Manufacturing: 1976-1987." *Brookings Papers: Microeconomics* .

Bliss, Christopher and Nicholas Stern. 1978. "Productivity, wages and nutrition." *Journal of Development Economics* (4):331–362.

Bloom, Nicholas, Fatih Guvenen, David Price and Jae Song. 2015. "Evolution of Inequality within Firms (SED Abstract).".

Bonhomme, Stephane and Laura Hospido. 2013. "The Cycle of Earnings Inequality: Evidence from Spanish Social Security Data." (June 2009).

Bonhomme, Stéphane, Thibaut Lamadon and Elena Manresa. 2015. "A Distributional Framework for Matched Employer Employee Data.".

Buhai, Sebastian and Miguel Portela. N.d. "Cambridge Working Papers in Economics Forthcoming in Econometrica." . Forthcoming.

Bustos, Paula and Caprettini Bruno. 2015. "Agricultural Productivity and Structural Transformation: Evidence from Brazil." *Working Paper* .

Cai, Wenbiao. 2012. "Skill accumulation and international productivity differences across sectors." *Working Paper* .

Card, D., F. Devicienti and a. Maida. 2013. "Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data." *The Review of Economic Studies* 81(1):84–111.

Card, David. N.d. "The Return of the Firm to Labor Economics David Card UC Berkeley." . Forthcoming.

Card, David, Ana Rute Cardoso and Patrick Kline. 2015. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women.".

Card, David, Bernd Fitzenberger, Thomas Lemieux and Alexandra Spitz-oener. 2009. "Revisiting the german wage structure." *Quarterly Journal of Economics* (May).

Card, David, Joerg Heining and Patrick Kline. 2013. "Workplace heterogeneity and the rise of West German wage inequality." *Quarterly Journal of Economics* 128:967–1015.

Card, David and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20(4):733–783.

Carvalho Filho, Irineu E. and Marcello M. Estevão. 2012. "Institutions, Informality, and Wage Flexibility: Evidence From Brazil." *IMF Working Papers* 12(84):1.

Caselli, Francesco. 2005. "Accounting for cross-country income differences." *Handbook of Economic Growth* 1:679–741.

Cavalcanti, Tiago. 2014. "Misallocation in the brazilian manufacturing sector.".

Chauvet, Marcelle. 2011. "Dating Business Cycles in Emerging Market Economies: The Case of Brazil." (June).

Córdova, José Ernesto López and Mauricio Mesquita Moreira. 2003. "Regional Integration and Productivity: The Experiences of Brazil and Mexico." *IDB Publications (Working Papers)* .

de Araujo, Ana Luisa Pessoa. 2014. "Wage Inequality and Job Stability." pp. 1–53.

de Vries, Gaaitzen J., Abdul A. Erumban, Marcel P. Timmer, Ilya Voskoboynikov and Harry X. Wu. 2012. "Deconstructing the BRICs: Structural transformation and aggregate productivity growth." *Journal of Comparative Economics* 40(2):211–227.

Deaton, Angus. 1997. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy.*

Demirgüç-Kunt, Asli and Leora Klapper. 2012. "Measuring financial inclusion: The global findex database." *World Bank Policy Research Working Paper* 6025.

Deninger, Klaus and Lyn Squire. 1998. "New ways of looking at old issues: inequality and growth." *Journal of Development Economics* 57.2:259–287.

Dix-Carneiro, Rafael. 2014. "Trade Liberalization and Labor Market Dynamics." *Econometrica* 82(3):825–885.

Dix-Carneiro, Rafael and Brian K. Kovak. 2015. "Trade Reform and Regional Dynamics: Evidence From 25 Years of Brazilian Matched Employer-Employee.".

Donovan, Kevin. 2011. "Risk , Farm Ownership , and International Productivity Differences." *Meeting Papers. Society for Economic Dynamics.* (1088):1–18.

Donovan, Kevin. 2012. "Agricultural Risk , Intermediate Inputs and Cross-Country Productivity Differences." *Working Paper* (October).

Dunne, Timothy, Lucia Foster, John Haltiwanger and Kenneth R. Troske. 2004. "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment." *Journal of Labor Economics* 22(2):397–430.

Dupas, P., S. Green, A. Keats and J. Robinson. 2016. "Challenges in banking the rural poor: Evidence from Kenya's western province." *African Successes: Modernization and Development* 3.

Engbom, Niklas and Christian Moser. 2015. "Earnings Inequality and the Minimum Wage: Evidence from Brazil.".

Faggio, G., K. G. Salvanes and J. Van Reenen. 2010. "The evolution of inequality in productivity and wages: panel data evidence." *Industrial and Corporate Change* 19(6):1919–1951.

Ferreira, Francisco H. G., Phillippe G. Leite and Matthew Wai-Poi. 2007. *Trade Liberalization, Employment Flows, And Wage Inequality In Brazil.* Policy Research Working Papers The World Bank.

Ferreira, Pedro Cavalcanti and Jose Luiz Rossi. 2003. "New Evidence from Brazil on Trade Liberalization and Productivity Growth." *International Economic Review* 44(4):1383–1405.

Friedrich, Benjamin and Luigi Pistaferri. 2014. "Earnings Dynamics and Firm-Level Shocks." pp. 1–45.

Gibbons, Robert and Lawrence Katz. 1992. "Does Unmeasured Ability Explain Inter-industry Wage Differentials?" *The Review of Economic Studies* 59(3):515–535.

Gollin, Douglas, David Lagakos and Michael Waugh. 2014. "The Agricultural Productivity Gap in Developing Countries." *The Quarterly Journal of Economics* 129(2):939–993.

Gollin, Douglas, Stephen L Parente and Richard Rogerson. 2004*a*. "Farm Work, Home Work and International Productivity Differences." *Review of Economic Dynamics* 7(4):827–850.

Gollin, Douglas, Stephen L Parente and Richard Rogerson. 2007. "The Food Problem and the Evolution of International Income Levels." *Journal of Monetary Economics* (54.4):1230–1255.

Gollin, Douglas, Stephen Parente and Richard Rogerson. 2002. "The Role of Agriculture in Development." *American Economic Review* pp. 160–164.

Gollin, Douglas, Stephen Parente and Richard Rogerson. 2004*b*. "Farm Work, Home Work and International Productivity Differences." *Review of Economic Dynamics* 7.4:827–850.

Greenwood, Jeremy, Juan M. Sanchez and Cheng Wang. 2013. "Quantifying the impact of financial development on economic development." *Review of Economic Dynamics* 16(1):194–215.

Gruetter, Max and Rafael Lalive. 2009. "The importance of firms in wage determination." *Labour Economics* 16(2):149–160.

Guvenen, Fatih, Serdar Ozkan and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy* 122(3):621–660.

Hall, Robert and Charles Jones. 1999. "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" *The Quarterly Journal of Economics* 114(February):83–116.

Heathcote, Jonathan, Fabrizio Perri and Giovanni L. Violante. 2010. "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967-2006." *Review of Economic Dynamics* 13(1):15–51.

Helpman, Elhanan, Oleg Itskhoki, Marc-Andreas Muendler and Stephen Redding. 2016. "Trade and Inequality: From Theory to Estimation." *Forthcoming at the Review of Economic Studies (Forthcoming)* .

Herrendorf, Berthold and Arilton Teixeira. 2011. "Barriers to Entry and Development." *International Economic Review* 52(2):573–602.

Herrendorf, Berthold and Todd Schoellman. 2011. "Why is Measured Productivity so Low in Agriculture?" *Working Paper* .

Herrendorf, Berthold and Todd Schoellman. 2014. "Why is Measured Productivity So Low in Agriculture?" *Review of Economic Dynamics* .

Just, Richard E. and Rulon D. Pope. 1979. "Production function estimation and related risk considerations." *American Journal of Agricultural Economics* 61.2:276–284.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto and Christopher Udry. 2014. "Agricultural Decisions after Relaxing Credit and Risk Constraints." *Quarterly Journal of Economics* 129(2):597–652.

Kim, Yong. 2014. "The Price of Experience.".

Klenow, P J and a Rodriguez-Clare. 1997. *The Neoclassical Revival in Growth Economics: Has It Gone Too Far?* Vol. 12 of *NBER Chapters* MIT Press.

Kongsamut, P., S. Rebelo and D. Xie. 2001. "Beyond Balanced Growth." *Review of Economic Studies* 68(4):869–882.

Kopczuk, Wojciech, Emmanuel Saez and Song Jae. 2010. "Earnings inequality and mobility in the United States: evidence from social security data since 1937." *Quarterly Journal of Economics* 125(1):91–128.

Krueger, Alan B and Lawrence H Summers. 1988. "Efficiency Wages and The Inter-Industry Wage Structure." *Econometrica: Journal of the Econometric Society* pp. 259–293.

Kuznets, Simon. 1971. *Economic Growth of Nations: Total Output and Production Structure*. Cambridge: Harvard University Press.

La Porta, Rafael, Florencio Lopez-de Silanes and Andrei Shleifer. 1997. "Legal determinants of external finance." *The Journal of Finance* 52(3):1131–1150.

La Porta, Rafael, Florencio Lopez-de Silanes and Andrei Shleifer. 2008. "The Economic Consequences of Legal Origins." *Journal of Economic Literature* 46(2):285–332.

Lagakos, David and Michael E Waugh. 2013. "Selection, Agriculture and Cross-country Productivity differences." *American Economic Review* 103:948–980.

Laitner, J. 2000. "Structural Change and Economic Growth." *Review of Economic Studies* 67(3):546–561.

Lewis, W Arthur. 1955. *Theory of Economic Growth*. Routledge.

Lopes De Melo, Rafael. 2013. "Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence." pp. 1–39.

Lopez, J. Humberto and Guillermo Perry. 2008. "Inequality in Latin America: Determinants and Consequences.".

Medeiros, Marcelo, Pedro H G Ferreira De Souza and Fábio Avila de Castro. 2014. "O topo da distribuição de renda no Brasil: primeiras estimativas com dados tributários e comparação com pesquisas domiciliares , 2006 - 2012." pp. 2006–2012.

Menezes-Filho, Naércio Aquino, Marc-Andreas Muendler and Garey Ramey. 2008. "The Structure of Worker Compensation in Brazil, With a Comparison to France and the United States." *The Review of Economics and Statistics* 90(2):324–346.

Michelacci, Claudio and Vincenzo Quadrini. 2009. "Financial Markets and Wages." *Review of Economic Studies* 76(2):795–827.

Milanovic, Branko L. 2014. "Database: All the Ginis, 1950-2012.".

Moreira, Mauricio Mesquita. 2004. "Brazil's Trade Liberalization and Growth: Has it Failed?" *International Trade* .

Mueller, Holger M and Paige P Ouimet. 2015. "Wage Inequality and Firm Growth.".

Muendler, Marc-Andreas. 2004. "Trade, Technology, and Productivity: A Study of Brazilian Manufacturers, 1986-1998." *CESifo Working Paper Series* .

Murphy, Kevin M and Robert H Topel. 1987. "Unemployment, Risk, and Earnings: Testing for Equalizing Wage Differences in The Labor Market." pp. 103–140.

Murphy, Kevin M and Robert H Topel. 1990. "Efficiency Wages Reconsidered: Theory and Evidence." pp. 204–240.

Oaxaca, Ronald. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review* 14(3):693–709.

Ogaki, Massao and Qiang Zhang. 2001. "Decreasing relative risk aversion and tests of risk sharing." *Econometrica* 69(2):515–526.

Pavcnik, Nina, Andreas Blom, Pinelopi Goldberg and Norbert Schady. 2004. "Trade Liberalization and Industry Wage Structure: Evidence from Brazil." *World Bank Economic Review* 18(3):319–344.

Peterson, Willis L. 1987. "International Land Quality Indexes." *Staff Papers 13877, University of Minnesota, Department of Applied Economics* .

Restuccia, Diego. 2011. "Recent Developments in Economic Growth." *FRB Richmond Economic Quarterly* 97(3):329–357.

Restuccia, Diego, Dennis Tao Yang and Xiaodong Zhu. 2008*a*. "Agriculture and Aggregate Productivity: A Quantitative Cross-country Analysis." *Journal of Monetary Economics* 55(2):234–250.

Restuccia, Diego, Dennis Tao Yang and Xiaodong Zhu. 2008*b*. "Agriculture and aggregate productivity: A quantitative cross-country analysis." *Journal of Monetary Economics* 55(2):234–250.

Restuccia, Diego and Richard Rogerson. 2008. "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments." *Review of Economic Dynamics* 11(4):707–720.

Restuccia, Diego and Tasso Adamopoulos. 2014. "The Size Distribution of Farms and International Productivity Differences." *The American Economic Review* 104(6):1667–1697.

Rosenzweig, Mark R. and Kenneth I. Wolpin. 1993. "Credit market consraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investment in bullocks in India." *Journal of Political Economy* 101(2):223–244.

Rostow, Walt Whitman. 1960. *The Stages of Economic Growth: A Non-communist Manifesto*. Cambridge University Press.

Roy, Andrew Donald. 1951. "Some Thoughts on The Distribution of Earnings." *Oxford Economic Papers* 3(2):135–146.

Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom and Till von Wachter. 2015. "Firming Up Inequality." *NBER Working Paper 21199* .

Stiglitz, Joseph E. 1976. "The efficiency wage hypthesis, surplus labor, and the distribution of income in L.D.C.s." *Oxford Economic Papers* 28:185–207.

Stoyanov, Andrey and Nikolay Zubanov. 2012. "Productivity Spillovers Across Firms through Worker Mobility." 4(April):168–198.

Traxler, Greg, Jose Falck-Zepeda and Ken Sayre. 1995. "Production risk and the evolution of varietal technology." *American Journal of Agricultural Economics* 77(1):1–7.

Tsounta, Evridiki and Anayochukwu I. Osueke. 2014. "What is behind Latin America's Declining Income Inequality?".

Udry, Christopher. 2011. "Land Tenure." *The Oxford Companion to the Economics of Africa* .

Ulyssea, Gabriel. 2014. "Firms, Informality and Development: Theory and evidence from Brazil.".

Vollrath, Dietrich. 2009. "How important are dual economy effects for aggregate productivity?" *Journal of Development Economics* 88(2):325–334.

Vollrath, Dietrich. 2012. "Land Tenure, Population, and Long-run Growth." *Journal of Population Economics* 25(3):883–852.

Vollrath, Dietrich. 2014. "The Efficiency of Human Capital Allocations in Developing Countries." *Journal of Development Economics* 108:106–118.

World Bank. 2015. "World Development Indicators.".
**URL:** *http://data.worldbank.org/data-catalog/world-development-indicators*

Young, Alwyn. 2013. "Inequality, the Urban-Rural Gap and Migration." *The Quarterly Journal of Economics* p. qjt025.

Zerfu, Daniel and Donald Larson. 2010. "Incomplete markets and fertilizer use: evidence from Ethipia." *World Bank* .