

KAUFFMAN DISSERTATION EXECUTIVE SUMMARY

KAUFFMAN ENTREPRENEURSHIP SCHOLARS

KAUFFMAN DISSERTATION FELLOWSHIP

Part of the Ewing Marion Kauffman Foundation's Entrepreneurship Scholars initiative, the Kauffman Dissertation Fellowship recognizes exceptional doctoral students and their universities. The annual program awards Dissertation Fellowship grants to Ph.D., D.B.A., or other doctoral students at accredited U.S. universities to support dissertations in the area of entrepreneurship.

Since its establishment in 2003, this program has helped to launch world-class scholars into the exciting and emerging field of entrepreneurship research, thus laying a foundation for future scientific advancement. The findings generated by this effort will be translated into knowledge with immediate application for policymakers, educators, service providers, and entrepreneurs as well as high-quality academic research.

Ewing Marion
KAUFFMAN
Foundation

www.kauffman.org/kdf

ESSAYS IN THE ECONOMICS OF INEQUALITY

CHRISTIAN ALEXANDER MOSER

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

ADVISERS: MIKHAIL GOLOSOV AND NOBUHIRO KIYOTAKI

SEPTEMBER 2016

Dissertation Abstract

What are the drivers of recent trends in earnings inequality? What role have economic policies played in this evolution? How should optimal policy relating to inequality and redistribution be designed? Finding answers to these questions is central to improving economic welfare in a world that is experiencing rapid shifts in the distribution of income. This collection of essays contributes to our understanding of the determinants of earnings inequality and the design of optimal redistributive policies.

Executive Summary

The central theme of this doctoral dissertation is to dissect the sources of earnings inequality, with a special focus on firms. To this end, four dissertation chapters focus on different aspects of earnings inequality in different contexts.

Chapter 1

Chapter 1, co-authored with Jorge Alvarez and Niklas Engbom, uses administrative matched employer-employee data to decompose a decline in earnings inequality in Brazil from 1996 to 2012.

Since the mid-1990s, Brazil has experienced a large reduction in earnings inequality resembling the experience of other Latin American economies during this period. This decline in earnings inequality stands in stark contrast to that of the U.S. and many developed countries, which saw inequality steadily increasing over the past two decades. This chapter studies the sources of this decline.

To this end, we exploit a large administrative linked employer-employee dataset containing information on hundreds of millions of job spells between 1988 and 2012. By repeatedly estimating an additive worker and firm fixed effects model due to Abowd et al. (1999)—henceforth AKM—we are able to separately identify the contributions of firm and worker-specific factors towards changes in Brazilian inequality. In a second stage, we link this earnings decomposition to confidential data on a rich set of worker demographics and firm financial data on Brazilian manufacturing and mining enterprises. The merged data allow us to study the link between worker characteristics, firm characteristics, and pay in a large developing economy. In doing so, we distinguish between changes in the distribution of worker and firm characteristics on the one hand and changes in the returns to these characteristics in pay.

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 attributable to a decline in the covariance between worker and firm fixed effects and the residual. Given that worker heterogeneity is the most important component in explaining pay levels throughout this period, the compression in firm-specific pay contributed more than proportionately towards Brazil's inequality decline.

Second, changes in the link between firm performance and pay account for a significant fraction of the compression in the firm component of workers' earnings. We first show that a substantial share of the cross-sectional variation in the firm component of pay is explained by differences in observable firm characteristics, with more productive and larger 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 firm characteristics to pay, while none is due to firms becoming more similar in observable characteristics over time. Altogether, a weaker link between observable firm characteristics and worker pay explains 30 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 and education explain over 35 percent of the variance of the worker component of pay. However, we do not observe a large compression in the underlying distributions of such characteristics over time. Instead, the decline in worker pay heterogeneity is driven by a rapid fall in the returns to observable measures of worker ability, particularly the education premia. Lower returns to worker age and education explain 14 percent of the overall fall in the variance of log earnings over the 1996-2012 period.

This decomposition of the sources of Brazil's earnings inequality decline informs our understanding of various commonly proposed explanations of the decline. On the worker side, our results do not support an often articulated view that changes in educational attainment accounted for a significant share of Brazil's inequality evolution over the period. While educational attainment increased rapidly over this period, those gains were approximately offset by concurrent declines in the high school and college education premia. 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 other factors affecting the productivity evolution during this period could have been an important driver behind changes to the earnings distribution. Yet, in line with U.S. trends, we find that the Brazilian productivity distribution actually grew more dispersed over this period.

Chapter 2

Chapter 2, co-authored with Niklas Engbom, quantifies the contribution of a rise in the minimum wage to Brazil's inequality decline in an equilibrium search model with heterogeneity in worker ability and firm productivity.

To what extent does economic policy shape earnings inequality? Given that inequality has increased significantly in many economies over the past decades, a quantitative answer to this question has become more urgent. Aiming to boost earnings at the bottom of the distribution, many countries have advocated a minimum wage. While this may come at a cost, including increased unemployment, proponents of the minimum wage defend the policy as an effective way to reduce labor income inequality. Despite its importance, given an ongoing debate in the empirical literature and limited theoretical guidance on how to reconcile different findings, the quantitative effect of a minimum wage on earnings inequality is far from clear.

Skeptics of the benefits associated with a minimum wage point to the small fraction of workers bound by the wage floor as evidence that its impact is likely to be limited. Furthermore, spillover effects of the minimum wage (i.e. effects higher up in the earnings distribution), while potentially promising, have been hard to identify given data limitations and methodological disagreements. Previous work in this area primarily builds on reduced-form evidence from household survey data as in Lee (1999) and Autor et al. (2016), with the latter concluding that spillover effects are indistinguishable from measurement error. Complementing this literature, we use large administrative data combined with a structural and testable model to quantify the effects of a minimum wage throughout the earnings distribution. The size and nature of the administrative data allow us to exploit more detailed variation with higher estimation precision than has previously been possible. The model lets us quantify the causal equilibrium effects of the minimum wage on inequality and unemployment in a counterfactual policy experiment, enabling us to discuss welfare implications.

To address this problem, we study the case of Brazil between 1988 and 2012, which has two key advantages for our purpose. First, Brazil has exceptional data availability, with administrative matched employer-employee data that we merge with administrative firm financial data covering a long time horizon. Second, there was a large policy change implemented in Brazil, with the real minimum wage increasing by 119 percent in real terms, starting out at 30 percent of median earnings and reaching 60 percent by the end of the period. The combination of a large policy change and detailed microdata provide us with an ideal testing ground for quantifying the effects of the minimum wage on earnings inequality.

To this end, we carry out the following three steps. In the first step, we use matched employer-employee data on workers and firms in Brazil to document a 26 log points drop in the variance of earnings between 1996 and 2012. We show that this decrease is characterized by three key facts: (i) the decrease was bottom-driven yet pervasive throughout large parts of the earnings distribution; (ii) reductions in the firm productivity-pay premium and in the worker skill premium were the key drivers behind the decrease; and (iii) the bindingness of the minimum wage is correlated across Brazilian states and over time with compression up to the 75th percentile of the earnings distribution.

In the second step, we build an equilibrium model of frictional wage dispersion based on the canonical framework by Burdett and Mortensen (1998). Motivated by our empirical findings, we extend this framework to tractably feature heterogeneity in worker ability in addition to firm productivity differentials described in the original paper. We close the model by introducing a vacancy margin, allowing job creation to respond to the minimum wage increase. Theoretically and in line with our empirical facts, we show that minimum wage effects ripple through the earnings distribution and cause a decrease in the pass-through from firm productivity and worker skill to pay. Spillover effects arise because firms compete for workers by

setting wages strategically relative to one another and in reference to the minimum wage. Therefore, the effects of the minimum wage reach above the wage floor but slowly fade toward the top of the earnings distribution by reducing the productivity-pay gradient across firms and the skill premium across workers.

In the third step, we use our model to quantify the causal effect of the minimum wage on earnings inequality. We estimate the model on Brazilian matched employer-employee data from 1996-2000 and use it to conduct a counterfactual policy experiment, simulating the equilibrium effects of the observed minimum wage increase. In line with our three stylized facts characterizing Brazil's inequality decrease, the estimated model predicts that the rise in the minimum wage caused: (i) a 21 log points decrease in the P50-P10 earnings ratio, or 68 percent of its empirical counterpart, but only a six log points decrease in the P90-P50 earnings ratio, or 46 percent of the empirical change; (ii) essentially all of the explained decline as a consequence of a lower firm productivity-pay gradient and lower worker skill premium; and (iii) significant spillover effects reaching up to the 75th percentile of the earnings distribution. Due to large effects of the minimum wage higher up the earnings distribution, the model attributes 70 percent of the total decrease in the variance of log earnings observed in the data over this period to the rise in the minimum wage. Half of the total inequality decrease in the data and in the model are due spillover effects of the minimum wage reaching up in the earnings distribution.

In contrast to a competitive theory of labor markets, our model also predicts modest disemployment effects of the minimum wage, consistent with the data. Thus, a general insight from our analysis is that frictional labor markets can propagate effects of policies like the minimum wage on the inequality while also buffering negative employment effects.

Chapter 3

Chapter 3 examines the interaction between financial frictions, barriers to entry, and firms' extensive margin decisions in a dynamic occupational choice model.

There is significant variation in per-capita income across countries and the largest share of it can be attributed to differences in total factor productivity (TFP). Differences in TFP may stand in for many factors, including differences in the economic environment affecting business operations. For instance, imperfect credit markets may impede entrepreneurship, reducing both aggregate productivity and output. On the one hand, a lack of financing may lead active entrepreneurs to operate below the optimal scale of production or, when individuals foresee this inefficiency, even prevent them from starting a business in the first place. On the other hand, less productive but wealthy entrepreneurs may continue to operate in less than competitive markets that arise with imperfect access to credit. If cost of creating a new business is large, this misallocation of resources may be further amplified by increasing entrepreneurs' dependence on external financing.

But the story does not end at the individual's decision to enter the market or not. Once active, entrepreneurs may be exposed to business risks and find themselves relatively unproductive for an uncertain amount of time. At that point, individuals face a choice between exercising a relatively attractive outside option, such as working in the labor market, or staying in business in hopes of a quick recovery. The option value of remaining active will be greater the higher the entrepreneur's expectation of future productivity and the higher the barriers to future re-entry. With perfect credit markets, entrepreneurs weigh current losses associated with running a relatively unprofitable business against the expectation of future profits and pick the more attractive option. When access to credit is limited, however, entrepreneurs may decide to liquidate their business prematurely, leading to an inefficiently high rate of firm exits.

In this paper, I first assesses the size of distortions due to entry costs combined with financial frictions and then considers one particular type of industrial policy targeted at mitigating such distortions. I then proceed to calibrate my model to assess the impact of corporate bailouts, or soft budget constraints, on a dynamic economy with occupational choice in the presence of entry costs and credit constraints. The focus of the investigation is the role of distortions to existing firms' extensive margins and the effects of government intervention in the form of subsidies to least productive active entrepreneurs. Thus, the analysis is primarily concerned with the premature exit of existing firms rather than barriers to entry of new entrepreneurs. I will argue that distortions due to premature firm exits can be sizable and that the effects of government intervention depend heavily on the range of underlying parameters.

The spirit of my investigation is to set up the economic environment in favor of finding a positive effect of soft budget constraints. I then look at remaining sources of inefficiencies due to government intervention and assess which of these are the most significant. I find that the persistence of productivity shocks is a key determinant of the efficiency of an economy with credit frictions and entry costs in general, and of soft budget constraints in particular. My main result is that distortions due to government intervention are largest when the stochastic process guiding firm productivities is persistent.

Credit markets and entry costs are two natural factors affecting firm dynamics. On the one hand, the better the access to credit the more businesses per capita are started on average. Intuitively, running a business is more attractive when operating at the optimal scale and as such better credit markets promote firm entry. On the other hand, higher credit market quality is significantly and negatively related to firm exit rates, which suggests that financing needs play an important role in shaping firm exit decisions. Firms that have undertaken large investments in the past, potentially related to industry entry, consider this cost as sunk when making their current production decisions but also take into consideration that this investment may be partially lost in case of firm liquidation. Firms with the greatest

fixed cost investments that find themselves with easy access to credit markets should, all else equal, be the most reluctant to liquidate their business. A problem arises when firms undergo temporary hardship and simultaneously run into financing constraints. In such a case, information about the prospects of companies may be imperfect and people's perception about the future profitability of the enterprise may be at odds. For example, while General Motors may claim to have prospects of a quick recovery and profitable future business operation, public and private lenders may wonder if their cash injections will help their creditor sustain business in the long-term or whether increased lending just keeps the "zombie" firm alive for a little longer before the next big collapse. Besides informational asymmetries, in such circumstances the government faces what Kornai (1980) termed the soft budget constraint problem: Ex-post aid may be efficient but to avoid adverse selection problems the government would like to commit ex-ante to not intervening. Such a commitment, of course, is non-credible and as firms foresee this they will make their extensive and intensive margin decisions accordingly. This industrial policy problem, among others, motivates my analysis of the interaction between financial frictions and entry-costs as well as potential government intervention.

Chapter 4

Chapter 4, co-authored with Pedro Olea de Souza e Silva, analyzes retirement savings policies in a model of optimal taxation with unobservable differences in earnings ability and heterogeneity in time preferences.

With a budget of more than \$700 billion dollars in 2014, the social security survivors and old-age benefits program is the largest U.S. federal government social policy. The public finance literature has suggested this program is a paternalistic policy that aims at helping individuals save for retirement. The large body of evidence for present bias and heterogeneity in present bias reinforces the importance of corrective retirement savings policies, while the rise in labor income inequality in the U.S. has brought redistributive policies to the center of the policy discussion in recent years. An under appreciated fact by both economists and policy makers is that savings policy and redistributive policy are likely to interact in non-trivial ways. For example, a policy that correct savings might reduce incentives for working as behavioral workers cannot allocate consumption the way they think is best, making it harder for the government to redistribute. Thus understanding these interactions is of key importance to the public economics and behavioral economics literature.

In this paper, we develop a two-period model with unobservable heterogeneity in labor earnings ability and in the level of present bias to study the interaction between policies aiming at correcting savings choices and policies aiming at reducing economic inequality. Our main finding concerns the shape of optimal retirement savings policies across different earnings levels when the government has a redistributive motive. At low earnings, policy forces individuals with the same

ability but different present bias to save at the same rate. Thus optimal policy is paternalistic and leaves little flexibility for agents in savings decisions. At high earnings, optimal policy is less paternalistic and it offers individuals with the same ability but different present bias more flexibility on their savings choices. As a result, at low earnings optimal policy resembles forced savings through social security and at high earnings there are tailored subsidies or taxes on savings that resemble the availability of multiple individual retirement savings accounts, similar to 401(k)s or IRAs in the United States.

Three key forces in the model interact to generate our main finding. First, concern for present bias leads the government to help people increase their retirement savings. Second, the government uses all available tools to reward hard work so as to improve redistribution. There are two different ways to achieve this in our model: allowing individuals to keep some of their earnings and allowing individuals to indulge to their present bias. Third, as there is unobservable heterogeneity in present bias, the government wants to curtail flexibility on savings. By restricting the menu of savings options available to individuals, the government provides social insurance against the behavioral bias. This is true even if retirement savings are not at the first best level. Because of the interaction of those three forces, there is a push for paternalistic savings policies, whereby the government favors forced savings at low earnings, but at the same time a driving force for more flexibility so that high earnings individuals can indulge to their heterogeneous biases.

As usual in the public economics literature, there are different mechanisms the government can use to implement optimal policy, however we provide an implementation that uses three sets of policy tools. First, the government grants retirees a retirement benefit whose level depends on labor earnings during working life. In addition, young people cannot borrow against their retirement benefits. Second, there is both a regular savings account and multiple special retirement savings accounts. There is a cap on contributions to each special retirement savings account, and the cap depends on the labor earnings level of the person. Lastly, there are taxes on labor earnings and on savings returns. The labor income tax is non-linear in labor earnings and depends on which special retirement savings accounts the person contributes to. Taxes on the regular savings account are linear on savings but depend on earnings. Finally, taxes on special retirement savings accounts are lower than the tax on the regular savings account, and these taxes can also vary by the earnings level. This implementation is particularly suitable for a comparison with current policy. The set of instruments we choose is very close to the set of instruments actually used by the U.S. government. For instance, retirement benefits directly translate into social security benefits and special retirement accounts resemble defined contribution plans such as a 401(k) account and individual retirement accounts (IRA). To the best of our knowledge, this is the first paper to highlight that a multitude of special retirement accounts can be used to implement optimal retirement savings policies when there is redistribution.

We find that an optimal retirement savings policy offers both social security old-age benefits and 401(k)-like retirement accounts to high earners, and only social security old-age benefits to low earnings individuals. Importantly, 401(k)-like retirement savings accounts are not available at low labor earnings. Indeed, low earnings individuals rely heavily on retirement benefits. High earnings individuals that have less present bias use retirement savings accounts, but high earnings individuals that have severe present bias rely more on social security.

In order to compare current U.S. policy and the normative model prescriptions, we calibrate the distribution of discounting preferences and labor earnings ability in the data. We use current policies and data on wealth at retirement from the Health and Retirement Survey to calibrate for the joint distribution of discount factors and labor earnings. We find a considerable level of heterogeneity in discount factors and a small positive correlation with earnings. The level of heterogeneity we find in discounting of retirement savings is comparable to the one obtained in the benchmark calibration. Therefore, for the sake of not introducing a more complicated model only for the calibration section, in the main text we present only the calibration of the two-period model. These findings are consistent with estimations of heterogeneity in discount factors in the literature and also with a high heterogeneity in present bias found in the behavioral economics literature.

Simulation of optimal policy in our model is complicated because of the government's two-dimensional screening problem. It is well known that results for multidimensional screening problems are hard to prove in closed form, it is also true that the presence of thousands of incentive constraints makes numerical solutions equally difficult (Judd and Su, 2006). To this end, we develop a stable numerical solution algorithm to solve our model. Our algorithm searches for the smallest subset of incentive constraints that are relevant to the global solution of the problem.

Using calibrated parameters, one still needs to fix government preferences to solve for optimal policy. The government has a discount factor and a redistributive motive. The government discounts retirement consumption of individuals using its discount factor (paternalism) and assigns welfare weights to different individuals depending on their earnings ability (redistribution). We use two alternative procedures to choose the government preferences. First, we develop a benchmark procedure in which we choose government preferences by approximating the normative model and the calibrated model allocation of consumption and earnings under current policy. This follows the approach in Heathcote et al. (2014) and guarantees that the level of redistribution in the benchmark normative model we choose is similar to the one calibrated for the U.S. economy. In this benchmark, the government has a discount factor that is in the intermediary range of calibrated discount factors for individuals, so that there are both individuals that save too little and individuals that save too much from the government's perspective. Furthermore, the government redistributive motive is considerably less progressive than utilitarian. This is consistent with the findings in the literature. However, this

procedure does not guarantee a perfect approximation, accordingly we find that there is a gain available to the government of 18% in consumption equivalent terms at the benchmark normative model.

Moreover, we find in striking difference between optimal retirement benefits and current social security benefits in the United States. We find that social security benefits are considerably lower than retirement benefits prescribed by the benchmark normative model, particularly so at high earnings. As the government is less progressive than utilitarian in the benchmark normative model, there is an important concern about the savings choices of higher earnings individuals. Therefore, government may find it optimal to force a higher minimum level of savings even at high earnings.

Concluding Remarks

In summary, this doctoral dissertation highlights the role of firms in determining earnings inequality (Chapters 1 and 2), wealth inequality (Chapter 3), and the role of public policy in this context (Chapters 3 and 4).