

Supply, Demand, Institutions, and Firms: A Theory of Labor Market Sorting and the Wage Distribution

Daniel Haanwinckel *

August 10, 2020

Abstract

This paper builds a task-based, imperfectly competitive labor market model and estimates it using linked employer-employee data from Brazil. The model matches several measures of wage inequality and generates realistic firm-worker sorting patterns, firm wage premiums, and minimum wage spillovers. I decompose changes in wages and sorting into contributions of education, technology, minimum wage, and other shocks. The minimum wage is the main driver of falling inequality, while rising assortativeness is due to skill-biased technical change. I also show that firm heterogeneity and imperfect competition can qualitatively alter the effect of supply, demand, and institutional shocks on the wage distribution.

*Department of Economics, UCLA. Contact information: haanwinckel@econ.ucla.edu

I would like to thank David Card, Fred Finan, Andrés Rodríguez-Clare, and especially Pat Kline for their guidance and support. Brian Kovak provided valuable feedback as a discussant. The paper also benefited from conversations with Ben Faber, Thibault Fally, Cecile Gaubert, Nicholas Li, Juliana Londoño-Vélez, Piyush Panigrahi, Raffaele Saggio, Yotam Shem-Tov, Avner Strulov-Shlain, José P. Vásquez-Carvajal, Christopher Walters, and seminar participants at Berkeley, UCL, CMU (Heinz College), Northwestern, UCLA, UW-Madison, UPenn, Harvard, Penn State's New Faces in International Economics Conference, the BFI Workers and Firms Workshop, and the Saieh Family Fellows Workshop. I also thank Lorenzo Lagos and David Card for providing code to clean the RAIS data set, and Gustavo de Souza for providing task content data for Brazilian occupations.

1 Introduction

A central task in labor economics is identifying the source of changes in wage distributions. Many papers that address this issue focus on the interplay between the supply of skills and shocks that affect relative demand for certain types of workers, such as skill-biased technical change (henceforth SBTC).¹ Other papers study the role of labor market institutions such as minimum wages.² Finally, more recent literature suggests that when labor markets are not perfectly competitive, wage inequality trends might reflect changes in cross-firm wage dispersion for similar workers. These trends can also reflect changes in firm-worker sorting patterns; specifically, which types of workers are more likely to be matched to firms offering high wage premiums.³

These three potential explanations—supply-demand interactions, institutions, and cross-firm wage differentials—are not mutually exclusive. Each factor may be more relevant in particular periods or for some part of the wage distribution. They may also interact with each other. For example, [Engbom and Moser \(2018\)](#) find that a rising minimum wage decreases the dispersion of firm wage premiums, strengthening the inequality-reducing effects of that shock.

However, the literature still lacks a quantitative framework that combines these three possible causes. This is a limitation for applications where all of these factors are potentially important, because model-based decompositions of inequality trends into their underlying causes can be biased if a relevant factor is omitted. In addition, a unified approach can reveal previously undocumented interactions. For example, it remains unknown whether technical change, labor supply shocks, or minimum wages could be the reason behind changes in labor market sorting that have been observed in some countries, including the US.

In this paper, I propose a tractable model that captures the equilibrium effects of labor supply shocks, technical change, and minimum wages on the wage distribution, while allowing for realistic firm-worker sorting patterns and firm wage premiums. The first half of the paper describes the model and discusses its theoretical properties, which I summarize below. The second half presents a quantitative exercise based on linked employer-employee data from

¹Leading examples in this literature are competitive models with an aggregate CES production function ([Katz and Murphy, 1992](#); [Bound and Johnson, 1992](#); [Krusell et al., 1999](#)) and task-based assignment models of the wage distribution ([Sattinger, 1975](#); [Teulings, 1995](#); [Acemoglu and Autor, 2011](#)).

²See, e.g., [DiNardo, Fortin and Lemieux \(1996\)](#); [Lee \(1999\)](#); [Harasztosi and Lindner \(2019\)](#).

³[Card, Heining and Kline \(2013\)](#) and [Song et al. \(2018\)](#), using reduced-form decompositions, find that changes in sorting account for one third of recent increases in inequality in Germany and the US, respectively.

Brazil. The estimated model successfully replicates a rich set of transformations in the data: falling wage inequality within education groups, wage polarization along education groups (workers with complete high school losing relative to both lower- and higher-educated workers), decreasing dispersion of cross-firm wage differentials, and rising assortativeness.

In the main quantitative exercise, I simulate counterfactuals that isolate the labor market effects of six time-varying factors: workforce composition along educational levels, the skill bias of technology, the minimum wage, a cross-firm wedge that generates wage premiums (linked to differences in entry costs and workplace amenities across firms producing different goods), a change in relative demand for skill-intensive goods (proxying for trade shocks), and a trend in the dispersion of worker productivity (proxying for other unmodeled factors). I find that the minimum wage is the main cause behind falling wage inequality in my setting. It accounts for a little less than half of the decline in the variance of log wages, and 80% of the reduction of the gap between the 50th and 10th percentiles of the wage distribution. SBTC is responsible for increases in the college wage premium and labor market sorting. The large observed increases in schooling achievement have relatively small effects on the variance of log wages, but they reduce between-group wage differentials and make the labor market less assortative.

The paper makes conceptual and technical contributions on two fronts. The first is the task-based production function. Production requires combining tasks of different complexity levels, with task requirements depending on the good the firm decides to sell. More skilled workers have a comparative advantage in more complex tasks. The production function is defined as the output produced by a group of workers when they are optimally assigned to tasks. I show that these assumptions impose strong but reasonable constraints on substitution patterns across worker types: workers who are close in skill level are substitutes, while those who are far apart in skill are complements. This result extends the work of [Teulings \(2000\)](#), who named this property "distance-dependent substitution."

This production function provides a tractable, intuitive and parsimonious way to model firm-worker sorting patterns and technical change. In most models of labor market sorting, either there is no distinction between a firm and a job, or workers are assumed to be perfect substitutes within firms—meaning that jobs at the firm are all the same. Labor market imperfections are the only reason for observing within-firm dispersion in skills or wages in these models.⁴ In contrast, firms in my model hire workers of multiple types to benefit from the

⁴Papers in this group include [Shimer \(2005\)](#); [Gautier, Teulings and van Vuuren \(2010\)](#); [Eeckhout and](#)

division of labor, even when labor markets are competitive. This distinction is important for modeling changes in sorting over time and decompositions of wage inequality in between-firm and within-firm components. I provide a parsimonious and computationally efficient parameterization that allows the use of this production function in complex quantitative applications. In that parameterization, SBTC is defined as increased demand for more complex tasks in the production of all goods.⁵

The task-based production function does not require data on tasks. Nevertheless, I use a measure of task content of occupations to test the microfoundation described above. Schooling correlates with analytical task content both within and between firms. This result is consistent with cross-firm differences in skill intensity being driven by task requirements. In addition, when workers transition to firms where their new colleagues are more educated than the previous ones, they move to more analytical occupations. That result mirrors cross-firm differences in optimal assignment when labor markets are imperfectly competitive.

The other conceptual contribution of the paper is combining monopsony power, endogenous firm entry, and minimum wages in a general equilibrium model of wages. As in [Card et al. \(2018\)](#), firms are horizontally and vertically differentiated in terms of workplace amenities. Firms choose the good they produce upon entry and can set wages below the marginal product of labor, extracting rents from infra-marginal employees who enjoy working there. Markets for goods are competitive.

In equilibrium, firms producing different goods might differ in wages offered to similar workers for three reasons. The first is good-specific entry costs, such that some firms operate at larger scales and post higher wages to attract more workers. The second is differences in average workplace amenities by good, which generate compensating wage differentials. The third is differences in relative task requirements, which cause firms to differ in skill

[Kircher \(2011\)](#); [Gautier and Teulings \(2015\)](#); [Lise, Meghir and Robin \(2016\)](#); [Grossman, Helpman and Kircher \(2017\)](#); [Hagedorn, Law and Manovskii \(2017\)](#); [de Melo \(2018\)](#); [Eeckhout and Kircher \(2018\)](#); [Lise and Postel-Vinay \(2020\)](#). A few papers, such as [Kremer and Maskin \(1996\)](#), [Herkenhoff et al. \(2018\)](#), and [Mak and Siow \(2018\)](#), model firms as teams of two workers executing different jobs.

⁵Models of hierarchical firms in the tradition of [Garicano \(2000\)](#), [Garicano and Rossi-Hansberg \(2006\)](#), and [Antràs, Garicano and Rossi-Hansberg \(2006\)](#) can, in principle, have an arbitrary, endogenously determined set of jobs performed by workers of different skills, such as in this paper. However, they are in general less tractable; quantitative applications resort to simplifications even in competitive environments, such as limiting hierarchies to two levels in [Caliendo and Rossi-Hansberg \(2012\)](#). My production structure can be seen as a version of hierarchical firm models where problems are reinterpreted as tasks and tractability is gained by eliminating the cost of information transmission within the firm. [Eeckhout and Pinheiro \(2014\)](#) and [Trottner \(2019\)](#) also model large firms with multiple jobs, but with common elasticities of substitution across all pairs of worker types.

composition and pay more to the worker types in which they are more intensive. Because of the latter component and the minimum wage, firm wage premiums vary by worker type. Thus, under the lens of the model, two-way fixed effects regression models of log wages (henceforth AKM regressions, after [Abowd, Kramarz and Margolis, 1999](#)) are misspecified. Nevertheless, the widely used decomposition of the variance of log wages based on AKM regressions is still useful, because it helps identify parameters governing between-firm wage dispersion and labor market sorting.

I show that assumptions about consumer behavior have implications for long-run comparative statics. Consider, for example, an increase in the number of college-educated workers. Labor costs decrease more in skill-intensive firms, affecting prices and reallocating consumption toward the goods these firms produce. That, in turn, increases demand for skilled labor and partially offsets the negative impact of the supply shock on skilled worker's wages. If the elasticity of substitution between goods is high enough, this supply shock can even increase mean log wages among college-educated workers relative to lower educated workers. That counterintuitive result is possible if skill-intensive firms also pay higher wages conditional on skill. Skilled wages decrease within firms, but mean log wages rise because a larger share of skilled workers benefit from high wage premiums. The idea that supply, demand, and institutional shocks have secondary impacts on the wage distribution when they change the composition of jobs in the economy is not new.⁶ My contribution in this area is offering a multi-factor quantitative framework that includes this channel.

Finally, the effects of the minimum wage in the model are realistic. Empirical studies of the minimum wage typically find small unemployment effects, spillovers (wage increases on quantiles of the wage distribution where the minimum wage does not bind), and the existence of spikes in histograms of log wages (sometimes called bunching at the minimum wage). The estimated model replicates these patterns, though model-simulated spillovers do not precisely match reduced-form estimates based on Brazilian data. To my knowledge, the only other equilibrium models consistent with these three empirical findings are [Flinn \(2006\)](#) and [Butcher, Dickens and Manning \(2012\)](#).

The paper is organized as follows. The remainder of the introduction compares my paper to recent work on wage inequality in Brazil. The next section presents the task-based pro-

⁶Examples of such papers, organized by type of aggregate shock, are [Kremer and Maskin \(1996\)](#), [Acemoglu \(1998\)](#), [Acemoglu \(1999\)](#), [Mak and Siow \(2018\)](#), [Shephard and Sidibe \(2019\)](#), and [Blundell, Green and Jin \(2020\)](#) (labor supply); [Acemoglu and Restrepo \(2018\)](#) (automation); [Acemoglu \(2001\)](#) (minimum wage); [Sampson \(2014\)](#) and [Davis and Harrigan \(2011\)](#) (trade liberalization).

duction function. The third section describes the general equilibrium model. The fourth section contains the quantitative exercises. The final section concludes with a discussion of directions for further research.

Comparison to recent papers that study the Brazilina case: In this brief section, I restrict attention to papers that use equilibrium models to identify causes of falling inequality in Brazil (see [Firpo and Portella, 2019](#), for a broader review). [Manacorda, Sanchez-Paramo and Schady \(2010\)](#), [Fernández and Messina \(2018\)](#), and [Acosta et al. \(2019\)](#) focus on supply-demand dynamics using models with CES aggregate production functions. These papers find that labor demand has either been roughly stable or has trended toward low-skilled labor. Because their models do not include cross-firm wage differentials or minimum wages, these factors can be absorbed into the estimated demand trend, blurring its interpretation.

[Mak and Siow \(2018\)](#) develop a model combining occupational choice and matching of workers executing different functions in two-worker teams. They fit the model to Brazilian data, simulate a counterfactual by imputing observed changes in the skill composition of the workforce, and find that this labor supply shock reduces inequality.

My paper is closer to previous work by [Engbom and Moser \(2018\)](#), who also build an imperfectly competitive labor market model with firm and worker heterogeneity and a minimum wage. Like me, they use moments from AKM decompositions as targets in an indirect inference estimation procedure. They also compare simulated minimum wage spillovers to reduced-form estimates obtained from specifications similar to those in [Autor, Manning and Smith \(2016\)](#). Their frictional model with on-the-job search makes several predictions that do not exist in mine, regarding worker transitions across firms and the existence of a job ladder. On the other hand, they do not model a market for goods or technical change, their counterfactual exercises focus exclusively on the role of the minimum wage, and wage histograms simulated from their model do not have spikes at the minimum wage. Despite those differences, and the fact that they use data for the whole country instead of a single state, their main conclusion is consistent with my results: minimum wages account for about half of the decline in the variance of log wages in Brazil.

Compared to that literature, my paper provides a more complete and nuanced account of recent labor market transformations in Brazil, because it uniquely disentangles the equilibrium effects of labor supply, two types of demand shocks, minimum wages, and cross-firm wedges that generate firm wage premiums.

2 The task-based production function

Task-based models of comparative advantage are increasingly used to model wage inequality. [Acemoglu and Autor \(2011\)](#) show that these models are better suited than the "canonical" constant elasticity of substitution (CES) model of labor demand to study inequality trends in the US. [Teulings \(2000, 2003\)](#) shows that substitution patterns implied by assignment models make them particularly suitable for studying minimum wages. [Costinot and Vogel \(2010\)](#) develop a task-based model to study the consequences of trade integration and offshoring, finding that it offers new perspectives relative to workhorse models of international trade.

In this section, I show an additional advantage of the task-based approach: it allows for intuitive, tractable, and parsimonious modeling of firm heterogeneity, whereby firms have production functions with imperfect substitution and differ in their demand for skill.

The production structure in this paper is built upon four assumptions. First, final goods embody a set of tasks that vary in complexity, combined in fixed proportions. Second, tasks cannot be traded. Third, workers are perfect substitutes in the production of any particular task, but with different productivities. Fourth, some worker groups have comparative advantage in the production of complex tasks relative to others.

I start this section by defining the production function and solving the managerial problem of assigning workers to tasks. The second subsection discusses cost minimization and shows how this structure generates differences in skill intensity between firms. The third subsection derives and explains distance-dependent substitution. The final subsection presents the parametric version that is employed in the quantitative exercises of this paper. All proofs are in Appendix A.

2.1 Setup, definitions, and the assignment problem

Workers are characterized by their type $h \in \{1, \dots, H\}$ and the amount of labor efficiency units they can supply, $\varepsilon \in \mathbb{R}_{>0}$. They use their labor to produce tasks which are indexed by their complexity $x \in \mathbb{R}_{>0}$. All labor types are perfect substitutes in the production of any particular task, but their productivities are not the same:

Definition 1. *The **comparative advantage function** $e_h : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ denotes the rate of conversion of worker efficiency units of type h into tasks of complexity x . It is continuously differentiable and log-supermodular: $h' > h \Leftrightarrow \frac{d}{dx} \left(\frac{e_{h'}(x)}{e_h(x)} \right) > 0 \forall x$.*

To fix ideas, consider two workers, whom I will refer to as Alice and Bob. Alice, characterized by h, ε , can use a fraction $r \in [0, 1]$ of her time to produce $re_h(x)$ tasks of complexity x . Bob (h', ε'), who belongs to a lower type ($h' < h$), can still produce more of those tasks than Alice, so long as his quantity of efficiency units is high enough relative to hers ($\varepsilon' > \varepsilon e_h(x)/e_{h'}(x)$). But Alice has a comparative advantage: moving towards more complex tasks increases her productivity relative to Bob's.

The interpretation of task complexity depends on how worker groups are defined. In the quantitative exercise of this paper, workers are grouped by educational achievement, and thus more complex tasks are those that benefit from formal education (or intrinsic characteristics that correlate with formal education). The assumption that all tasks are ordered in a single dimension of complexity is strong, but useful as an approximation that allows for complexity in other dimensions. Quantitative models using multi-dimensional skills and tasks include [Lindenlaub \(2017\)](#) and [Lise and Postel-Vinay \(2020\)](#).

Because workers in the same group differ only in a proportional productivity shifter, the sum of efficiency units of each type is a sufficient statistic for analyzing production. Thus, throughout this section, definitions and results are in terms of total efficiency units of each type available to the firm, which I denote by $\mathbf{l} = \{l_1, \dots, l_H\}$ (bold-faced symbols denote vectors over worker types throughout the paper). The distinction between labor efficiency units and workers will be relevant in the next section, when discussing labor markets and the wage distribution.

There is a discrete number of final consumption goods, $g = 1, \dots, G$. Each good is produced by combining tasks in fixed proportions:

Definition 2. *The **blueprint** $b_g : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ is a continuously differentiable function that denotes the density of tasks of each complexity level x required for the production of a unit of consumption good g . Blueprints satisfy $\int_0^\infty b_g(x)/e_H(x)dx < \infty$ (production is feasible given a positive quantity of the highest labor type).*

Tasks cannot be traded; firms must use their internal workforce to produce them. The justification for this assumption is that there are unmodeled costs that make task exchange between firms unprofitable, in the spirit of [Coase \(1937\)](#).⁷ I assume that firms are allowed to split worker's time across tasks in a continuous way by choosing assignment functions

⁷If tasks are freely traded, the model makes no predictions about sorting of workers to firms. A less extreme assumption — e.g. formally modeling output losses from assembling tasks produced at different firms — could be used for studying the boundaries of the firm and the effects of outsourcing.

$m_h : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{\geq 0}$, where $m_h(x)$ denotes the intensity of use of efficiency units of labor type h on tasks of complexity x . The only restriction imposed on $m_h(\cdot)$ is that these functions are right continuous.⁸ That formulation of the assignment problem is very general, allowing firms to use multiple worker types to produce the same task, the same worker type in disjoint sets of tasks, and discontinuities in assignment rules.

Given a blueprint $b(\cdot)$ and l efficiency units of labor, firms choose these assignment functions with the goal of maximizing output. In this problem, they are subject to two constraints: producing the required amount of tasks of each complexity level x and using no more than l_h units of labor of type h .

Definition 3. *The task-based production function $f : \mathbb{R}_{\geq 0}^{H-1} \times \mathbb{R}_{>0} \times \{b_1(\cdot), \dots, b_G(\cdot)\} \rightarrow \mathbb{R}_{\geq 0}$ is the value function of the following assignment problem:⁹*

$$\begin{aligned} f(\mathbf{l}; b_g) &= \max_{q \in \mathbb{R}_{\geq 0}, \{m_h(\cdot)\}_{h=1}^H \subset RC} q \\ \text{s.t.} \quad qb_g(x) &= \sum_h m_h(x) e_h(x) \quad \forall x \in \mathbb{R}_{>0} \\ l_h &\geq \int_0^\infty m_h(x) dx \quad \forall h \in \{1, \dots, H\} \end{aligned}$$

where q is output and m_h is an assignment function denoting the density of labor efficiency units of type h used in the production of each task x . RC is the space of right continuous functions $\mathbb{R}_{>0} \rightarrow \mathbb{R}_{\geq 0}$.

Comparative advantage implies that the optimal assignment of workers to tasks is assortative:

Lemma 1 (Optimal allocation is assortative). *For every combination of inputs $\mathbf{l}, b_g(\cdot)$, there exists a unique set of $H - 1$ complexity thresholds $\bar{x}_1 < \dots < \bar{x}_{H-1}$ that defines the range of tasks performed by each worker type in an optimal allocation: $m_h(x) > 0 \iff x \in [\bar{x}_{h-1}, \bar{x}_h)$, with $\bar{x}_0 = 0$ and $\bar{x}_H = \infty$. Thresholds satisfy:*

$$\frac{e_{h+1}(\bar{x}_h)}{e_h(\bar{x}_h)} = \frac{f_{h+1}}{f_h} \quad h \in \{1, \dots, H-1\} \quad (1)$$

⁸ $\forall x, \tau \in \mathbb{R}_{>0}, \exists \delta \in \mathbb{R}_{>0}$ such that $x' \in [x, x + \delta) \Rightarrow |m_h(x) - m_h(x')| < \tau$.

⁹The definition of the production function assumes positive input of the highest worker type. This assumption simplifies proofs and ensures well-behaved derivatives, while not being restrictive for the applications in this paper. In general, blueprints might require at least one worker of a minimum worker type \underline{h} — if none is available, lower types have zero marginal productivity. This property might be useful for models of endogenous growth and innovation.

where $f_h = \frac{d}{dl_h} f(\mathbf{l}, b_g(\cdot))$ denotes marginal product of labor h , which is strictly positive.

Lower types specialize in low complexity tasks and vice-versa. Equation (1) means that the shadow cost of using neighboring worker types is equalized at the task that separates them. This result is useful for obtaining compensated labor demands, as described in the next subsection.¹⁰

2.2 Compensated labor demand and sorting of workers to firms

To study the properties of this production function, I start by considering its implications in a competitive labor market, where the cost of acquiring efficiency units of each type is given by $\mathbf{w} = \{w_1, \dots, w_H\}$. When firms choose labor quantities by minimizing production costs, marginal productivity ratios equal wage ratios. It then follows from Equation (1) that:

$$\frac{e_{h+1}(\bar{x}_h)}{e_h(\bar{x}_h)} = \frac{w_{h+1}}{w_h}$$

Because the ratio on the left-hand side is strictly increasing in \bar{x}_h , this expression pins down all task thresholds as functions of wage ratios and comparative advantage functions. Since neither are firm-specific, thresholds are common across firms in competitive economies.

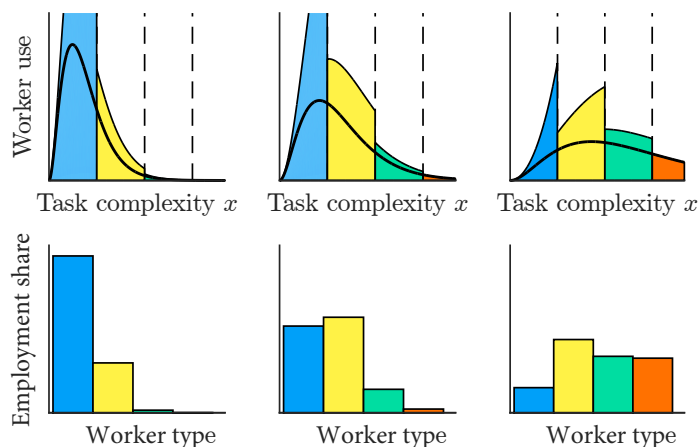
The compensated labor demand is then given by:

$$l_h(q, b_g, \mathbf{w}) = q \int_{\bar{x}_{h-1}(w)}^{\bar{x}_h(w)} \frac{b_g(x)}{e_h(x)} dx \quad (2)$$

Figure 1 illustrates how differences in blueprints reflect into differences in the internal workforce composition of firms. The graphs at the top show the compensated labor demand integral above. The heavy, continuous line is the blueprint, which varies across graphs (becoming more intensive in high complexity tasks from left to right). The vertical dashed lines are the thresholds defining ranges of tasks assigned to each worker type. The colored areas represent the labor demand integrals from Equation 2. The bottom panels show corresponding factor intensities as histograms.

¹⁰In general, the task-based production function and its derivatives do not have simple closed-form representations. If one needs to evaluate output and marginal productivities as a function of labor inputs, first solve the system of H compensated labor demand equations (2) on q and the $H - 1$ thresholds. Next, use equation (1) to calculate marginal productivity gaps. Finally, use the constant returns relationship $q = \sum_h l_h f_h$ to normalize marginal productivities.

Figure 1: Compensated labor demand in competitive labor markets



If labor markets are not competitive, as in labor market model described in the next section, thresholds might differ across firms. Firms using different blueprints will still differ in the skill composition of their internal workforce, though possibly less so than in the competitive benchmark.

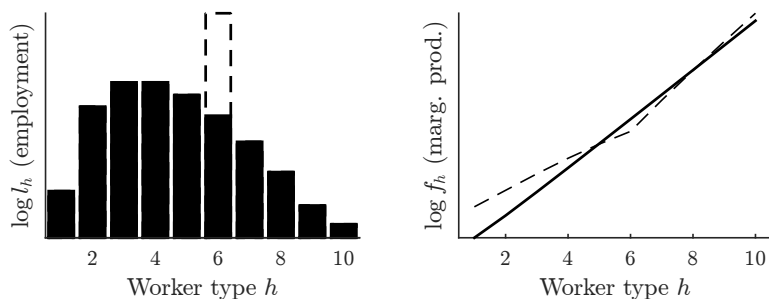
The concept of firms in this model is significantly different from that in the literature on labor market sorting. Most models in this literature focus on sorting of workers to *jobs* (or, equivalently, to firms that employ exactly one worker). Even in the ones with a concept of large firms, such as [Eeckhout and Kircher \(2018\)](#), any degree of within-firm wage dispersion is a sign of inefficiencies introduced by search frictions; if markets are competitive, each firm hires workers of a single type. I contribute to this literature by introducing a more realistic concept of firms as bundles of jobs (tasks in my model), coupled with a technology to acquire workers. In addition to having welfare implications, this distinction is relevant for quantitative studies where model predictions are matched to firm-related moments, such as the between-firm share of wage inequality or variance decompositions from AKM regressions.

2.3 Substitution patterns and distance-dependent complementarity

The task-based structure might appear exceedingly flexible at first glance, due to the infinite-dimensional blueprints and efficiency functions. [Proposition 1](#) extends the results in [Teulings \(2005\)](#) and shows that, on the contrary, there are strong constraints on substitution patterns.¹¹

¹¹[Teulings \(2005\)](#) derives elasticities of complementarity for a similar model, but using parametric efficiency functions and taking a limit where the number of worker types grows to infinity. In an application of

Figure 2: Distance-dependent complementarity



Locally, the $H \times (H - 1)/2$ partial elasticities of complementarity or substitution depend only on factor shares and at most $H - 1$ scalars ρ_h — the same number of elasticity parameters in an equally-sized nested CES structure. However, unlike with a CES, there is a straightforward way to impose further restrictions on the number of parameters (both elasticities and productivity shifters for each worker type), via parameterization of blueprints and efficiency functions.

Proposition 1 (Curvature). *The task-based production function is concave, has constant returns to scale, and is twice continuously differentiable with strictly positive first derivatives. Appendix A provides formulas for elasticities of complementarity and substitution.*

The curvature of the task-based production function reflects division of labor within the firm. Suppose that, initially, a firm only employs Alice, who belongs to the highest type H . In that case, output is linear in the quantity of labor bought from Alice. Adding another worker, Bob, of a lower type increases Alice’s productivity, because she can now specialize in complex tasks while Bob takes care of the simpler ones. At that point, decreasing returns to Alice’s hours reflect a reduction in those gains from specialization.

The impact of adding a third worker, Carol, on the marginal productivities of Alice and Bob depends on Carol’s skill level (in terms of comparative advantage), relative to Alice’s and Bob’s:

Corollary 1 (Distance-dependent complementarity). *For a fixed h , the partial elasticity of complementarity is strictly increasing in h' for $h' \geq h$ and strictly decreasing in h' for $h' \leq h$.*

Close types perform similar tasks and are net substitutes; distant types perform different tasks and are complements. The distance-dependent complementarity pattern is illustrated

assignment models to optimal taxation, [Ales, Kurnaz and Sleet \(2015\)](#) derive elasticities of substitution in a model of production where the unique output is CES in tasks, instead of Leontief.

in Figure 2. The left panel shows baseline log employment by worker type (black bars) and a shock to employment of workers of type 6 (dashed contour). The right panel shows baseline log marginal productivities (solid line) and marginal productivities after the employment shock (dashed). Workers of type 6 suffer the largest relative decline in marginal productivity, followed by neighbor types 7 and 5. Marginal productivities increase for types that are further away, both low-skilled and high-skilled.

2.4 Exponential-Gamma parameterization

Consider the following parameterization, used in the quantitative exercises of this paper:

$$\begin{aligned}
 e_h(x) &= \exp(\alpha_h x) & -1 &= \alpha_1 < \alpha_2 < \dots < \alpha_{H-1} < \alpha_H = 0 \\
 b_g(x) &= \frac{x^{k_g-1}}{z_g \Gamma(k_g) \theta_g^{k_g}} \exp\left(-\frac{x}{\theta_g}\right) & (z_g, \theta_g, k_g) &\in \mathbb{R}_{>0}^3
 \end{aligned}$$

The exponential function is a straightforward way to generate log-supermodularity. Differences in the α_h coefficients determine the degree of comparative advantage between any two worker types. The expression for blueprints is the probability density function of a Gamma distribution divided by a "productivity" term z_g . Doubling z_g divides the quantity of tasks needed per unit of output by two, effectively doubling physical productivity.

Appendix C presents the mapping between marginal productivity gaps and task thresholds in this parametrization, as well as formulas for compensated labor demand integrals in terms of incomplete Gamma functions. These formulas are useful because they dispense with numerical integration, improving computational performance. Incomplete Gamma functions are readily available in software packages commonly used by economists.

The parameter θ_g is related to average task complexity. All else equal, goods with higher θ_g require more complex tasks, and firms producing these goods find it optimal to employ workers of higher types. Increases in task complexity over time, modeled as changes in θ_g , provide an intuitive way to model skill-biased technical change because higher complexity is linked to increasing returns to skill (measured as the worker group h). The shape parameter k_g determines the dispersion of tasks. If two firms differ only in this parameter, the one with the smallest k_g has fatter tails. Thus, differences in k_g in the cross-section translate into some firms being more specialized than others in their hiring patterns.

This approach allows for modeling firm-level differences in skill intensity, skill dispersion,

and productivity with a small number of parameters, while ensuring sensible substitution patterns within all firms. Consider an example of two firms in the retail sector. One is a small local shop, while the other is a large online retailer. In the first one, most tasks are of low complexity, measured in terms of how they benefit from schooling: stocking shelves, operating the register, and cleaning. In those tasks, workers with little formal education can easily substitute for others with a college degree. Because workers with a college degree cost much more, that first firm mostly hires less educated workers. In contrast, the online retailer is intensive in tasks such as web design, system administration, and business analytics, where college-graduated workers usually perform much better. This is why those firms find it profitable to use a more skilled workforce. All of those differences in skill intensity and elasticities of substitution are encoded by a single parameter, θ_g , regardless of the number of worker types in the model.

3 Markets and wages

This section builds a general equilibrium model with monopsonistic firms and free entry. The first subsection lays out the structure of the economy. The second subsection describes the functioning of labor markets, solves the problem of the firm, and shows an important property of the model: goods encapsulate firm heterogeneity in skill intensity and wages. The third subsection derives analytical results on what determines wage differentials between firms and how the wage distribution changes over time.

This is the point of departure from other task-based assignment models of comparative advantage. The contributions of the previous section fit inside that literature: new formulas for elasticities of complementarity and substitution, along with the convenient exponential-gamma parametrization. This section introduces more significant deviations: imperfect competition in labor markets and aggregate demand defined in terms of goods, not individual tasks. Both have implications for comparative statics.

3.1 Factors, goods, technology, and preferences

Consider an economy with $N = \{N_1, \dots, N_H\}$ workers of each type h , and a large number of entrepreneurs. Entrepreneurs own entrepreneurial talent, whose total stock in the economy is T and which is used to create firms. The model is static.

There are G final goods in this economy, which can be interpreted as either different in-

dustries or differentiated varieties within industries.¹² An entrepreneur j may set up a firm producing one good $g \in \{1, \dots, G\}$ or not enter at all. Setting up a firm requires a fixed cost F_g , paid in units of entrepreneurial talent. Once that cost is paid, the entrepreneur receives the blueprint b_g and a random draw of workplace amenities a_j from a good-specific distribution with strictly positive support and a finite mean \bar{a}_g . The role of workplace amenities will be explained below. Hiring and production decisions are done after the amenities draw is observed.

I assume that there is a competitive market for entrepreneurial talent and that entrepreneurs can form coalitions to insure against risk associated with the draw of firm amenities a_j . These assumptions allow me to abstract from the distribution of entrepreneurial talent and to pin down firm entry by equating expected profit and entry costs for each good g :

$$E_{a_j|g} [\pi_g(a_j)] = F_g p_T = F_g \quad \forall g$$

where $\pi_g(a_j)$, defined below, denotes profits achieved by a firm with amenities a_j producing good g . The second equality follows from assuming that entrepreneurial talent is the numeraire in this economy. This choice of numeraire is valid because firms have positive profits, as I will show below, and so the price of entrepreneurial talent cannot be zero. A positive price for entrepreneurial talent also implies that all of it is used up in equilibrium:

$$\sum_g J_g F_g = T \tag{3}$$

where J_g is total entry of firms producing good g . When there is a single good $g = 1$ in this economy, the number of firms is fixed at $J_1 = T/F_1$. But with multiple goods, the number of firms producing each good might respond to shocks.

The utility of entrepreneurs, U^E , is a constant elasticity aggregate of consumption Q_1, \dots, Q_G .

¹²In the empirical exercise, I do not map goods to industries because the within-industry dimension is important. In many contexts, changes in inequality happen within industries (see [Card, Heining and Kline, 2013](#); [Song et al., 2018](#)). The first validation exercise in Subsection 4.6 suggests substantial task heterogeneity within finely defined sectors.

Worker preferences U_i^L depend on consumption and the firm j where they are employed:

$$U^E \left(\{Q_g\}_{g=1}^G \right) = \left[\sum_{g=1}^G \gamma_g Q_g^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{1-\sigma}}$$

$$U_{hi}^L \left(\{Q_g\}_{g=1}^G, j \right) = \left[\sum_{g=1}^G \gamma_g Q_g^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{1-\sigma}} \left[a_j \exp(\eta_{ij}) \right]^{\frac{1}{\beta}} \quad \text{with} \quad \eta_{ij} \sim \text{Extreme Value Type I}$$

Firms matter to workers not only due to their overall level of amenities a_j , but also because of an idiosyncratic component η_{ij} . This component captures match-specific features such as distance to the workplace or personal relationships with the manager or other coworkers. The parameter β measures the importance of consumption relative to these non-pecuniary elements. Higher β implies that labor markets are closer to competitive, as discussed below.

Markets for goods are competitive. Thus, any equilibrium will feature prices p_g equal to the marginal cost of good g at all firms producing that good. There is a price index $P = \left[\sum_{g=1}^G \gamma_g^\sigma p_g^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$ such that consumption level u costs $u \times P$. Because $u(\cdot)$ is homothetic, aggregate consumption is only a function of prices and aggregate income.

Continuing with the example from Section 2.4, the small local shop and the large online retailer are interpreted as differentiated varieties in the retail sector, with elasticity of substitution σ . In addition to task requirements, these firms might differ in entry costs and the average level of amenities. The online retailer might require substantial capital investment or managerial input to set up, justifying high entry costs F_g . If \bar{a}_g is higher for the large retailers, then they are also more desirable workplaces on average.

3.2 Labor markets, the problem of the firm, and equilibrium

Labor markets are based on Card et al. (2018), to which I add minimum wages and use a two-dimensional concept of worker heterogeneity. Each worker is characterized by its type $h \in \{1, \dots, H\}$ and a quantity of efficiency units of labor ε . The distribution of efficiency units of labor across workers of type h is continuous with density $r_h(\cdot)$ and support over the positive real line.¹³ Throughout this section, it is important to distinguish between quantities of workers, denoted by n , and quantities of labor, denoted by l . Worker earnings are denoted

¹³I employ lognormal distributions of ε in the quantitative exercise. Counterfactual exercises require a parametric assumption for $r_h(\cdot)$, which is used to obtain the number of workers driven to unemployment because of the minimum wage and the distribution of ε in that unobserved population.

by y , while prices for efficiency units of labor are denoted by w .

Labor regulations prevent firms from paying a total compensation of less than \underline{y} to any worker. I refer to \underline{y} as the minimum wage; the model has no variation in hours worked, so earnings and wages are interchangeable. Workers with low ε might have a marginal product of labor lesser than \underline{y} at some firms, in which case hiring those workers would be unprofitable. Thus, I allow firms to reject workers with productivity below some minimum value $\underline{\varepsilon}_{hj}$, generating involuntary non-employment.

3.2.1 Firm-level labor supply and labor costs

There are separate labor markets for each worker group h . The timing of each of these labor markets is as follows:

1. Each firm j posts a rejection cutoff $\underline{\varepsilon}_{hj}$ and earning schedules $y_{hj}(\varepsilon) : [\underline{\varepsilon}_{hj}, \infty) \rightarrow [\underline{y}, \infty)$.
2. Workers observe all $\underline{\varepsilon}_{hj}$ and $y_{hj}(\varepsilon)$. Based on that information, they choose firms that maximize their indirect utility. If no firm is chosen, the worker earns zero income.
3. Firms observe (h, ε) of workers who applied to them (but not idiosyncratic preference shifters η_{ij}) and hire those with $\varepsilon \geq \underline{\varepsilon}_{hj}$.
4. Production occurs and hired workers are paid. Rejected workers, if any, earn zero income.

To study worker choices in step 2, consider the indirect utility of a worker i characterized by (h, ε) , if this worker chooses firm j . It can be written as:

$$V_{ih}(\varepsilon, j) = \frac{\mathbf{1}\{\varepsilon \geq \underline{\varepsilon}_{hj}\}}{P} \exp\left(\beta \log y_{hj}(\varepsilon) + \log a_j + \eta_{ij}\right)^{\frac{1}{\beta}}$$

Because η_{ij} is drawn from a Type I Extreme Value distribution, the probability of a worker (h, ε) choosing a particular firm j is given by:

$$P\left(j = \arg \max_{j' \in \{1, \dots, J\}} V_{ih}(\varepsilon, j')\right) = \mathbf{1}\{\varepsilon > \underline{\varepsilon}_{hj}\} a_j \left(\frac{y_{hj}(\varepsilon)}{\omega_h(\varepsilon)}\right)^{\beta}$$

$$\text{where } \omega_h(\varepsilon) = \left(\sum_{j'} \mathbf{1}\{\varepsilon > \underline{\varepsilon}_{hj'}\} a_{j'} y_{hj'}(\varepsilon)^{\beta}\right)^{\frac{1}{\beta}}$$

The "inclusive value" $\omega_h(\varepsilon)$ is a measure of demand for skills in this model. A high value

means that many firms are posting high wages for type h and willing to hire that particular ε , despite the minimum wage. That makes those workers harder to attract for any individual firm because they have good outside options at other firms. As in [Card et al. \(2018\)](#), I assume that firms ignore their own contribution to $\omega_h(\varepsilon)$, an approximation that is adequate when firms are small relative to the size of the labor market. Mechanically, $\omega_h(\varepsilon)$ is a cost shifter taken as given by firms that ensures market clearing, similar to wages in competitive models.

The number of workers choosing a particular firm, the resulting supply of labor, and total labor costs are increasing in posted earnings, decreasing in rejection cutoffs, and linear in amenities:

$$n_h(y_{hj}, \underline{\varepsilon}_{hj}, a_j) = N_h a_j \int_{\underline{\varepsilon}_{hj}}^{\infty} \left(\frac{y_{hj}(\varepsilon)}{\omega_h(\varepsilon)} \right)^\beta r_h(\varepsilon) d\varepsilon \quad (4)$$

$$l_h(y_{hj}, \underline{\varepsilon}_{hj}, a_j) = N_h a_j \int_{\underline{\varepsilon}_{hj}}^{\infty} \varepsilon \left(\frac{y_{hj}(\varepsilon)}{\omega_h(\varepsilon)} \right)^\beta r_h(\varepsilon) d\varepsilon \quad (5)$$

$$C_h(y_{hj}, \underline{\varepsilon}_{hj}, a_j) = N_h a_j \int_{\underline{\varepsilon}_{hj}}^{\infty} \frac{y_{hj}(\varepsilon)^{\beta+1}}{\omega_h(\varepsilon)^\beta} r_h(\varepsilon) d\varepsilon \quad (6)$$

3.2.2 Problem of the firm

Firms maximize profit by choosing posted wages and rejection cutoffs:

$$\pi_g(a_j) = \max_{\mathbf{y}_j, \underline{\varepsilon}_j} p_g f(\mathbf{l}(\mathbf{y}_j, \underline{\varepsilon}_j, a_j), b_g) - \sum_{h=1}^H C_h(y_{hj}, \underline{\varepsilon}_{hj}, a_j)$$

The following Lemma shows that this problem has intuitive solutions and that the model admits a representative firm for each good:

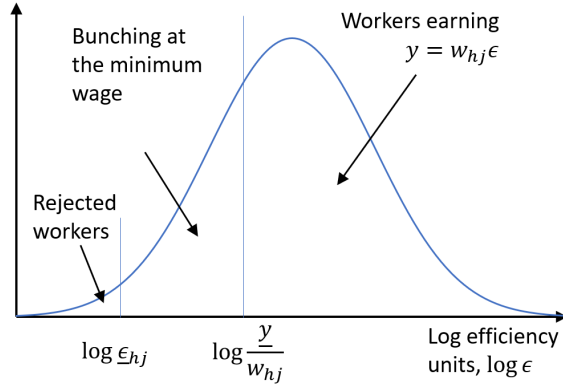
Lemma 2. *Optimal earnings schedules are of the form $y_{hj}(\varepsilon) = \max\{w_{hj}\varepsilon, \underline{y}\}$ for a positive scalar w_{hj} . The solution of the problem of the firm is characterized by the following first order conditions:*

$$p_g f_h(\mathbf{l}(\mathbf{w}_j, \underline{\varepsilon}_j, a_j), b_g) \frac{\beta}{\beta+1} = w_{hj} \quad h = 1, \dots, H \quad (7)$$

$$p_g f_h(\mathbf{l}(\mathbf{w}_j, \underline{\varepsilon}_j, a_j), b_g) \underline{\varepsilon}_{hj} = \underline{y} \quad h = 1, \dots, H \quad (8)$$

Additionally, firms producing good g choose the same wages \mathbf{w}_g and rejection criteria $\underline{\varepsilon}_g$. Output and employment are linear in firm amenities: $q_j = \frac{a_j}{\bar{a}_g} \bar{q}_g$ and $\mathbf{l}_j = \frac{a_j}{\bar{a}_g} \bar{\mathbf{l}}_g$, where \bar{q}_g and

Figure 3: Firm choices and bunching at the minimum wage



\bar{l}_g denotes mean output and mean labor demand for all firms producing good g , respectively.

Figure 3 illustrates how the earnings schedules divide workers into three groups according to their level of efficiency units. Those to the left of $\underline{\epsilon}_{hj}$ are rejected. Those with $\epsilon > \underline{y}/w_{hj}$ earn the wage posted by the firm times their quantity of labor units. Finally, those in the intermediate range earn the minimum wage. Log wage histograms simulated from the model have spikes at the minimum wage corresponding to the mass of workers between the two vertical lines. Bunching at the minimum wage is often observed in the data (DiNardo, Fortin and Lemieux, 1996; Harasztosi and Lindner, 2019) but is not a common feature in models of wage inequality.

The first order conditions represent trade-offs along two different margins: workers above the minimum wage and workers around the rejection threshold. To build intuition on the optimality condition on wages, start by noting that a marginal increase in w_{hj} has no bearing on workers with $\epsilon \in [\underline{\epsilon}_{hj}, \underline{y}/w_{hj})$, because they earn exactly the minimum wage. Denote by l_{hj}^+ the sum of efficiency units at firm j supplied by workers earning more than the minimum wage. A proportional increase in posted wages $d \log w_{hj}$ brings in $(\beta d \log w_{hj}) l_{hj}^+$ additional labor units, generating $(\beta d \log w_{hj}) l_{hj}^+ p_g f_h(\cdot)$ in additional revenues. Labor costs increase for two reasons. First, the firm pays $(\beta d \log w_{hj}) l_{hj}^+ w_{hj}$ for the additional labor purchased. Second, a higher wage increases the wage bill for current workers by $d \log w_{hj} l_{hj}^+ w_{hj}$. Setting added revenues equal to additional costs yields Equation 7.

Equation 8 is the first order condition on the rejection cutoffs. A lower cutoff brings in additional workers with $\epsilon = \underline{\epsilon}_{hj}$, each of which increases revenues by $p_g f_h \underline{\epsilon}_{hj}$. When firms chose thresholds optimally, that additional revenue equals the minimum wage \underline{y} , which is the

cost of labor at that margin.

Lemma 2 also shows that firms producing the same good are equal in wages and input intensities. Dispersion in amenities within good only scales the firm up or down. This result simplifies the analysis of between-firm wage differentials and sorting in this model by restricting the sources of these patterns to differences in blueprints, entry costs, or mean amenities \bar{a}_g . It also simplifies the expression for $\omega(\varepsilon)$, making the computation of labor demands feasible:

$$\omega_h(\varepsilon) = \left(\sum_g J_g \mathbf{1}\{\varepsilon > \underline{\varepsilon}_{hg}\} \bar{a}_g \max\{\varepsilon w_{hg}, \underline{y}\}^\beta \right)^{\frac{1}{\beta}} \quad (9)$$

3.2.3 Equilibrium

An equilibrium of this model is a set of prices $\{p_g\}_{g=1}^G$, aggregate consumption $\{Q_g\}_{g=1}^G$, firm entry $\{J_g\}_{g=1}^G$, and choices by representative firms $\{\mathbf{w}_g, \underline{\varepsilon}_g\}_{g=1}^G$ such that:

1. Markets for goods clear:

$$Q_g = \left[\frac{p_g}{P} \right]^{-\sigma} \frac{I}{P} = J_g \bar{q}_g \quad \forall g \quad (10)$$

$$\text{where } I = T + \sum_{g=1}^G J_g \sum_{h=1}^H C_h(w_{hg}, \underline{\varepsilon}_{hg}, \bar{a}_g)$$

2. For all g , firm choices solve the set of equations (7) and (8).
3. Entrepreneurs have zero ex-ante expected profits:

$$E_{a_j|g} [\pi_g(a_j)] = p_g f(\mathbf{l}(\mathbf{w}_g, \underline{\varepsilon}_g, \bar{a}_g), b_g) - \sum_{h=1}^H C_h(w_{hg}, \underline{\varepsilon}_{hg}, \bar{a}_g) = F_g \quad \forall g \quad (11)$$

4. The market for entrepreneurial talent clears (Equation 3).

Labor market clearing is implied by the definition of $\omega_h(\varepsilon)$, which ensures that the number of job applicants to all firms (calculated using Equation 4) is equal to total number of workers N_h .

Appendix C presents a numerical algorithm to solve for equilibrium efficiently.

3.3 Firm wage premiums

We know from the labor market structure that log earnings of a worker i of type h at a firm producing good g take the form $\log y_{ihg} = \max\{\log w_{hg} + \log \varepsilon_i, \log \underline{y}\}$. Within-group variation of log wages has three components: the dispersion of efficiency units, differences in mean log wages across goods for the same worker type, and censoring by the minimum wage.

The following proposition describes how wages vary across firms producing different goods:

Proposition 2. *1. If $b_g(x) = b(x)/z_g$ and $\frac{F_g}{\bar{a}_g}$ is common across goods, then there are no firm-level wage premiums:*

$$\log y_{ihg} = \max\{\lambda_h + \log \varepsilon_i, \log \underline{y}\}$$

where $\lambda_1, \dots, \lambda_H$ are scalar functions of parameters.

2. If there is no minimum wage ($\underline{y} = 0$) and $b_g(x) = b(x)/z_g$, then wages are log additive in worker type and firm type:

$$\log y_{ihg} = \lambda_h + \frac{1}{1+\beta} \log\left(\frac{F_g}{\bar{a}_g}\right) + \log \varepsilon_i$$

3. If there is no minimum wage and there are firm types g, g' and worker types h, h' such that $\ell_{h'g'}/\ell_{hg'} > \ell_{h'g}/\ell_{hg}$ (that is, good g' is relatively more intensive in h'), then:

$$y_{ih'g'}/y_{ihg'} > y_{ih'g}/y_{ihg}$$

The first part of Proposition 2 shows that wage dispersion for similar workers exists only if there are differences in the shapes of blueprints (such that firms differ in skill intensity) or in the ratio of entry costs to mean amenities. Notably, differences in physical productivity across goods (z_g) or in taste shifters (γ_g) are not enough to generate wage differentials between firms. The reason is that, if the ratio of entry costs to firm amenities is the same, differences in physical productivity or tastes lead to additional entry and reduced marginal utility of consumption for the good with more productivity, up to the point where marginal revenue product of labor is equalized across firms.

The second part highlights the role of entry costs and mean amenities in generating wage

differences across firms. The zero profits condition implies that firms producing goods with higher entry costs need to operate at larger scale. To hire more workers, these firms need to post higher wages, unless the differences in entry costs are exactly offset by differences in mean amenities.

The third part of Proposition 2 shows how heterogeneity in skill intensity generates differential wage gaps across firms. Firms using some factors more intensively than others must pay a relative premium to that factor. Thus, in general, the model cannot generate log-additive wages as in [Abowd, Kramarz and Margolis \(1999\)](#), except when factor intensities do not vary.

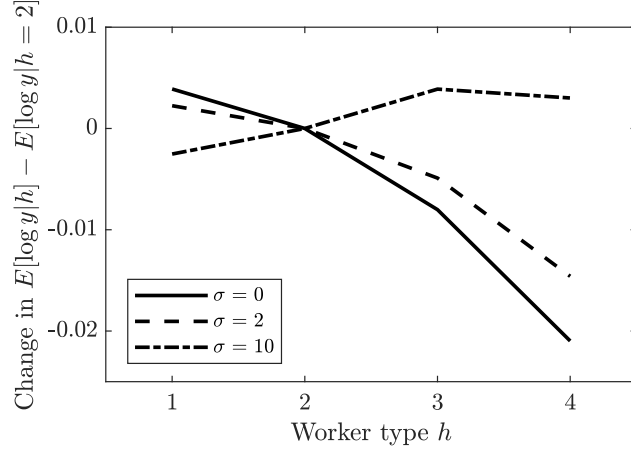
The inability of this model to simultaneously generate log-additive wages and sorting echoes some results in the literature on labor market sorting, such as those in [Eeckhout and Kircher \(2011\)](#). But it is possible that skill-intensive firms pay a positive wage premium for all worker types if those firms have high entry costs relative to amenities. The quantitative exercise shows that this flexibility is necessary for fitting the data.

To provide a concrete example of how firms differ in equilibrium, consider the Exponential-Gamma parametrization introduced in Subsection 2.4. Under that parametrization, goods are fully described by six scalars: blueprint complexity θ_g , blueprint shape k_g , blueprint productivity z_g , taste shifter γ_g , mean amenities \bar{a}_g , and entry costs F_g . These scalars map into six empirical measures regarding firms producing that good: mean worker education, dispersion in worker education, share of workers in the economy employed by those firms, share of those firms in aggregate consumption, mean firm size, and firm-level wage premiums. In the quantitative exercise, we will make further restrictions to make the model more parsimonious, such as assuming that the blueprint shape is the same for all firms and across periods.

3.4 Changes in the wage distribution over time

The final step in the theoretical analysis is understanding how the wage distribution changes over time, given shocks to labor supply, labor demand, and minimum wages. In theory, these shocks may act in concert, making the economy more productive while leaving the wage distribution unchanged (see Proposition 4 in Appendix B). If, however, there are imbalances in the race between education, technology, and minimum wages, then relative prices and the allocation of resources might change.

Figure 4: Simulated effects of an increase in the share of high-skilled workers.



The model is too complex to yield closed-form comparative statics for realistic parameter sets. Counterfactual analysis is the best way to disentangle the role of each shock. It is possible, however, to obtain some intuition about how firm heterogeneity and imperfect competition might amplify or attenuate the impact of specific shocks on the wage distribution, relative to a framework with a representative firm.

As a starting point, the following proposition shows that changes in relative consumption have consequences for wage inequality, even if technology, labor supply and institutions do not change:

Proposition 3 (Changes in relative output affects the returns to skill). *Consider a competitive version of this economy ($\beta = \infty$, $F_g = 0$) with two goods ($G = 2$) and no minimum wages ($\underline{y} = 0$). Assume good $g = 2$ is relatively more intensive in high-complexity tasks, such that $b_2(x)/b_1(x)$ is increasing in x . Then, an increase in the relative taste for the second good (γ_2/γ_1) causes increases in all wage gaps w_{h+1}/w_h .*

Proposition 3 has a more general implication. Supply, demand, and institutional shocks change relative costs for goods. If there is enough substitutability in consumption (i.e., if σ is high), then that substitution will cause a secondary effect on wages. A high value of σ may represent an economy where firms choose different technologies to produce similar goods or a small open economy where all goods are traded.

Figure 4 illustrates this effect by presenting simulations of an increase in supply of highly skilled workers, for varying levels of the elasticity of substitution between goods. In the Leontief case, the impact of labor supply shock on wages is a smooth reduction in inequality

between groups. That result reflects the distance-dependent substitution property of the production function. A moderate increase in substitutability across goods attenuates the wage effects. That's because the supply shock reduces the price of the high-complexity good, reallocating consumption in a way that increases the returns to skill.

The full effect of a shock to labor supply, technical change, or minimum wages includes not only this secondary demand effect but also impacts on wage premiums and sorting. With high substitution, mean log wages for types 3 and 4 increase relative to types 1 and 2. That result, which could not happen in a competitive economy, is due to changes in sorting. The high-types benefit from being reallocated to the high-wage firms, while the lower types remain at the low-wage firms. On average, high-skilled workers benefit from the increase in the supply, even though skilled wages fall conditional on firm types. In the main quantitative exercise, I show that this theoretical result is relevant for understanding changes in wage inequality in Brazil.

4 Wage inequality and sorting in Rio Grande do Sul, Brazil

While the literature covering countries such as the US and Germany often tries to explain increasing wage disparities in recent decades, the academic debate in Brazil attempts to rationalize a decline in inequality starting in the 1990s. The most salient facts in the Brazilian context are significant increases in both the minimum wage and educational achievement, following policies aimed at universal primary schooling in the 1980s and 1990s and expansion in access to college-level education in the 2000s. In this section, I take a sparsely parameterized, over-identified version of the model to the data and study whether it can rationalize changes from 1998 through 2012. Next, I use the estimated model to generate counterfactuals that isolate the labor market impacts of education, skill-biased technical change, minimum wages, and other time-varying factors.

The exercise focuses on the formal sector in the southernmost state of Brazil, Rio Grande do Sul. I use a single state because that is a better empirical counterpart to the well-connected labor market in the model. This is particularly relevant for measuring the role of firms and labor market sorting. Taking the country as a whole, much of firm-worker sorting could plausibly be attributed to regional differences and geographical barriers, factors that are absent from the model. Furthermore, [Kline, Saggio and Sølvssten \(2018\)](#) show that metrics of labor market sorting can be imprecise if the data is composed of multiple regions which are poorly

connected in terms of firm-to-firm job transitions. I choose the particular state of Rio Grande do Sul to limit the relevance of the informal sector. I discuss the implications of ignoring the informal sector after presenting the results of the counterfactual exercises.¹⁴

I employ data from the RAIS (*Relação Anual de Informações Sociais*), a confidential linked employer-employee dataset maintained by the Brazilian Ministry of Labor. Firms are mandated by law to report to RAIS at the establishment level. The dataset I use contains information about both the establishment (including legal status, economic sector, and the municipality in which it is registered) and each worker it formally employs (including education, age, earnings in December, contract hours, and hiring and separation dates).

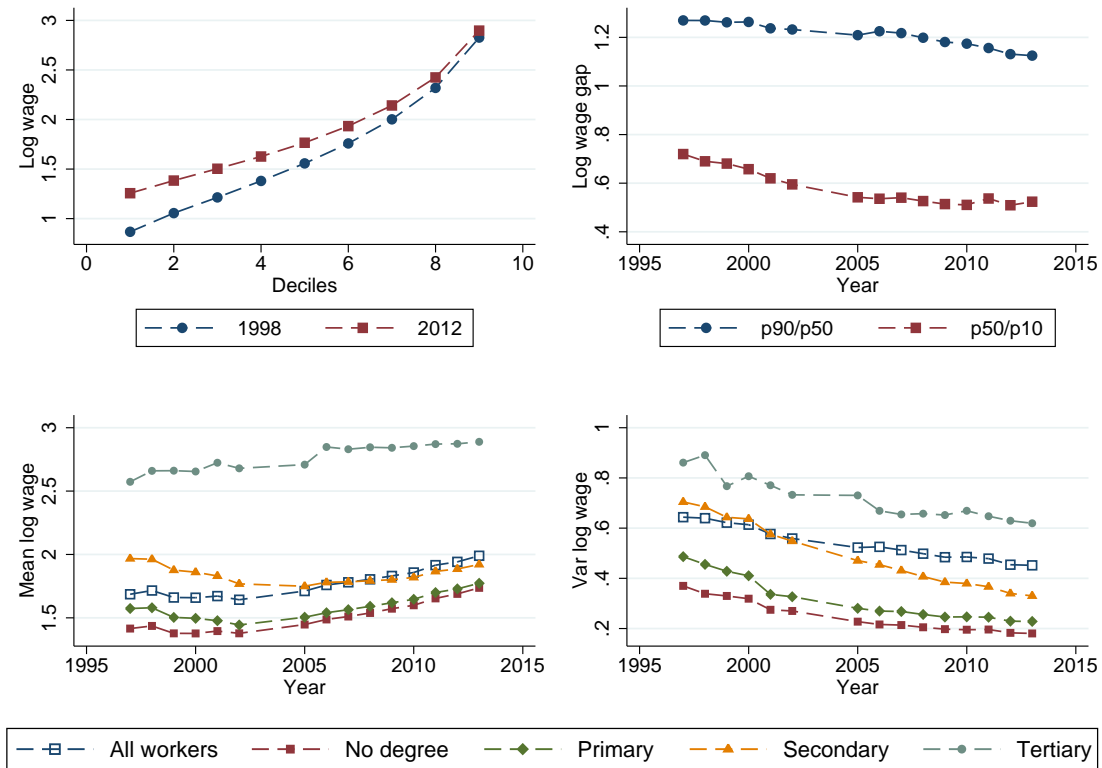
Because I am interested in equilibrium effects, I use a broad sample. I select adults between 18 and 54 years of age, who are not currently in school, and who are working in December having been hired in November or earlier. I only consider the highest-paying job per worker in each year. The resulting data set has 1,494,186 workers and 148,203 establishments in 1998, and 2,398,391 workers and 238,545 establishments in 2012. For each worker, I calculate the hourly wage based on their monthly earnings and contract hours, before winsorizing the bottom and top percentiles of the wage distribution. Summary statistics are provided in Table D1, located in Appendix D.1.

4.1 Empirical context and target features in the data

Figure 5 demonstrates the evolution of wages in the *Rio-Grandense* economy. The top left panel shows that, from 1998 to 2012, real wages have increased for all deciles of the log wage distribution, and particularly so for the lowest deciles. Almost all commonly used measures show a reduction in inequality: upper-tail or lower-tail percentile gaps (top-right panel), differences in mean log wage between workers with secondary education (that is, those complete high-school and college dropouts) and less educated workers, and the variance of log wages — for the sample as a whole and within each educational group. The one

¹⁴Similarly to many other developing countries, a substantial share of the Brazilian workforce is informal (employed at firms that evade regulations such as payroll taxes and minimum wages). Including informal workers is difficult because of data limitations. Households surveys measure employment and wages in the informal sector, but I require matched employer-employee data to gauge between-firm wage dispersion for similar workers and labor market assortativeness. Except for the Southeast, the Brazilian South is the region with the lowest rates of employment informality in the country. The Southeast is less interesting for this exercise, however, because higher wages in that region make the minimum wage shock less relevant. In that regard, Rio Grande do Sul is closer to a "typical" Brazilian state than states in the Southeast. Appendix D.2 contains a thorough analysis of inequality and education patterns using a different data source that includes the informal sector.

Figure 5: Measures of wage dispersion in Rio Grande do Sul, Brazil

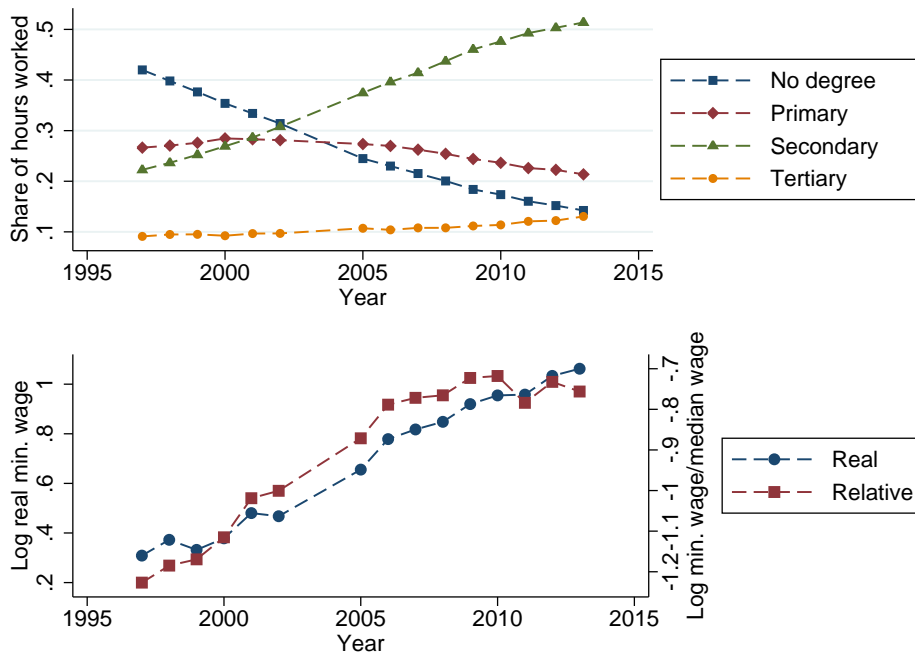


Notes: RAIS data for the formal sector in Rio Grande do Sul, Brazil. Data from 2003 and 2004 are unavailable.

exception is the gap between secondary and tertiary education (workers that had completed college and beyond), which rose until 2006 and subsequently remained stable through to the end of the period studied. In Appendix D.2, I show that wage inequality trends are similar in a different data set that includes informal workers.

Among potential causes behind falling inequality in Brazil, the most conspicuous are increased educational achievement and rising minimum wages. Figure 6 shows that both factors are relevant in Rio Grande do Sul. The first graph displays the fraction of hours worked by employees in each educational group. The pattern is striking: workers with less than a complete primary education (that is, less than eight years of schooling) supply 40 percent of work hours in 1998, but only around 15 percent in 2012. On the other hand, the group with a complete secondary education (high school and college dropouts) increased its participation level by almost 30 percentage points. Moreover, there is a substantial increase in college

Figure 6: Changes in educational achievement and minimum wages



Notes: RAIS data for the formal sector in Rio Grande do Sul, Brazil. Data from 2003 and 2004 are unavailable.

completion in relative terms (from 9.4 percent to 12.2 percent), though they remain a small fraction of the formal workforce.¹⁵

The bottom graph in Figure 6 shows the large and steady increase in the national minimum wage in Brazil. The same figure shows that the minimum wage increased much faster than median wages in Rio Grande do Sul until 2006. These increases point to an important role of the minimum wage, consistent with the fact that lower-tail inequality fell more than upper-tail inequality,

Next, I use the panel structure of the matched employer-employee data to gauge the importance of cross-firm wage differentials and labor market sorting in Rio Grande do Sul. A widely used approach in this literature is a variance decompositions based on AKM regres-

¹⁵Figure D3 in Appendix D.2 shows similar trends in the share of all adults belonging to each of these educational groups, regardless of whether they participate in the labor force or not. That is consistent with the changes in employment shares by educational group being caused by government policies and development, instead of changes in selection patterns into employment.

sions. Let the log wage of worker i at time t be written as:

$$\log y_{it} = v_i + \psi_{J(i,t)} + \delta_t + u_{it}$$

where v_i is worker i 's fixed effect, ψ_j is establishment j 's fixed effect, $J(i,t)$ represents the establishment employing worker i at time t , δ_t is a time effect, and u_{it} is a residual that is uncorrelated with all fixed effects. Based on this model, the total variance of log wages can be decomposed in the following way:

$$\text{Var}(\log y_{it}) = \text{Var}(v_i) + \text{Var}(\psi_{J(i,t)}) + 2\text{Cov}(v_i, \psi_{J(i,t)}) + \text{Var}(\delta_t) + 2\text{Cov}(v_i + \psi_{J(i,t)}, \delta_t) + \text{Var}(u_{it}) \quad (12)$$

If wages vary substantially across establishments for similar workers, the variance of establishment effects will be large, adding to overall wage dispersion. If high-wage workers are more likely to work at high-wage establishments, then the first covariance term will be positive, further increasing overall inequality. Based on this logic, the correlation between establishment and worker fixed effects has been used as a dimensionless scalar measure of sorting in the labor market.

Estimating the variance decomposition (12) is not trivial. I use the method and code of [Kline, Saggio and Sølvesten \(2018\)](#) (henceforth KSS), which is not subject to the limited mobility bias discussed by [Andrews et al. \(2008\)](#). Their estimator also provides standard errors for variance decomposition components. Appendix D.3 provides details about the procedure.

Table 1 reports the KSS estimates of variance components. The variance of both establishment and worker effects decline over time, helping to explain the fall of wage inequality. There is also a sizable and statistically significant increase in the covariance of worker effects and establishment effects. The labor market is becoming more assortative, a force that works against the overall trend of falling inequality.¹⁶

The interpretability of AKM decompositions relies on establishments being categorized as high- or low-wage. But in my model (and many other structural models of sorting), wages are not log-additive in worker and establishment components: some establishments may pay relatively more to some worker types and less to others. Still, if the model is able to replicate

¹⁶[Alvarez et al. \(2018\)](#) previously estimated AKM decompositions using Brazilian data. They used the whole sample instead of a single state and did not correct for limited mobility bias. The only significant difference in results is that the correlation between worker and firm fixed effects is roughly constant in their decompositions, such that the covariance term falls (in line with the declines in the variances of worker and firm effects).

Table 1: Variance decomposition from two-way fixed effects model

Component	1998		2012		Change	
$\text{Var}(\log y_{it})$	0.675	100%	0.544	100%	-0.131	100%
$\text{Var}(v_i)$ (worker effects)	0.391 (0.003)	57.9%	0.324 (0.001)	59.5%	-0.067 (0.003)	51.1%
$\text{Var}(\psi_{J(i,t)})$ (establishment effects)	0.139 (0.002)	20.6%	0.071 (0.001)	13.1%	-0.067 (0.003)	51.5%
$2\text{Cov}(v_i, \psi_{J(i,t)})$	0.073 (0.003)	10.8%	0.103 (0.001)	18.9%	0.030 (0.003)	-22.6%
Other terms	0.072	10.7%	0.046	8.5%	-0.026	20.0%
$\text{Corr}(v_i, \psi_{J(i,t)})$	0.157		0.337		0.180	

Notes: This table shows the variance decomposition (12) estimated using [Kline, Saggio and Sølvssten \(2018\)](#). Numbers in parentheses are asymptotic standard errors. Columns labeled 1998 use data from 1997 and 1999. The 2012 specification uses data from 2011 and 2013. Each worker-year observation, including both stayers and movers, has the same weight in the decomposition. See [Appendix D.3](#) for details and sample sizes.

the data well enough, indirect inference can be used extract identifying information from the AKM decomposition. This is the strategy I employ in this paper.

The quantitative exercise focuses on the following features of the data: differences in mean log wages between the four educational groups; variances of log wages within each of the four groups; shares of hours worked by each educational group; levels of the minimum wage relative to the mean wage; and variances of establishment fixed effects and covariance between worker and establishment fixed effects from the AKM decomposition. These moments offer a broad picture of wage inequality, the supply of skills, how binding the minimum wage is, and the role of cross-firm wage differentials. In the estimation procedure described below, I will attempt to match those moments for two different periods: circa 1998 (formally, the years 1997 and 1999) and circa 2012 (2011 and 2013). Other aspects of the data will be used for model validation.

4.2 Parameterization and identification

I propose a simple empirical specification that can capture the rich patterns in the data and generate meaningful, credible counterfactuals. This section describes that parameterization and which features of the data allow for identification of the estimated parameters.

4.2.1 Workers

I use four worker types ($H = 4$), which will be linked to the four educational groups in the data. That will allow the model to capture "wage polarization" across educational groups, whereby high-school workers lose relative to both less and more educated workers. Having two low-education groups helps with capturing nuanced labor supply changes, the role of the minimum wage, and wage compression in the lower tail. Workforce composition parameters $N_{h,t}$ will be identified by matching the share of hours worked by each group.

Efficiency units of labor are normally distributed within each worker group, with mean zero. The standard deviation is parameterized as:

$$S_{h,t} = S_{h,1998} \exp(\mathbf{1}\{t = 2012\} \hat{S})$$

The type-specific parameters $S_{h,1998}$ are identified from overall within-group variances of log wages over both periods. The \hat{S} parameter is identified from overall changes in within-group inequality between periods. This parameter is one of two structural trends in the model intended to capture unmodeled factors that might have explained falling wage inequality in Brazil.

In the data, within-group inequality falls more for some groups than for others. The empirical model is not hardwired to capture those differential reductions, since the structural residual \hat{S} applies equally to all groups. However, they can be plausibly explained by changes in cross-firm wage differentials and in the minimum wage. Thus, within-group variances of log wages offer a total of three overidentifying restrictions that help identify the rest of the model, as described below.

4.2.2 Preferences and goods

The preference parameter β , which determines the slope of the firm-level labor supply curve, is set to 4. That value implies a markdown of wages relative to marginal revenue products of labor of 20% for skilled workers, and slightly less for less skilled workers (since a binding minimum wage makes the labor supply curve more elastic). That value is considered a "reasonable near-competitive benchmark" by [Card et al. \(2018\)](#) based on their review of the literature. It is also close to the value of 4.99 estimated by [Lamadon, Mogstad and Setzler \(2019\)](#). I choose not to estimate β because I would not have a design that is as credible as the best ones in that literature.

I set $G = 2$, the minimum number of goods such that the model can generate sorting and cross-firm wage differentials. Because I am not interested in predictions about establishment sizes, entry costs F_g cannot be separately identified from mean workplace amenities by good, \bar{a}_g . Moreover, since there is no cardinal meaning to amenities or the entry input, only relative differences in entry costs and amenities have empirical content. Thus, I define the cross-firm labor cost wedge as:

$$\Delta_t = \log \left(\frac{F_{2,t}}{\bar{a}_{2,t}} \right) - \log \left(\frac{F_{1,t}}{\bar{a}_{1,t}} \right)$$

From Proposition 2, we know that this parameter is strongly linked to the overall size of wage differences across firms producing the different goods, driven either by differences in firm sizes (representing movements along the labor supply curve) or compensating differentials (vertical translations of the curves). I allow this parameter to change over time, capturing changes in either of these factors. They will be identified from the variance of establishment fixed effects in each period, coming from the AKM decompositions.¹⁷

The relative taste shifter $\gamma_{2,t}/\gamma_{1,t}$ is allowed to vary over time. This serves two purposes. First, it provides the second structural trend capturing the effect of alternative explanations for the fall of inequality in Brazil. It might, for instance, absorb the effect of trade liberalization and commodity shocks, factors that have been documented as relevant in Brazil (see, e.g., the review in [Firpo and Portella, 2019](#)).

In addition, having a flexible demand shifter allows me to be agnostic about the elasticity of substitution σ at the estimation stage. That's because changes in σ can be counteracted by changes in $\gamma_{2,t}/\gamma_{1,t}$, leaving the allocation of inputs and consumption goods untouched. Thus, for estimation purposes, I set $\sigma = 1$. However, different values of σ will imply different counterfactual predictions, as discussed in Subsection 4.5. The demand shifter is identified jointly with the technology parameters, as described below.

4.2.3 Technology

I use the exponential-Gamma parameterization of the task-based production function, described in Subsection 2.4. The productivity parameters are not separately identified from the relative demand shifters $\gamma_{2,t}/\gamma_{1,t}$. Changes in z_g affect the allocation of labor across firms and tasks if $\sigma \neq 1$. But the ensuing effects in wages and sorting can be offset by corresponding

¹⁷In the numerical implementation of the model, I set $\bar{a}_1 = \bar{a}_2 = F_2 = T = 1$ and $F_{1,t} = \exp(-\Delta_t)$. Different normalizations would only change the nominal level of prices in the model (including the estimated minimum wage).

changes in relative demand for goods. I thus normalize all physical productivities z_g to one.¹⁸

I make two further restrictions on technology parameters. First, the blueprint shape parameter k_g is assumed to be the same for all goods. Second, all technology parameters are stable over time, except for a common change in blueprint scale for all goods:

$$\theta_{g,t} = \theta_{g,1998} \exp(\mathbf{1}\{t = 2012\} \hat{\theta})$$

The shock $\hat{\theta}$ represents a change in the skill bias of technologies. If positive, it increases the demand for high-complexity tasks in the production of all goods, making skilled workers relatively more productive. It can also be negative, in which case it would cause a reduction in between-group inequality.

There are eight technology and taste parameters left to be identified: α_2 , α_3 , $\theta_{1,1998}$, $\theta_{2,1998}$, $\hat{\theta}$, k , $\gamma_{2,1998}/\gamma_{1,1998}$, and $\gamma_{2,2012}/\gamma_{1,2012}$. They will be jointly identified from a total of 11 moments: six between-group gaps in mean log wages, two covariances between establishment and worker effects from the AKM decomposition, and three relative changes in within-group variances of log wages (the over-identifying restrictions from Subsection 4.2.1). The technology and taste parameters have implications for all of these moments, and their effects are different enough to satisfy a rank condition for identification.¹⁹

4.3 Estimation procedure

I estimate the model via optimal minimum distance. For each candidate vector of parameters ϕ , I solve for equilibria for both periods (circa 1998 and circa 2012) and simulate the corresponding moments. The estimated parameters are the ones that minimize $(M(\phi) - \hat{M})' \hat{V}^{-1} (M(\phi) - \hat{M})$, where $M(\phi)$ is the vector of simulated moments implied by ϕ , \hat{M} is the vector of estimated moments from the data, and \hat{V} is the estimated covariance matrix associated with those moments.

To reduce the dimensionality of the minimization problem, labor supply parameters $N_{h,t}$ are not included in ϕ . Correspondingly, observed employment shares by educational group are

¹⁸I would need to observe prices for goods to separately identify these supply and demand shocks.

¹⁹For example, it would be difficult to identify biased technical change $\hat{\theta}$ separately from relative demand for the complex good in 2012, $\gamma_{2,2012}/\gamma_{1,2012}$, using only data on between-group inequality. That's because both parameters shift the economy-wide task demand towards more complexity. But the demand shock also causes a substantial reduction in within-group inequality for college-educated workers in the neighborhood around the estimated parameters, as these workers become more concentrated in the complex good.

not included in \hat{M} . Instead, I use a modified equilibrium condition in estimation: employment in the model must match observed employment in the data, rather than the total number of employed and unemployed workers matching labor supply. That improves computational efficiency at the cost of not taking into account sampling variation in the employment shares. Since that sampling variation is extremely small given the size of the data set, the quantitative implications are negligible.

The simulated variance decomposition is based on a regression of model-predicted log wages on worker type dummies (representing pairs of h, ε) and firm type dummies (g). In other words, I find the best approximation for wages in the model under the assumption that firms pay a common premium to all workers, the core assumption behind the AKM decomposition. This approach is much more computationally efficient than simulating panel data at the individual-firm level and then running the KSS estimator on that data. It also precludes the need for assumptions about how job-to-job transition patterns respond to changes in model parameters. Appendix D.4 presents assumptions under which both approaches are equivalent.

I use a large number of starting points randomly drawn from a broad region of the parametric space to increase the likelihood of finding a global optimum. Most starting points converge to the estimated parameters, suggesting that the objective function is well behaved. Implementation details are laid out in Appendix D.4.

4.4 Estimated model and goodness of fit

The model fits the data well. Table 2 shows that all of the simulated moments are within 0.01 of the targets in the data, with the exception of the covariance of worker and establishment effects. There, the model still captures the qualitative features of the data. The simulated covariance is positive, and labor market sorting — measured by the correlation of worker and establishment fixed effects — is increasing from 1998 to 2012. However, the degree of assortativeness is higher in the data (from 0.157 to 0.337) than in the model (from 0.102 to 0.120).

The fit of the model can be verified visually by comparing observed wage distributions with model-generated ones in Appendix Figure D4. The model captures the most salient features of the data, though it over-predicts bunching and the share of workers close to the minimum wage.

Table 2: Target moments and model fit

Target	Data		Model	
	1998	2012	1998	2012
<i>Mean log wage gaps:</i>				
Primary / No degree	0.139 (0.001)	0.041 (0.001)	0.140	0.041
Secondary / Primary	0.381 (0.002)	0.156 (0.001)	0.379	0.155
Tertiary / Secondary	0.701 (0.002)	0.984 (0.001)	0.703	0.993
<i>Within-group variance of log wages:</i>				
No degree	0.352 (0.001)	0.188 (0.001)	0.348	0.189
Primary	0.458 (0.001)	0.237 (0.001)	0.466	0.233
Secondary	0.675 (0.001)	0.348 (0.001)	0.668	0.347
Tertiary	0.813 (0.002)	0.634 (0.001)	0.815	0.633
<i>Log min. wage - mean log wage</i>	-1.352 (0.001)	-0.942 (0.000)	-1.352	-0.942
<i>Variance decomposition from two-way fixed effects model:</i>				
Variance of establishment effects	0.139 (0.002)	0.071 (0.001)	0.142	0.078
Cov. worker and establishment effects	0.037 (0.002)	0.052 (0.001)	0.024	0.019

Notes: Standard errors in parentheses.

Table 3 shows the estimated parameters and their standard errors. Overall, parameters are precisely estimated. Comparative advantage is increasing in education, as expected. The dispersion of efficiency units of labor is higher for more educated workers and falls over time. Relative to good $g = 1$, good $g = 2$ requires more complex tasks in production, is deemed as more desirable by consumers, and has substantially higher entry costs to mean amenity ratios in both periods. But differences in consumer preferences and in the cross-firm labor cost wedge become smaller from 1998 to 2012. Finally, technical change is skill-biased. That result is consistent with a micro-level study of the impact of digital technologies in Brazil (Almeida, Corseuil and Poole, 2018).

Figure 7 shows employment patterns and posted wages (per efficiency unit of labor) in the estimated equilibria. Firms producing the second good are more skill intensive in both pe-

Table 3: Estimated parameters

Parameter	Value	S.E.
S_1 (st. dev. of eff. units, no degree)	0.557	(0.003)
S_2 (, primary)	0.615	(0.002)
S_3 (, secondary)	0.692	(0.003)
S_4 (, tertiary)	0.824	(0.003)
\hat{S} (log change in S , all workers)	-0.199	(0.004)
α_2 (comparative advantage, primary)	-0.599	(0.015)
α_3 (, secondary)	-0.247	(0.008)
$\theta_{1,1998}$ (blueprint complexity, $g = 1$)	0.016	(0.009)
$\theta_{2,1998}$ (, $g = 2$)	1.279	(0.243)
$\hat{\theta}$ (bias of technical change)	0.923	(0.030)
k (blueprint shape)	1.121	(0.158)
$\log \gamma_{2,1998} / \gamma_{1,1998}$ (product demand shifter)	0.801	(0.028)
$\log \gamma_{2,2012} / \gamma_{1,2012}$	0.221	(0.017)
Δ_{1998} (cross-firm labor cost wedge)	4.451	(0.025)
Δ_{2012}	3.979	(0.026)
$\log y_{1998}$ (min. wage to entry input)	-0.336	(0.001)
$\log y_{2012}$	0.161	(0.001)

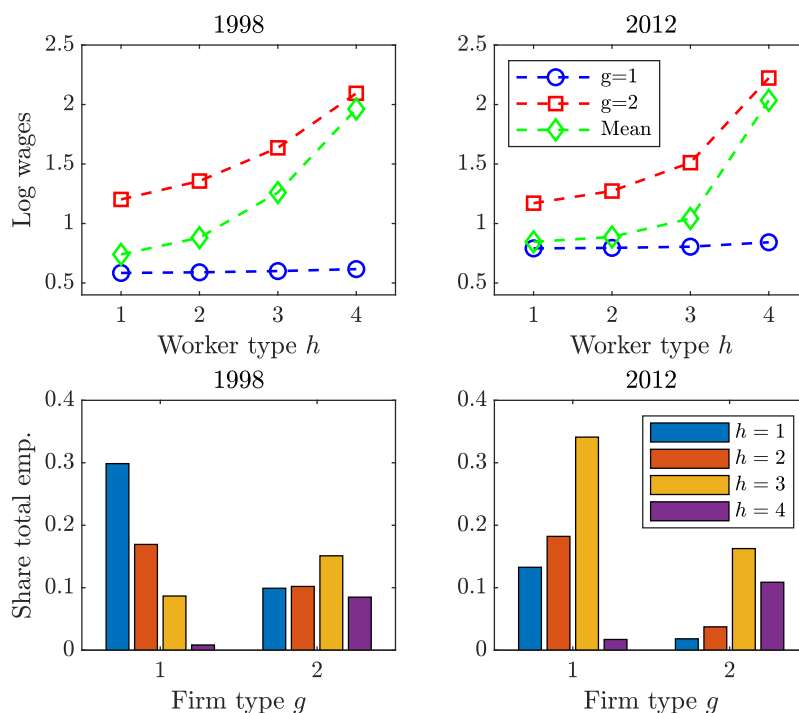
riods. In the first period, firms producing $g = 2$ pay higher wages to all worker types, but the premiums are higher for more skilled workers. The slopes differ because of demand for skills, while the vertical difference between the curves is due to the cost wedge Δ_t (stemming from differences in entry costs or mean amenities).

In the second period, firms producing $g = 2$ still pay more to all workers. But the vertical distance between the wage curves become smaller due to the estimated reduction in the cross-firm labor cost wedge. At the same time, firms producing good $g = 2$ become relatively more intensive in college-educated workers. The combination of these effects explains why firm fixed effects become less relevant as a share of the total variance of log wages, while at the same time the correlation between worker and firm fixed effects increases.

4.5 The role of supply, demand, minimum wages, and other shocks

In this section, I perform a counterfactual analysis to isolate the individual impact of each of the six time-varying factors in the model: labor supply, blueprint complexities, the minimum wage, the cross-firm labor cost wedge, and the two trends capturing alternative explanations for the fall of wage inequality. I split the path from 1998 to 2012 into 300 steps. In each step, I change one of the time-varying parameters by 1/50th of the total change and solve

Figure 7: Wages and employment by firm type and worker type



Notes: The top panel shows log wages per efficiency unit of labor for each worker type h posted by firms producing good g , $\log w_{hg}$. It also shows mean log earnings for each worker type h . The bottom panel shows the distribution of employment in the economy between worker types and firm types. For each year, all bars sum to one.

for the counterfactual equilibrium. Then, I calculate the implied change in each outcome variable of interest relative to the previous step and store the results. The total contribution of each factor is the sum across the 50 steps where the corresponding parameter(s) has(have) changed.²⁰

In this exercise, I need to take a stance on the elasticity of substitution between goods. Table 4 shows decompositions for three cases: Cobb-Douglas ($\sigma = 1$), high substitutability ($\sigma = 10$), and low substitutability ($\sigma = 0.1$).²¹

²⁰I use small changes for two reasons. First, it ensures that the model is always close to a parameter set that is a linear combination of 1998 and 2012 estimated parameters. This is useful because model predictions far estimated parameters might be unrealistic and unreliable. Second, it makes the choice of ordering of the shocks irrelevant.

²¹As discussed in Subsection 4.2, the baseline estimation of the taste parameters uses $\sigma = 1$. When a different value of σ is used, the taste parameters are adjusted in the following way. First, γ_2 is set to 1 in all exercises. Letting γ_1 denote the estimated value with $\sigma = 1$ considering that normalization, the adjusted value for an alternative σ' is $\log(\gamma'_1) = \log(\gamma_1)/\sigma' + (1 - 1/\sigma') \log(p_1/p_2)$, where p_g denote prices for goods in the

Table 4: Model-based decomposition of trends in wage inequality and sorting

Outcome	All changes	Labor supply	Tech. change	Min. wage	Δ wedge	Taste shock	Var. of eff. units
<i>Panel A: Cobb-Douglas demand for goods ($\sigma = 1$)</i>							
Primary / No degree	-0.099	-0.058	-0.004	0.006	-0.003	-0.038	-0.003
Secondary / Primary	-0.224	-0.201	0.077	-0.001	-0.004	-0.102	0.007
Tertiary / Secondary	0.291	-0.064	0.465	-0.028	-0.022	-0.093	0.033
Var. log earnings	-0.167	-0.014	0.071	-0.080	-0.021	-0.045	-0.078
Upper tail: p90/p50	-0.098	-0.035	0.093	-0.036	-0.030	-0.044	-0.046
Lower tail: p50/p10	-0.179	0.023	-0.030	-0.144	-0.002	-0.011	-0.015
AKM sorting	0.019	-0.083	0.120	-0.052	0.010	-0.042	0.066
<i>Panel B: High substitution ($\sigma = 10$)</i>							
Primary / No degree	-0.099	-0.010	-0.081	0.009	0.020	-0.029	-0.007
Secondary / Primary	-0.224	-0.080	-0.124	0.007	0.055	-0.077	-0.005
Tertiary / Secondary	0.291	0.028	0.297	-0.019	0.027	-0.065	0.023
Var. log earnings	-0.167	0.037	-0.016	-0.076	0.005	-0.034	-0.083
Upper tail: p90/p50	-0.098	0.012	0.011	-0.032	-0.006	-0.032	-0.050
Lower tail: p50/p10	-0.179	0.040	-0.056	-0.143	0.005	-0.009	-0.016
AKM sorting	0.019	-0.036	0.040	-0.048	0.033	-0.031	0.061
<i>Panel C: Low substitution ($\sigma = 0.1$)</i>							
Primary / No degree	-0.099	-0.070	0.015	0.006	-0.008	-0.040	-0.002
Secondary / Primary	-0.224	-0.230	0.127	-0.004	-0.018	-0.108	0.009
Tertiary / Secondary	0.291	-0.089	0.510	-0.030	-0.035	-0.101	0.036
Var. log earnings	-0.167	-0.026	0.093	-0.081	-0.027	-0.048	-0.077
Upper tail: p90/p50	-0.098	-0.047	0.114	-0.037	-0.036	-0.047	-0.044
Lower tail: p50/p10	-0.179	0.019	-0.024	-0.145	-0.004	-0.011	-0.015
AKM sorting	0.019	-0.095	0.140	-0.053	0.004	-0.044	0.076

Notes: *All changes* denotes changes predicted by the estimated model, comparing the 1998 equilibrium to the 2012 equilibrium. Each of the other columns represent the contributions of the six time-varying factors in the model, which add up to the total change. The six factors are (in the order they appear): workforce composition along education groups, the skill bias of technology, the minimum wage, the difference across goods in the ratio of entry costs to mean amenities, the relative weight of each good in consumption, and the variance of efficiency units of labor among workers.

Many of the results are robust to the choice of elasticity of substitution across goods. For example, the minimum wage is consistently found to reduce the total variance of log wages by a little less than half of the total decline. That makes it the most important factor according to that metric, along with the reduction in the dispersion of efficiency units of labor. The minimum wage has particularly strong effects in reducing the p50/p10 percentile gap, but it has limited influence in between-group inequality measures. Finally, the minimum wage reduces the correlation between establishment and worker fixed effects in the AKM decomposition, regardless of σ .

Some of the effects of SBTC are also robust to different specifications of σ . It is by far the most important driver of the rising college premium. In addition, it always contributes to AKM sorting, being the most important positive contribution when $\sigma = 1$ and $\sigma = 0.1$.

However, the impact of SBTC on the variance of log wages changes from positive to negative when σ becomes large. To understand why, note that good $g = 1$ is estimated to be very intensive in low-complexity goods, with a mean task complexity of 0.018 (compared to 1.434 for good 2). Thus, the SBTC shock (a proportional increase in task complexities) makes $g = 2$ more costly, but has minimal effects on $g = 1$. With $\sigma = 0.1$ the net effect is mechanical: a shift towards more complex tasks in the whole economy, widening all between-group wage gaps and overall inequality. But when $\sigma = 10$, consumption shifts toward good 1. There is still space for college-educated workers at the skill-intensive firms. However, high school workers who would be employed in skill-intensive firms before the shock end up at low-skill firms, losing their firm premiums. That explains the reductions in some between-group wage gaps and in the variance of log wages.

The effects of the supply shock on inequality also depend on the specification of σ . With low substitution, the increase in schooling achievement reduces between-group wage gaps. The effects on the variance of log earnings are negative, but small in magnitude. That's because the shock moves workers towards groups with larger variance of efficiency units and for whom cross-firm wage differentials are larger. When the elasticity of substitution is high, the supply shock stimulates production of good $g = 2$, offsetting the reductions in between-group inequality (and reverting them in the case of the college premium). The total effect on the variance of log wages becomes positive. These results are consistent with the theoretical discussion in Subsection 3.4.

Rising schooling achievement makes the economy less assortative. In the initial period, there

estimated equilibrium.

are very few highly skilled workers, almost all of which are employed by firms producing good 2. As the supply of these workers increases, firms become less segregated by skill, reducing the degree of sorting in the economy.

The reduction of the cross-firm labor cost wedge has modest effects on inequality. The strongest and most consistent result is a reduction in upper-tail wage inequality. That happens because cross-firm wage differentials are most relevant to skilled workers.

Finally, the two trends intended to capture unmodeled factors help explain falling inequality in Rio Grande do Sul. I find that the taste shock benefits the low-skill good. That is consistent with trade shocks affecting prices for goods in a way that benefits low-skilled workers, a view supported by the summary of the literature in [Firpo and Portella \(2019\)](#).

The sizable reduction in the dispersion of efficiency units of labor can be a side effect of mapping comparative advantage groups in the model to educational groups in the data. If skill groups in the model were latent, imperfectly correlated with observed education, then reductions in between-group inequality would also cause a reduction in within-group inequality. That would reduce the scope for the residual factor \hat{S} . I opted to map types directly to observables to make the exercise simpler and more transparent, in the tradition of classic supply-demand studies such as [Katz and Murphy \(1992\)](#).

I finish this discussion with a caveat on the interpretation of these results. Because labor supply to the formal sector is exogenous in the model, the impact of other factors should be interpreted as net of potential labor supply responses. On the long run, shocks such as the minimum wage might affect labor force participation of some educational groups more than others. In addition, the informal sector provides an "outside option" for workers in the formal sector, and the gap in attractiveness between formal and informal jobs might be a complicated function of the workforce composition, technical change, and minimum wages ([Haanwinckel and Soares, 2020](#)). A model-based account of the informal sector and endogenous labor supply decisions is possible, but beyond the scope of this paper.

4.6 Model validation

In the last part of the quantitative section, I use additional data to provide support for the theory. I test predictions of the assignment structure and verify if the effects of the minimum wage in the model are consistent with the data.

Table 5: Validation of the task-based production function.

	(1)	(2)	(3)	(4)
Own schooling	0.0796 (0.0101)			
Mean schooling in establishment		0.0693 (0.000767)		
Mean schooling of coworkers in establishment			0.0147 (0.00118)	0.00862 (0.00131)
Unit of observation	Worker	Firm	Worker	Worker
Firm fixed effects	✓			
Worker fixed effects, sample of movers			✓	✓
Sector fixed effects				✓
r ²	0.525	0.239	0.819	0.836
N	947435	1013467	402540	402494

Notes: RAIS data for Rio Grande do Sul state. Columns (1) and (2) use data from 1997. Columns (3) and (4) use data from 1997 and 1999. The dependent variable is the analytical non-routine task content of the occupation (averaged across workers employed by the establishment in Column (2)). Standard errors (in parenthesis) are robust in Column (2) and two-way clustered at the worker and firm levels in columns (1), (3), and (4). The standard deviation of the dependent variable is approximately one.

4.6.1 Task content of occupations and the assignment model

In the model, worker types are assigned to specific ranges of tasks depending on their type, which I tie to schooling in the empirical exercise. The assignment rules may differ across firms when labor markets are not competitive. Proposition 2 shows that, with upward-sloping firm-specific labor supply curves, firms that use skilled workers more intensely pay relatively more to those workers. From Lemma 1, we can conclude that task thresholds that define assignment in more skill-intensive firms are all to the right of the thresholds in less skill-intensive ones. Thus, when workers move from low- to high-skill firms in the model, they will on average be allocated to more complex tasks.

To test those predictions, I need a proxy for task complexity (i.e., how much the job benefits from formal education). I use a measure of the nonroutine analytical task content of Brazilian occupations created by de Sousa (2020). That measure reflects whether O*NET survey respondents believe that the occupation requires mathematical reasoning, and was created following the methodology in Deming (2017).²²

²²The O*NET is a survey that asks workers in the US about their jobs, including skill requirements, the types of activities performed, and the degree of automation in the occupation. Deming (2017) describes how that survey is collected and processed, resulting in data that describes each occupation as a combination of tasks in

To start, I regress the nonroutine analytical task content of the worker’s occupation on his own schooling (defined in years), controlling for establishment fixed effects. That regression tests if, within establishments, more educated workers are assigned to more complex tasks. Results are reported in the first column of Table 5. The correlation is positive and precisely identified.

Next, I regress mean analytical task content on mean schooling at the establishment level. Again, the results—reported in column (2)—are positive, consistent with differences in task requirement generating cross-firm heterogeneity in skill intensity.

I use panel data to investigate changes in assignment following worker transitions across establishments. Specifically, I regress the analytical task content of the worker’s occupation on mean schooling among all other workers in the same establishment, while controlling for worker fixed effects. That regression uses data from 1997 and 1999. Column (3) shows that the estimate is positive and significant, though the correlation is weaker than in Column (2). Workers who move to firms with more educated colleagues tend to be assigned to more analytical occupations, consistent with differences in optimal assignment across firms in imperfectly competitive environments.

I also investigate whether changes in assignment are driven by workers moving across sectors. Column (4) shows results for a specification similar to that in (3), but with sector fixed effects.²³ I find that the coefficient falls by about half, but remains highly significant. That suggests sizable within-sector variation in skill intensity and task content of occupations, consistent with the interpretation that goods in the model might represent differentiated varieties or technologies within industries.

4.6.2 Unemployment effects of the minimum wage and spillovers

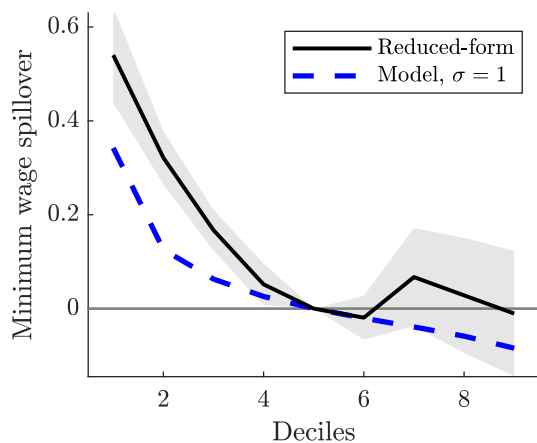
This subsection tests whether impact of the minimum wage in the model is consistent with reduced form findings. I start with employment effects. A large empirical literature documents that minimum wages have small or negligible employment effects.²⁴ The monopsony model is consistent with these results. Using parameters as of 2012, the elasticity of employ-

varying intensities. [de Sousa \(2020\)](#) builds a crosswalk between SOC occupation codes and occupations codes in the RAIS data, and then calculates the task content of occupations using O*NET data and the procedures in [Deming \(2017\)](#).

²³CNAE10 sectors. There are 544 sectors in the regression sample, with 187 of them containing at least 400 observations (0.1% of the sample).

²⁴[Cengiz et al. \(2019\)](#) and [Dustmann et al. \(2020\)](#) are recent examples studying the US and Germany, respectively.

Figure 8: Minimum wage spillovers



Notes: This figure shows the impact of changes in the "effective minimum" (statutory minimum wage minus median wage) on different deciles of the wage distribution, relative to the median. The shaded area represents 95 percent confidence intervals associated with reduced-form estimates.

ment with respect to changes in the effective minimum wage (minimum wage minus median wage) is only -0.028 (-0.042 for workers in groups $h = 1$ and $h = 2$). With parameters as of 1998, the elasticities are even smaller in magnitude, because the minimum wage was less binding at that time.

Next, I verify if reduced-form measures of minimum wage spillovers in Brazil are similar to spillovers implied by the model. I use the procedures in [Autor, Manning and Smith \(2016\)](#) (henceforth AMS) to obtain reduced-form estimates using data from all Brazilian states. The empirical model measures how changes in the "effective minimum wage" — that is, the minimum wage minus the median wage in state-year — affects different quantiles of the wage distribution, relative to the median. [Appendix D.6](#) describes the empirical model, the identification strategy using instrumental variables, and reports regression results. It also explains how model-based spillovers are simulated.

[Figure 8](#) shows that model-based spillovers, measured consistently with the AMS specification, are in the same order of magnitude compared to the reduced-form ones, but significantly smaller in the lower tail. This does not imply that spillovers in the model are weaker than in the data. AMS normalize wages relative to the median, under the assumption that those workers are not likely to be affected by the minimum wage. This stance is sensible for the US but not necessarily for Brazil. If wages were renormalized relative to, say, percentiles 70, 80, or 90, model-based spillovers would be similar to reduced form ones in the lower tail,

and larger for median workers. These differences might arise from misspecification in the distribution of efficiency units of labor (generating a simulated wage distribution with more mass close to the minimum, as shown in Appendix Figure D4). Alternatively, the reduced-form estimates, which use data for the whole country, might not provide a perfect measure of spillovers in Rio Grande do Sul state.²⁵

5 Conclusion

I showed that the task-based production function is a tractable, intuitive, and parsimonious way to model cross-firm differences in employment patterns and SBTC. Adding capital to the theory is a potential path for further research. In the model of routine-biased technical change in [Acemoglu and Autor \(2011\)](#), machines are particularly cost-effective on tasks executed by middle-skill workers. That idea can be generalized in my framework by modeling vintages of capital as inputs similar to labor, with different productivities at each task. The combination of firm heterogeneity, product demand, and imperfect competition could offer insights beyond those in [Acemoglu and Autor \(2011\)](#). For example, what determines whether job polarization happens within firms or between firms, and does that matter for the wage distribution?

Embedding that production function in a general equilibrium, imperfectly competitive model revealed novel results, such as a technological shock making the labor market in my setting more assortative. But it also raised new questions. I found that the cross-firm wedge in labor costs shrank over time, reducing the dispersion of firm wage premiums. It is not clear what caused that reduction. Changes in technology, market structure, or regulations may have reduced entry costs for skill-intensive firms. Targeted studies of particular regions or industries, using additional data sources, may shed light on that point.

²⁵ Simulated spillovers in Figure 8 correspond to $\sigma = 1$. Values are almost identical for $\sigma = 0.1$ and $\sigma = 10$. That is because the minimum wage does not introduce large changes in costs across goods. The minimum wage does lead to different responses in wage posting and employment between firms, but those details are omitted for brevity.

References

- Abowd, John M, Francis Kramarz, and David N Margolis.** 1999. “High Wage Workers and High Wage Firms.” *Econometrica*, 67(2): 251–333.
- Acemoglu, Daron.** 1998. “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality.” *The Quarterly Journal of Economics*, 113(4): 1055–1089.
- Acemoglu, Daron.** 1999. “Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence.” *The American Economic Review*, 89(5): 1259–1278.
- Acemoglu, Daron.** 2001. “Good Jobs versus Bad Jobs.” *Journal of Labor Economics*, 19(1): 1–21.
- Acemoglu, Daron, and David Autor.** 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” In *Handbook of Labor Economics*. Vol. 4, Part B, , ed. David Card and Orley Ashenfelter, Chapter 12, 1043–1171. Elsevier.
- Acemoglu, Daron, and Pascual Restrepo.** 2018. “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review*, 108(6): 1488–1542.
- Acosta, Pablo, Guillermo Cruces, Sebastian Galiani, and Leonardo Gasparini.** 2019. “Educational upgrading and returns to skills in Latin America: evidence from a supply–demand framework.” *Latin American Economic Review*, 28(18).
- Ales, Laurence, Musab Kurnaz, and Christopher Sleet.** 2015. “Technical Change, Wage Inequality, and Taxes.” *American Economic Review*, 105(10): 3061–3101.
- Almeida, Rita, Carlos Henrique Leite Corseuil, and Jennifer Poole.** 2018. “The Impact of Digital Technologies on Worker Tasks: Do Labor Policies Matter?” CESifo Working Paper 6798.
- Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser.** 2018. “Firms and the Decline in Earnings Inequality in Brazil.” *American Economic Journal: Macroeconomics*, 10(1): 149–89.
- Andrews, M. J., L. Gill, T. Schank, and R. Upward.** 2008. “High wage workers and low wage firms: negative assortative matching or limited mobility bias?” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171(3): 673–697.

- Antràs, Pol, Luis Garicano, and Esteban Rossi-Hansberg.** 2006. “Offshoring in a Knowledge Economy.” *The Quarterly Journal of Economics*, 121(1): 31–77.
- Autor, David H., Alan Manning, and Christopher L. Smith.** 2016. “The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment.” *American Economic Journal: Applied Economics*, 8(1): 58–99.
- Autor, David H, Lawrence F Katz, and Melissa S Kearney.** 2008. “Trends in U.S. Wage Inequality: Revising the Revisionists.” *The Review of Economics and Statistics*, 90(2): 300–323.
- Blundell, Richard, David Green, and Wenchao Jin.** 2020. “The UK as a Technological Follower: Higher Education Expansion, Technological Adoption, and the Labour Market.” Unpublished manuscript.
- Bound, John, and George Johnson.** 1992. “Changes in the Structure of Wages in the 1980’s: An Evaluation of Alternative Explanations.” *The American Economic Review*, 82(3): 371–392.
- Butcher, Tim, Richard Dickens, and Alan Manning.** 2012. “Minimum wages and wage inequality: some theory and an application to the UK.” Centre for Economic Performance Discussion Papers 1177.
- Caliendo, Lorenzo, and Esteban Rossi-Hansberg.** 2012. “The Impact of Trade on Organization and Productivity.” *The Quarterly Journal of Economics*, 127(3): 1393–1467.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline.** 2018. “Firms and Labor Market Inequality: Evidence and Some Theory.” *Journal of Labor Economics*, 36(S1): S13–S70.
- Card, David, Jörg Heining, and Patrick Kline.** 2013. “Workplace Heterogeneity and the Rise of West German Wage Inequality.” *The Quarterly Journal of Economics*, 128(3): 967.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer.** 2019. “The Effect of Minimum Wages on Low-Wage Jobs.” *The Quarterly Journal of Economics*, 134(3): 1405–1454.
- Coase, R H.** 1937. “The Nature of the Firm.” *Economica*, 4(16): 386–405.
- Costinot, Arnaud, and Jonathan Vogel.** 2010. “Matching and Inequality in the World Economy.” *Journal of Political Economy*, 118(4): 747–786.

- Cragg, John G., and Stephen G. Donald.** 1993. “Testing Identifiability and Specification in Instrumental Variable Models.” *Econometric Theory*, 9(2): 222–240.
- Davis, Donald R., and James Harrigan.** 2011. “Good jobs, bad jobs, and trade liberalization.” *Journal of International Economics*, 84(1): 26 – 36.
- de Melo, Rafael Lopes.** 2018. “Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence.” *Journal of Political Economy*, 126(1): 313–346.
- Deming, David J.** 2017. “The Growing Importance of Social Skills in the Labor Market.” *Quarterly Journal of Economics*, 132(4): 1593–1649.
- De Negri, João Alberto, Paulo Furtado de Castro, Natalia Ribeiro de Souza, and Jorge Saba Arbache.** 2001. “Mercado Formal de Trabalho: Comparação entre os Microdados da RAIS e da PNAD.” IPEA Texto para Discussão 840.
- de Sousa, Gustavo.** 2020. “The Labor Market Consequences of (In)Appropriate Technology.” Unpublished manuscript.
- DiNardo, John, Nicole M Fortin, and Thomas Lemieux.** 1996. “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach.” *Econometrica*, 64(5): 1001–1044.
- Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge.** 2020. “Reallocation Effects of the Minimum Wage: Evidence from Germany.” Centre for Research and Analysis of Migration (CReAM), Department of Economics, University College London CReAM Discussion Paper Series 2007.
- Eeckhout, Jan, and Philipp Kircher.** 2011. “Identifying Sorting—In Theory.” *The Review of Economic Studies*, 78(3): 872–906.
- Eeckhout, Jan, and Philipp Kircher.** 2018. “Assortative Matching With Large Firms.” *Econometrica*, 86(1): 85–132.
- Eeckhout, Jan, and Roberto Pinheiro.** 2014. “Diverse Organizations and the Competition for Talent.” *International Economic Review*, 55(3): 625–664.
- Engbom, Niklas, and Christian Moser.** 2018. “Earnings Inequality and the Minimum Wage: Evidence from Brazil.” Federal Reserve Bank of Minneapolis, Opportunity and Inclusive Growth Institute Working Papers 7.

- Fernández, Manuel, and Julián Messina.** 2018. “Skill premium, labor supply, and changes in the structure of wages in Latin America.” *Journal of Development Economics*, 135(C): 555–573.
- Firpo, Sergio, and Alysson Portella.** 2019. “Decline in Wage Inequality in Brazil: A Survey.” World Bank Policy Research Working Paper 9096.
- Flinn, Christopher J.** 2006. “Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates.” *Econometrica*, 74(4): 1013–1062.
- Garicano, Luis.** 2000. “Hierarchies and the Organization of Knowledge in Production.” *Journal of Political Economy*, 108(5): 874–904.
- Garicano, Luis, and Esteban Rossi-Hansberg.** 2006. “Organization and Inequality in a Knowledge Economy.” *The Quarterly Journal of Economics*, 121(4): 1383–1435.
- Gautier, Pieter A, and Coen N Teulings.** 2015. “Sorting and the Output Loss due to Search Frictions.” *Journal of the European Economic Association*, 13(6): 1136–1166.
- Gautier, Pieter A, Coen N Teulings, and Aico van Vuuren.** 2010. “On-the-Job Search, Mismatch and Efficiency.” *The Review of Economic Studies*, 77(1): 245–272.
- Grossman, Gene M, Elhanan Helpman, and Philipp Kircher.** 2017. “Matching, Sorting, and the Distributional Effects of International Trade.” *Journal of Political Economy*, 125(1): 224–264.
- Haanwinckel, Daniel, and Rodrigo R. Soares.** 2020. “Workforce Composition, Productivity, and Labor Regulations in a Compensating Differentials Theory of Informality.” Becker Friedman Institute Working Paper 2020-45.
- Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii.** 2017. “Identifying Equilibrium Models of Labor Market Sorting.” *Econometrica*, 85(1): 29–65.
- Harasztosi, Peter, and Attila Lindner.** 2019. “Who Pays for the Minimum Wage?” *American Economic Review*, 109(8): 2693–2727.
- Herkenhoff, Kyle, Jeremy Lise, Guido Menzio, and Gordon M Phillips.** 2018. “Production and Learning in Teams.” National Bureau of Economic Research Working Paper 25179.
- Katz, Lawrence F, and Kevin M Murphy.** 1992. “Changes in Relative Wages, 1963-1987: Supply and Demand Factors.” *The Quarterly Journal of Economics*, 107(1): 35–78.

- Kline, P., R. Saggio, and M. S¸olvsten.** 2018. "Leave-out estimation of variance components." *ArXiv e-prints*.
- Kremer, Michael, and Eric Maskin.** 1996. "Wage Inequality and Segregation by Skill." National Bureau of Economic Research Working Paper 5718.
- Krusell, Per, Lee E Ohanian, Jos¸e-V¸ictor R¸ios-Rull, and Giovanni L Violante.** 1999. "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica*, 68(5): 1029–1053.
- Lamadon, Thibaut, Magne Mogstad, and Brian Setzler.** 2019. "Imperfect Competition, Compensating Differentials, and Rent Sharing in the U.S. Labor Market." Unpublished manuscript.
- Lee, David S.** 1999. "Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage?" *The Quarterly Journal of Economics*, 114(3): 977–1023.
- Lindenlaub, Ilse.** 2017. "Sorting Multidimensional Types: Theory and Application." *The Review of Economic Studies*, 84(2): 718–789.
- Lise, Jeremy, and Fabien Postel-Vinay.** 2020. "Multidimensional Skills, Sorting, and Human Capital Accumulation." *American Economic Review*, 110(8): 2328–76.
- Lise, Jeremy, Costas Meghir, and Jean-Marc Robin.** 2016. "Matching, sorting and wages." *Review of Economic Dynamics*, 19: 63 – 87. Special Issue in Honor of Dale Mortensen.
- Mak, Eric, and Aloysius Siow.** 2018. "Occupational Choice, Matching and Earnings Inequality." Unpublished manuscript.
- Manacorda, Marco, Carolina Sanchez-Paramo, and Norbert Schady.** 2010. "Changes in Returns to Education in Latin America: the Role of Demand and Supply of Skills." *Industrial and Labor Relations Review*, 63(2): 307–326.
- Sampson, Thomas.** 2014. "Selection into Trade and Wage Inequality." *American Economic Journal: Microeconomics*, 6(3): 157–202.
- Sattinger, Michael.** 1975. "Comparative Advantage and the Distributions of Earnings and Abilities." *Econometrica*, 43(3): 455–468.

- Shephard, Andrew, and Modibo Sidibe.** 2019. “Schooling Investment, Mismatch, and Wage Inequality.” PIER Working Paper 19-013.
- Shimer, Robert.** 2005. “The Assignment of Workers to Jobs in an Economy with Coordination Frictions.” *Journal of Political Economy*, 113(5): 996–1025.
- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter.** 2018. “Firming Up Inequality*.” *The Quarterly Journal of Economics*, 134(1): 1–50.
- Teulings, Coen N.** 1995. “The Wage Distribution in a Model of the Assignment of Skills to Jobs.” *Journal of Political Economy*, 103(2): 280–315.
- Teulings, Coen N.** 2000. “Aggregation Bias in Elasticities of Substitution and the Minimum Wage Paradox.” *International Economic Review*, 41(2): 359–398.
- Teulings, Coen N.** 2003. “The Contribution of Minimum Wages to Increasing Wage Inequality.” *The Economic Journal*, 113(490): 801–833.
- Teulings, Coen N.** 2005. “Comparative Advantage, Relative Wages, and the Accumulation of Human Capital.” *Journal of Political Economy*, 113(2): 425–461.
- Trottner, Fabian.** 2019. “Who Gains From Scale? Trade and Wage Inequality Within and Between Firms.” Unpublished manuscript.

Appendix — For Online Publication

A Proofs

Section 2: Task-based production function

Proof of Lemma 1: Allocation is assortative and labor constraints bind

I proceed by proving two lemmas that, together, imply the desired result. I use the term *candidate solution* to refer to tuples of output and schedules $\{q, \{m_h\}_{h=1}^H\}$ that satisfy all constraints in the assignment problem.

Lemma 3. *If there exists a candidate solution $\{q, \{m_h(\cdot)\}_{h=1}^H\}$ such that one can find two tasks $x_1 < x_2$ and two worker types $h_1 < h_2$ with $m_{h_1}(x_2) > 0$ and $m_{h_2}(x_1) > 0$, then there exists an alternative candidate solution $\{q', \{m'_h(\cdot)\}_{h=1}^H\}$ that achieves the same output ($q = q'$) but has a slack of labor of type h_1 ($l_{h_1} > \int_0^\infty m'_{h_1}(x) dx$).*

Proof. Let $\Delta = x_2 - x_1$ and pick $\tau \in (0, \min\{m_{h_1}(x_2), m_{h_2}(x_1)e_{h_2}(x_1 + \Delta)/e_{h_1}(x_1 + \Delta)\})$. Because $m_h(\cdot)$ is right continuous and the efficiency functions $e_h(\cdot)$ are strictly positive and continuous, I can find $\delta > 0$ such that $m_{h_1}(x) > \tau \forall x \in [x_2, x_2 + \delta)$ and $m_{h_2}(x_1)e_{h_2}(x_1 + \Delta)/e_{h_1}(x_1 + \Delta) > \tau \forall x \in [x_1, x_1 + \delta)$.

Now construct $\{q', \{m'_h(\cdot)\}_{h=1}^H\}$ identical to $\{q, \{m_h(\cdot)\}_{h=1}^H\}$, except for:

$$\begin{aligned} m'_{h_1}(x) &= m_{h_1}(x) - \tau, & x \in [x_2, x_2 + \delta) \\ m'_{h_2}(x) &= m_{h_2}(x) + \tau \frac{e_{h_1}(x)}{e_{h_2}(x)}, & x \in [x_2, x_2 + \delta) \\ m'_{h_2}(x) &= m_{h_2}(x) - \tau \frac{e_{h_1}(x + \Delta)}{e_{h_2}(x + \Delta)}, & x \in [x_1, x_1 + \delta) \\ m'_{h_1}(x) &= m_{h_1}(x) + \tau \frac{e_{h_1}(x + \Delta) e_{h_2}(x)}{e_{h_2}(x + \Delta) e_{h_1}(x)}, & x \in [x_1, x_1 + \delta) \end{aligned}$$

I need to prove that $\{q', \{m'_h(\cdot)\}_{h=1}^H\}$ satisfies all constraints in the assignment problem and has a slack of labor h_1 , and that $m'_h(\cdot) \in RC$. Starting with the latter, note that $m'_h(\cdot)$ is always identical to $m_h(\cdot)$ except in intervals of the form $[a, b)$. In those intervals, $m'_h(\cdot)$ is a continuous transformation of $m_h(\cdot)$. So, because $m_h(\cdot)$ is right continuous, so is $m'_h(\cdot)$. In addition, $m'_h(x) > 0 \forall x \in \mathbb{R}_{>0}$ by the condition imposed when defining δ . So $m'_h(\cdot) \in RC$.

Next, the blueprint constraints are satisfied under the new candidate solution because second and fourth rows increase task production of particular complexities in a way that exactly offsets decreased production due to the first and third rows, respectively. Total labor use of type h_2 is identical under both allocations, because the additional assignment in the second row is offset by reduced assignment in the third row. Finally, decreased use of labor type h_1 follows from log-supermodularity of the efficiency functions, which guarantees that the term multiplying τ in the fourth row is strictly less than one. So labor added in that row is strictly less than labor saved in the first row. \square

Lemma 4. *Any candidate solution with slack of labor is not optimal.*

Proof. Consider two cases:

If there is slack of labor of the highest type, $h = H$: By the feasibility condition in the definition of blueprints, $u_H = \int_0^\infty b(x)/e_H(x)dx$ is finite. Denote the slack of labor of type H in the original candidate solution by $S_H = l_H - \int_0^\infty m_H(x)dx$. Now consider an alternative candidate solution with $q' = q + S_H/u_H$, $m'_H(x) = m_H(x) + (S_H/u_H)b(x)/e_H(x)$, and $m'_h(\cdot) = m_h(\cdot) \forall h < H$. That candidate solution satisfies all constraints and achieves a strictly higher level of output. Thus, the original candidate solution is not optimal.

Otherwise: Then there is a positive slack $S_h = l_h - \int_0^\infty m_h(x)dx$ for some $h < H$, and no slack of type H . I will show that it is possible to construct an alternative allocation with the same output and positive slack of labor type H . Using that alternative allocation, one can invoke the first part of this proof to construct a third allocation with higher output.

Remember that the domain of f imposes $l_H > 0$. Because there is no slack of labor H , there must be some \underline{x} with $m_H(\underline{x}) > 0$. Pick an arbitrarily small $\tau > 0$. By right continuity of m_H , there is a small enough $\delta > 0$ such that $m_H(x) > \tau \forall x \in [\underline{x}, \underline{x} + \delta)$. Let $\tilde{u}_h = \int_{\underline{x}}^{\underline{x} + \delta} e_H(x)/e_h(x)dx < \infty$ and define $g = \min\{\tau, S_h/\tilde{u}_h\}$.

Now consider an alternative candidate solution identical to the original one, except that $m'_H(x) = m_H(x) - g$ in the interval $[\underline{x}, \underline{x} + \delta)$ and $m'_h(x) = m_h(x) + ge_H(x)/e_h(x)$ in the same interval. The new candidate solution satisfies all constraints, has right continuous and non-negative assignment functions, and has slack of labor of type H . \square

Proof of Lemma 1, except non-arbitrage condition. From Lemma 4, we know that any optimal solution must not have any slack. The same Lemma implies that any candidate solution satisfying the conditions in Lemma 3 is also not optimal. So any optimal solution must

be such that for any two tasks $x_1 < x_2$ and two types $h_1 < h_2$, $m_{h_2}(x_1) > 0 \Rightarrow m_{h_1}(x_2) = 0$ and $m_{h_1}(x_2) > 0 \Rightarrow m_{h_2}(x_1) = 0$. This property can be re-stated as: for any pair of types $h_1 < h_2$, there exists at least one number $h_1 \bar{x}_{h_2}$ such that $m_{h_2}(x) = 0 \forall x < h_1 \bar{x}_{h_2}$ and $m_{h_1}(x) = 0 \forall x > h_1 \bar{x}_{h_2}$. By combining all such requirements together, there must be $H - 1$ numbers $\bar{x}_1, \dots, \bar{x}_{H-1}$ such that, for any type h , $m_h(x) = 0 \forall x \notin [\bar{x}_{h-1}, \bar{x}_h]$ (where $\bar{x}_0 = 0$ and $\bar{x}_H = \infty$ are introduced to simplify notation).

Because there is no overlap in types that get assigned to any task (except possibly at the thresholds), the blueprint constraint implies that $m_h(x) = b(x)/e_h(x) \forall x \in (\bar{x}_{h-1}, \bar{x}_h)$. Right continuity of assignment functions means that the thresholds must be assigned to the type on the right.

It remains to be shown that the thresholds are unique and non-decreasing. To see that, recall that $b(x) > 0$ and $e_h(x) > 0 \forall h$. Now start from type $h = 1$ and note that the integral $\int_0^{\bar{x}_1} m_1(x) dx = \int_0^{\bar{x}_1} b(x)/e_1(x) dx$ is strictly increasing in \bar{x}_1 . Thus, there is only one possible $\bar{x}_1 \geq 0$ consistent with full labor use of type 1. One can then proceed by induction, showing that for any type $h > 1$, the thresholds \bar{x}_h is greater than \bar{x}_{h-1} and unique, for the same reason as in the base case.

Proof of the non-arbitrage condition (Equation 1) is provided in the next section of this Appendix. \square

Proposition 1, curvature of the production function: formulas for elasticities and proofs (including Equation 1)

Elasticities: I denote by $c = c(w, q)$ the cost function, use subscripts to denote derivatives regarding input quantities or prices, and omit arguments in functions to simplify the expressions. Then, for any pair of worker types h, h' with $h < h'$:

$$\frac{c c_{h,h'}}{c_h c_{h'}} = \begin{cases} \frac{\rho_h}{s_h s_{h'}} & \text{if } h' = h + 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{Allen partial elasticity of substitution})$$

$$\frac{f f_{h,h'}}{f_h f_{h'}} = \sum_{h=1}^{H-1} \xi_{h,h',h} \frac{1}{\rho_h} \quad (\text{Hicks partial elasticity of complementarity})$$

$$\begin{aligned}
\text{where } \rho_h &= b_g(\bar{x}_h) \frac{f_h}{e_h(\bar{x}_h)} \left[\frac{d}{d\bar{x}_h} \ln \left(\frac{e_{h+1}(\bar{x}_h)}{e_h(\bar{x}_h)} \right) \right]^{-1} \\
\xi_{h,h',\mathfrak{h}} &= \left(\mathbf{1}\{h \geq \mathfrak{h} + 1\} - \sum_{k=\mathfrak{h}+1}^H s_k \right) \left(\mathbf{1}\{\mathfrak{h} \geq h'\} - \sum_{k=1}^{\mathfrak{h}} s_k \right) \\
\text{and } s_h &= \frac{f_h l_h}{f} = \frac{c_h l_h}{c}
\end{aligned}$$

Proofs: Constant returns to scale and concavity follow easily from the definition of the production function. Let's start with concavity. Suppose that there are two input vectors l^1 and l^2 , achieving output levels q^1 and q^2 using optimal assignment functions m_h^1 and m_h^2 , respectively. Now take $\alpha \in [0, 1]$. Given inputs $\bar{l} = \alpha l^1 + (1 - \alpha)l^2$, one can use assignment functions defined by $\bar{m}_h(x) = \alpha m_h^1(x) + (1 - \alpha)m_h^2(x) \forall x, h$ to achieve output level $\bar{q} = \alpha q^1 + (1 - \alpha)q^2$, while satisfying blueprint and labor constraints. So $f(\bar{l}, b) \geq \bar{q}$. For constant returns, note that, given $\alpha > 1$, output αq^1 is attainable with inputs αl^1 by using assignment functions $\alpha m_h^1(x)$. Together with concavity, that implies constant returns to scale.

Lemma 1 implies that, given inputs $(l, b_g(\cdot))$, the optimal thresholds and the optimal production level satisfy the set of H labor constraints with equality. I will now prove results that justify using the implicit function theorem on that system of equations. That will prove twice differentiability and provide a path to obtain elasticities of complementarity and substitution.

Definition 4. The excess labor demand function $z : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^{H-1} \times \mathbb{R}_{\geq 0}^{H-1} \times \mathbb{R}_{> 0} \rightarrow \mathbb{R}^H$ is given by:

$$z_h(q, \bar{x}_1, \dots, \bar{x}_{H-1}; l) = q \int_{\bar{x}_{h-1}}^{\bar{x}_h} \frac{b_g(x)}{e_h(x)} dx - l_h$$

Lemma 5. The excess labor demand function is C^2 .

Proof. We need to show that, for all components $z_h(\cdot)$, the second partial derivatives exist and are continuous. This is immediate for the first derivatives regarding q and l , as well as for their second own and cross derivatives (which are all zero).

The first derivative regarding threshold $\bar{x}_{h'}$ is:

$$\frac{\partial z_h(\cdot)}{\partial \bar{x}_{h'}} = q \left[\mathbf{1}\{h' = h\} \frac{b_g(\bar{x}_h)}{e_h(\bar{x}_h)} - \mathbf{1}\{h' = h-1\} \frac{b_g(\bar{x}_h)}{e_{h+1}(\bar{x}_h)} \right]$$

Because blueprints and efficiency functions are continuously differentiable and strictly positive, this expression is continuously differentiable in \bar{x}_h . The cross-elasticities regarding q

and l also exist and are continuous. □

Lemma 6. *The Jacobian of the excess labor demand function regarding $(q, \bar{x}_1, \dots, \bar{x}_{H-1})$, when evaluated at a point where $z(\cdot) = \mathbf{0}_{H \times 1}$, has non-zero determinant.*

Proof. The Jacobian, when evaluated at the solution to the assignment problem, is:

$$J = \begin{bmatrix} \frac{l_1}{q} & q \frac{b_g(\bar{x}_1)}{e_1(\bar{x}_1)} & 0 & 0 & \dots & 0 & 0 \\ \frac{l_2}{q} & -q \frac{b_g(\bar{x}_1)}{e_2(\bar{x}_1)} & q \frac{b_g(\bar{x}_2)}{e_2(\bar{x}_2)} & 0 & \dots & 0 & 0 \\ \frac{l_3}{q} & 0 & -q \frac{b_g(\bar{x}_2)}{e_3(\bar{x}_2)} & q \frac{b_g(\bar{x}_3)}{e_3(\bar{x}_3)} & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \frac{l_{H-1}}{q} & 0 & 0 & 0 & \dots & -q \frac{b_g(\bar{x}_{H-2})}{e_{H-1}(\bar{x}_{H-2})} & q \frac{b_g(\bar{x}_{H-1})}{e_{H-1}(\bar{x}_{H-1})} \\ \frac{l_H}{q} & 0 & 0 & 0 & \dots & 0 & -q \frac{b_g(\bar{x}_{H-1})}{e_H(\bar{x}_{H-1})} \end{bmatrix}$$

The determinant is:

$$|J| = (-1)^{H+1} q^{H-2} \left[\prod_{h=1}^{H-1} \frac{b_g(\bar{x}_h)}{e_{h+1}(\bar{x}_h)} \right] \sum_{h=1}^H \left(l_h \prod_{i=2}^h \frac{e_i(\bar{x}_{i-1})}{e_{i-1}(\bar{x}_{i-1})} \right)$$

which is never zero, since $q > 0$ (from feasibility of blueprints and $l_H > 0$) and $b(x), e_h(x) > 0 \forall x, h$. □

Lemmas 5 and 6 mean that the implicit function theorem can be used at the solution to the assignment problem to obtain derivatives of the solutions to the system of equations imposed by the labor constraints. These solutions are $q(l) = f(l, b_g(\cdot))$ and $\bar{x}_h(l)$. Because z is C^2 , so are the production function and the thresholds as functions of inputs.

Obtaining the ratios of first derivatives in Lemma 1 and the elasticities of complementarity and substitution in Proposition 1 is a matter of tedious but straightforward algebra, starting from the implicit function theorem. For the non-arbitrage condition in Lemma 1, a simpler approach is to define the allocation problem in terms of choosing output and thresholds, and then use a Lagrangian to embed the labor constraints into the objective function. Then, the result of Lemma 1, along with the constant returns relationship $q = \sum_h l_h f_h$, emerge as first order conditions, after noting that the Lagrange multipliers are marginal productivities.

When working towards second derivatives, it is necessary to use the derivatives of thresholds regarding inputs. For reference, here is the result:

$$\frac{d\bar{x}_h}{dl_{h'}} = \frac{e_h(\bar{x}_h)}{qb_g(\bar{x}_h)} \frac{f_{h'}}{f_h} \left[\mathbf{1}\{h \geq h'\} - \sum_{i=1}^h s_i \right]$$

One can verify $\frac{d\bar{x}_h}{dl_{h'}} > 0 \Leftrightarrow h \geq h'$. Adding labor "pushes" thresholds to the right or to the left depending on whether the labor which is being added is to the left or to the right of the threshold in question.

Proof of Corollary 1: Distance-dependent complementarity

This is proven by inspecting the sign of the weights $\xi_{h,h',h}$ above. When $h = h'$, these terms are negative for all i . Changing h' by one, either up or down, changes one of the $\xi_{h,h',h}$ from negative to positive while keeping the others unchanged. So there must be an increase in the elasticity of complementarity since all of the $\rho_{h'}$ are positive. Every additional increment or decrement of h' away from h involves a similar change of sign in one of the $\xi_{h,h',h}$, leading to the same increase in complementarity.

Section 3: Markets and wages

Proof of Lemma 2: Firm problem and representative firms

I start by establishing that the solution must have positive employment of all types. The marginal product of an efficiency unit of labor of the highest type is bounded below by $1/\int_0^\infty b_g(x)/e_H(x)dx = \underline{f}_H$, which is strictly positive due to the feasibility condition imposed on blueprints. Consider the strategy of posting a fixed payment $y_{Hj}(\varepsilon) = \bar{y} \geq \underline{y}$ to all workers with $\varepsilon > \underline{\varepsilon}_{Hj}$. Profit from workers of type H associated with that strategy are bounded below by $\int_{\underline{\varepsilon}_{Hj}}^\infty N_{Haj} \bar{y}^\beta / \omega_H(\varepsilon)^\beta r_H(\varepsilon) (p_g \underline{f}_H \varepsilon - \bar{y}) d\varepsilon$. That expression is assured to be positive for high enough $\underline{\varepsilon}_{Hj}$ (note that $\omega_h(\varepsilon)$ is always finite in an equilibrium). Thus, positive employment of skilled workers following that strategy is more profitable than not employing any of those workers.

A positive amount of l_H ensures that all other types are employed as well. Consider a particular type $h < H$ and whether it is optimal to set $l_h = 0$, fixing employment of all other types. Because $l_H > 0$, \bar{x}_{H-1} is finite, and thus threshold \bar{x}_h (the highest task performed by h) is guaranteed to be finite as well. Then, from Equation 1, the marginal product of type

h is bound below by $f_{He_h}(\bar{x}_{H-1})/e_H(\bar{x}_{H-1})$. A similar reasoning as above establishes that employing small quantities of labor h is more profitable than setting $l_h = 0$.

The rest of the proof follows from the logic described in the text. The threshold $\underline{\epsilon}_{hj}$ is chosen so that the worker with the least amount of efficiency units pays for himself, bringing in revenue equal to the minimum wage. Below that, labor payments — which are bound by the minimum wage — will necessarily exceed marginal revenue from those workers. For every $\epsilon > \underline{\epsilon}_{hj}$, the firm chooses $y_{hj}(\epsilon)$ by equating marginal revenue from workers of that (h, ϵ) combination with their marginal cost. For high enough ϵ , that leads to the constant markdown rule, implying that earnings are proportional to marginal product of labor — and thus linear in ϵ . Workers close to the cutoff are still profitable, but for them, the minimum wage constraint binds.

To see why these solutions do not depend on amenities, such that there is a representative firm for each good g , first note that a_j is a multiplicative term in both $C_h(y_{hj}, \underline{\epsilon}_{hj}, a_j)$ and $l_h(y_{hj}, \underline{\epsilon}_{hj}, a_j)$. Now remember that the task-based production function has constant returns to scale. Thus, the profit function can be rewritten as $\pi(a_j) = a_j \pi(1)$. Amenities scale up employment and production while keeping average labor costs constant.

Proof of Proposition 2: Wage differentials across firms

I start by proving a useful Lemma that shows how proportional terms dividing task requirements can be interpreted as physical productivity shifters.

Lemma 7. *If $b_g(x) = b(x)/z_g$ for a blueprint $b(\cdot)$ and scalar $z_g > 0$, then $f(\mathbf{l}, b_g(\cdot)) = z_g f(\mathbf{l}, b(\cdot))$.*

Proof. Plug $b_g(x) = b(x)/z_g$ into the assignment problem defining the task-based production function. Change the choice variable to $q' = q/z_g$. The z_g terms in the task constraint cancel each other and the maximand changes to $z_g q'$. The result follows from noting that $\max_{\{.\}} z_g q' = z_g \max_{\{.\}} q'$ and that the resulting value function is $f(\mathbf{l}, b(\cdot))$ by definition. \square

Now I proceed to the proof of each statement of Proposition 2 separately.

Proof of part 1: From Lemma 7, $f_h(\mathbf{l}, b_g(\cdot)) = z_g f_h(\mathbf{l}, b(\cdot))$. Also note $\mathbf{l}(\mathbf{w}_g, \underline{\epsilon}_g, \bar{a}_g) = \bar{a}_g \mathbf{l}(\mathbf{w}_g, \underline{\epsilon}_g, 1)$ and $\mathbf{C}(\mathbf{w}_g, \underline{\epsilon}_g, \bar{a}_g) = \bar{a}_g \mathbf{C}(\mathbf{w}_g, \underline{\epsilon}_g, 1)$, and remember that the task-based production function has constant returns to scale (and so marginal productivities are homogeneous of degree

zero). Now let $\tilde{F} = F_g/\bar{a}_g$ and rewrite the first order conditions of the firm (7), (8) and the zero profits condition (11) imposing the conditions from this proposition:

$$\begin{aligned} p_g z_g f_h(\mathbf{l}(\mathbf{w}_g, \boldsymbol{\epsilon}_g, 1), b(\cdot)) \exp(\boldsymbol{\epsilon}_{hg}) &= \underline{y} & \forall h, g \\ p_g z_g f_h(\mathbf{l}(\mathbf{w}_g, \boldsymbol{\epsilon}_g, 1), b(\cdot)) \frac{\beta}{\beta + 1} &= w_{hg} & \forall h, g \\ \bar{a}_g \left[p_g z_g f(\mathbf{l}(\mathbf{w}_g, \boldsymbol{\epsilon}_g, 1), b(\cdot)) - \sum_{h=1}^H C_h(\mathbf{w}_g, \boldsymbol{\epsilon}_g, 1) \right] &= \bar{a}_g \tilde{F} & \forall g \end{aligned}$$

To see that these equations imply a representative firm for the economy, plug in $\boldsymbol{\epsilon}_g = \boldsymbol{\epsilon}$, $\mathbf{w}_g = \boldsymbol{\lambda} = \{\lambda_1, \dots, \lambda_H\}$, and $p_g = p/z_g$ for common $\boldsymbol{\epsilon}$, $\boldsymbol{\lambda}$, and p . All dependency on g is eliminated, showing that the solution of the problem of the firm is the same for all firms in the economy and that prices are inversely proportional to physical productivity shifters z_g (such that marginal revenue product of labor is equalized across firms). \square

Proof of part 2: Without a minimum wage, there is no motive for a cutoff rule: $\boldsymbol{\epsilon}_{hg} = 0$. In addition, the labor supply curve becomes isoelastic with identical elasticities for all worker types:

$$\begin{aligned} l_h(w_{hg}, \cdot, \bar{a}_g) &= \bar{a}_g \left(\frac{w_{hg}}{\boldsymbol{\omega}_h} \right)^\beta \\ C_h(w_{hg}, \cdot, \bar{a}_g) &= w_{hg} l_h(w_{hg}, \cdot, \bar{a}_g) \\ \text{where } \boldsymbol{\omega}_h &= \left(\sum_g J_g \bar{a}_g w_{hg}^\beta \right)^{\frac{1}{\beta}} \end{aligned}$$

Rewrite the first order conditions on wages as in the proof of part 1 above:

$$p_g z_g f_h(\mathbf{l}(\mathbf{w}_g, \cdot, 1), b(\cdot)) \frac{\beta}{\beta + 1} = w_{hg} \quad \forall h, g$$

Also, rewrite the zero profit condition as:

$$\begin{aligned} F_g &= p_g z_g f(\mathbf{l}(\mathbf{w}_g, \cdot, \bar{a}_g), b(\cdot)) - \sum_{h=1}^H C_h(\mathbf{w}_g, \cdot, \bar{a}_g) \\ &= p_g z_g \sum_{h=1}^H l_h(w_{hg}, \cdot, \bar{a}_g) f_h(\mathbf{l}(\mathbf{w}_g, \cdot, 1), b(\cdot)) - \sum_{h=1}^H w_{hg} l_h(w_{hg}, \cdot, \bar{a}_g) \end{aligned}$$

I claim that $\mathbf{w}_g = (F_g/\bar{a}_g)^{1/(\beta+1)}\boldsymbol{\lambda}$ for some vector $\boldsymbol{\lambda} = \{\lambda_1 \dots, \lambda_H\}$. From the labor supply equation, that implies $l_{hg} = F_g^{\beta/(\beta+1)}\bar{a}_g^{1/(\beta+1)}\ell_h$, where $\ell_h = \omega_h^{-\beta/(\beta+1)}$. Plugging these expressions in the rewritten zero profit condition yields $\sum_h \ell_h \lambda_h = 1 \forall g$, showing that the claim does not contradict optimal entry behavior; instead, optimal entry merely imposes a normalization on the $\boldsymbol{\lambda}$ vector.

The corresponding prices that lead to zero profits are:

$$\begin{aligned} \Rightarrow p_g &= \frac{(\beta+1)F_g}{z_g f(\mathbf{l}(\mathbf{w}_g, \cdot, \bar{a}_g), b(\cdot))} \\ &= \frac{\beta+1}{z_g f(\boldsymbol{\ell}, b(\cdot))} \left(\frac{F_g}{\bar{a}_g} \right)^{\frac{1}{\beta+1}} \end{aligned}$$

Finally, plugging these results into the first order conditions yields:

$$f_h(\boldsymbol{\ell}, b) \beta = \lambda_h \quad \forall h, g$$

Which again has no dependency on g , showing that the claimed solution solves the problem for all firms. \square

Proof of part 3: Under the conditions from this part, labor supply curves are isoelastic, as shown in the proof of part 2 above. It is easily shown, using that isoelastic expression for $l_h(\cdot)$, that:

$$\left(\frac{w_{h'g'}}{w_{hg'}} \right) / \left(\frac{w_{h'g}}{w_{hg}} \right) = \left[\left(\frac{l_{h'g'}}{l_{hg'}} \right) / \left(\frac{l_{h'g}}{l_{hg}} \right) \right]^{\frac{1}{\beta}}$$

Under the condition imposed on labor input ratios, the right hand side is positive. The proof follows from noting that the desired ratio of earnings is equal to the ratio of wages in the left hand side. \square

Proof of Proposition 3: Changes in firm costs affect the returns to skill

Before proving the Proposition, I derive a Lemma that states that blueprints that are more intensive in complex tasks lead to higher gaps in marginal productivity, holding constant the quantity of labor. This Lemma is conceptually similar to the monotone comparative statics

in Costinot and Vogel (2010).

Lemma 8. *Let b and b' denote blueprints such that their ratio $b'(x)/b(x)$ is strictly increasing. Then:*

$$\frac{f_{h+1}(\mathbf{l}, b')}{f_h(\mathbf{l}, b')} > \frac{f_{h+1}(\mathbf{l}, b)}{f_h(\mathbf{l}, b)} \quad h = 1, \dots, H-1$$

Proof. Fix \mathbf{l} , let $q = f(\mathbf{l}, b)$ and $q' = f(\mathbf{l}, b')$. Now construct $b''(x) = b'(x)q'/q$. From Lemma 7, it follows that $f(\mathbf{l}, b'') = q$ and $f_h(\mathbf{l}, b'') = f_h(\mathbf{l}, b') \forall h$. I will show that the statement holds for b and b'' , and since b'' and b' lead to the same marginal products, the desired result holds.

Because b and b'' lead to the same output given the same vector of inputs, but $b''(x)/b(x)$ is increasing, there must be a task x^* such $b''(x) < b(x) \forall x < x^*$ and $b''(x) > b(x) \forall x > x^*$. To see why they must cross at least once at x^* , suppose otherwise (one blueprint is strictly more than other for all x): there will be a contradiction since task demands are strictly higher for one of the blueprints, but they still lead to the same production q given the same vector of inputs. From this crossing point, differences before and after emerge from the monotonic ratio property.

Now note from the non-arbitrage condition (1) in Lemma 1, along with log-supermodularity of $e_h(x)$, that the statement to be proved is equivalent to

$$\bar{x}'_h \geq \bar{x}_h \quad h \in \{1, \dots, H-1\}$$

where \bar{x}'_h denotes thresholds under the alternative blueprint b'' .

I proceed by using compensated labor demand integrals to show that thresholds differ as stated above. Denote by h^* the type such that $x^* \in [\bar{x}_{h^*-1}, \bar{x}_{h^*})$. The proof will be done in two parts: starting from \bar{x}'_1 and ascending by induction up to \bar{x}_{h^*-1} , and next starting from \bar{x}_{h-1} and descending by induction down to \bar{x}_{h^*} . Note that if $h^* = 1$ or $h^* = H$, only one part is required.

Base case \bar{x}_1 : The equation for $h = 1$ is $\int_0^{\bar{x}_1} \frac{b(x)}{e_1(x)} dx = \frac{l_1}{q}$ under the original blueprint, and $\int_0^{\bar{x}'_1} \frac{b''(x)}{e_1(x)} dx = \frac{l_1}{q}$ under the new one. Equating the right hand side of both expressions and rearranging yields:

$$\int_{\bar{x}'_1}^{\bar{x}_1} \frac{b''(x)}{e_1(x)} dx = \int_0^{\bar{x}_1} \frac{b(x) - b''(x)}{e_1(x)} dx$$

Since $b(x) \geq b''(x)$ for $x < x^*$, the right-hand side is positive, and then the equality will only hold if $\bar{x}'_1 \geq \bar{x}_1$.

Ascending induction rule: Suppose $\bar{x}'_{h-1} \geq \bar{x}_{h-1}$ and $h < h^*$. I will prove that $\bar{x}'_h \geq \bar{x}_h$. To do so, use the fact that $\frac{l_h}{q}$ is the same under both the old and new blueprints to equate the labor demand integrals, as was done in the base case. This yields the following equivalent expressions:

$$\begin{aligned} \int_{\bar{x}_h}^{\bar{x}'_h} \frac{b''(x)}{e_h(x)} dx &= \int_{\bar{x}_{h-1}}^{\bar{x}'_{h-1}} \frac{b(x)}{e_h(x)} dx + \int_{\bar{x}'_{h-1}}^{\bar{x}_h} \frac{b(x) - b''(x)}{e_h(x)} dx \\ &= \int_{\bar{x}_{h-1}}^{\bar{x}_h} \frac{b(x)}{e_h(x)} dx + \int_{\bar{x}_h}^{\bar{x}'_{h-1}} \frac{b''(x)}{e_h(x)} dx \end{aligned}$$

It is enough to show that the expression is positive, ensuring that $\bar{x}'_h \geq \bar{x}_h$. Consider two cases. If $\bar{x}'_{h-1} \leq \bar{x}_h$, then use the first expression. The induction assumption guarantees positivity of the first term, and the integrand of the second term is positive because $\bar{x}_h < z^*$. If instead $\bar{x}'_{h-1} > \bar{x}_h$, the second expression is more convenient. There, all integrands are positive and the integration upper bounds are greater than the lower bounds.

Base case \bar{x}_{H-1} and descending induction rule: Those are symmetric to the cases above. \square

In a competitive economy, thresholds are the same for all firms. Given total endowments of labor efficiency units L and aggregate demand for tasks $B(x) = Q_1 b_1(x) + Q_2 b_2(x)$ (where Q_g denotes aggregate demand for good g before the shock), wages w_h must be proportional to marginal productivities $f_h(L, B(\cdot))$, because the labor constraints that determine thresholds and marginal productivities in the task-based production function are the labor clearing conditions for this economy.

Aggregate demand for tasks following the shock is $B'(x) = Q'_1 b_1(x) + Q'_2 b_2(x)$. As noted above, wages after the shock are proportional to $f_h(L, B'(\cdot))$. But $B(x, Q'_1, Q'_2)/B(x, Q_1, Q_2)$ is increasing in x if $Q'_2/Q'_1 > Q_2/Q_1$. And an increase in relative taste for good 2, holding all else equal, necessarily implies an increase in aggregate consumption of good 2 relative to good 1. Thus, Lemma 8 implies that wage gaps increase as stated in the Proposition.

B Tinbergen's race

The following proposition considers a case in which the supply of skill, demand for task complexity, and minimum wages rise in tandem:

Proposition 4 (Race between technology, education, and minimum wages). *Start with a baseline economy characterized by parameters $(\{e_h, N_h, \beta\}_{h=1}^H, \{b_g, F_g, \bar{a}_g\}_{g=1}^G, T, \sigma, \underline{y})$ and consider a new set of parameters denoted with prime symbols. Assume e_h are decreasing functions to simplify interpretation (more complex tasks are harder to produce). Let Δ_0, Δ_1 and Δ_2 denote arbitrary positive numbers and consider the following conditions:*

1. $N'_h = \Delta_0 N_h \forall h$ and $T' = \Delta_0 T$: *The relative supply of factors remains constant.*
2. $e'_h(x) = e_h\left(\frac{x}{1+\Delta_1}\right) \forall h$: *Workers become better at all tasks and the degree of comparative advantage becomes smaller for the current set of tasks (e.g. both high school graduates and college graduates improve at using text editing software, but the improvement is larger for high school graduates).*
3. $b'_g(x) = \frac{1}{(1+\Delta_1)(1+\Delta_2)} b_g\left(\frac{x}{1+\Delta_1}\right) \forall g$: *Production requires fewer tasks, but the composition of tasks moves towards increased complexity.*
4. $\underline{y}' = \underline{y}$: *The minimum wage stays constant relative to the price of entrepreneurial talent.*

*If these conditions are satisfied, the equilibrium under the new parameter set is identical to the initial equilibrium, except that prices for goods are uniformly lower: $p'_g = p_g/(1+\Delta_2)$ and $P' = P/(1+\Delta_2)$.*²⁶

Proof. The proof is simple once one notes that the difference between the two economies is a linear change of variables in the task space $x' = (1+\Delta_1)x$, coupled with a reduction in task demand by a factor of $(1+\Delta_2)$. Let \bar{x}_h^g denote task thresholds for firm g in the original equilibrium. Thresholds $(1+\Delta_1)\bar{x}_h^g$ lead to exactly the same unit labor demands, except for a proportional reduction:

$$\int_{(1+\Delta_1)\bar{x}_{h-1}^g}^{(1+\Delta_1)\bar{x}_h^g} \frac{b'_g(x')}{e'_h(x')} dx' = \int_{(1+\Delta_1)\bar{x}_{h-1}^g}^{(1+\Delta_1)\bar{x}_h^g} \frac{1}{(1+\Delta_1)(1+\Delta_2)} \frac{b_g(x'/(1+\Delta_1))}{e_h(x'/(1+\Delta_1))} dx' = \frac{1}{1+\Delta_2} \int_{\bar{x}_{h-1}^g}^{\bar{x}_h^g} \frac{b_g(x)}{e_h(x)} dx$$

So if firms use exactly the same labor inputs, they will produce $(1+\Delta_2)$ times more goods.

²⁶Using the exponential-gamma parametrization, changes in comparative advantage functions and blueprints are equivalent to $\alpha'_h = \alpha_h/(1+\Delta_1)$, $\theta'_g = (1+\Delta_1)\theta_g$, $k'_g = k_g$, and $z'_g = (1+\Delta_2)z_g$.

But because $p'_g = p_g/(1 + \Delta_2)$, total and marginal revenues are the same. Since all other equilibrium variables are the same, all equilibrium conditions are still satisfied. \square

Proposition 4 delineates balanced technological progress in this economy. Production becomes more efficient by using tasks that are more complex. At the same time, the skill of workers increases, changing the set of tasks where skill differences are relevant. If minimum wages remain as important, then there is a uniform increase in living standards. Wage differences between worker groups and across firms for workers in the same group remain stable.

C Numerical implementation

C.1 Task-based production function

The basic logic of obtaining compensated labor demands in this model is to use the non-arbitrage equation 1 from Lemma 1 to obtain thresholds as functions of marginal productivity gaps. Then, compensated labor demands can be obtained through numerical integration of Equation 2.

The exponential-Gamma parametrization is helpful because it provides a simple closed form solution for thresholds and the labor demand integrals. Let:

$$\begin{aligned}
 e_h(x) &= \exp(\alpha_h x) & -1 &= \alpha_1 < \alpha_2 < \dots < \alpha_{H-1} < \alpha_H = 0 \\
 b_g(x) &= \frac{x^{k_g-1}}{z^g \Gamma(k_g) \theta_g^{k_g}} \exp\left(-\frac{x}{\theta_g}\right) & (z_g, \theta_g, k_g) &\in \mathbb{R}_{>0}^3
 \end{aligned}$$

Then:

$$\bar{x}_h \left(\frac{f_{h+1}}{f_h} \right) = \frac{\log f_{h+1}/f_h}{\alpha_{h+1} - \alpha_h} \quad (13)$$

$$\begin{aligned} \ell_{hg}(\bar{x}_{h-1}, \bar{x}_h) &= \int_{\bar{x}_{h-1}}^{\bar{x}_h} \frac{b_g(x)}{e_h(x)} dx \\ &= \begin{cases} \frac{1}{z_g \Gamma(k_g)} \left(\frac{1}{\Upsilon_{hg} \theta_g} \right)^{k_g} [\gamma(\Upsilon_{hg} \bar{x}_h, k_g) - \gamma(\Upsilon_{hg} \bar{x}_{h-1}, k_g)] & \text{if } \Upsilon_{hg} \neq 0 \\ \frac{1}{z_g k_g \Gamma(k_g)} [(\bar{x}_h/\theta_g)^{k_g} - (\bar{x}_{h-1}/\theta_g)^{k_g}] & \text{otherwise} \end{cases} \quad (14) \\ &= \begin{cases} \frac{\sum_{m=0}^{\infty} \frac{\bar{x}_h^{k_g} \exp(-\Upsilon_{hg} \bar{x}_h) (\Upsilon_{hg} \bar{x}_h)^m - \bar{x}_{h-1}^{k_g} \exp(-\Upsilon_{hg} \bar{x}_{h-1}) (\Upsilon_{hg} \bar{x}_{h-1})^m}{z_g \theta_g^{k_g} \Gamma(k_g + m + 1)}}{z_g \theta_g^{k_g} \Gamma(k_g + m + 1)} & \text{if } \Upsilon_{hg} \neq 0 \\ \frac{1}{z_g k_g \Gamma(k_g)} [(\bar{x}_h/\theta_g)^{k_g} - (\bar{x}_{h-1}/\theta_g)^{k_g}] & \text{otherwise} \end{cases} \quad (15) \end{aligned}$$

where $\Upsilon_{hg} = \alpha_h + \frac{1}{\theta_g}$, $\gamma(\cdot, \cdot)$ is the lower incomplete Gamma function, and $\Gamma(\cdot)$ is the Gamma function.

Expression 14 is simple to code and fast to run in software packages such as Matlab, where optimized implementations of the incomplete Gamma function are available.²⁷ When $\Upsilon_{hg} < 0$, that expression requires calculating complex numbers as intermediate steps. This is not a problem in Matlab. If using complex numbers is not convenient, then the power series representation in 15 should be used. Another option is to change the normalization of α_h such that they are all non-negative.

Calculating the production function and its derivatives — that is, solving for output and marginal productivities given labor inputs — is not needed in the equilibrium computation nor in estimation. However, it might be useful for other purposes. Those numbers are obtained from a system of H equations implied by requiring that labor demand equals labor available to the firm. The choice variables can be either $(q, \bar{x}_1, \dots, \bar{x}_{H-1})$ or f_1, \dots, f_H . Moving from thresholds and output to marginal productivities, or vice-versa, is a matter of applying the constant returns relation $\sum_h f_h = q$.

²⁷Note that Matlab's *gammainc* yields a normalized incomplete Gamma function, so dividing by $\Gamma(k_g)$ is not necessary.

C.2 Equilibrium

Solving for equilibrium can seem challenging at first glance. Using a convenient set of choice variables reduces the problem to solving a square system of $(H + 1) \times G$ equations where the choice variables are firm-specific task thresholds, firm-level output, and prices for each good. The procedure below describes how to calculate that system of equations:

1. Start with values for mean output \bar{q}_g and task thresholds $\bar{x}_g = \{\bar{x}_{1g}, \dots, \bar{x}_{Hg}\}$ for the representative firms of each type, along with prices for goods p_g .
2. Use the compensated labor demand integral for the task-based production function to find average labor demands \bar{l}_{hg} (Equation 2 in the text, or Equation 14 in Appendix C if using the exponential-Gamma parametrization).
3. Find marginal products of labor f_{hg} via the non-arbitrage conditions (1) and the constant returns to scale relationship $\sum_h f_{hg} \bar{l}_{hg} = \bar{q}_g$.
4. Employ the first order conditions of the firm (7) and (8) to find wages w_{hg} and rejection cutoffs $\underline{\varepsilon}_{hg}$, respectively.
5. Calculate relative consumption $Q_g/Q_1 = (p_g/p_1)^{-\sigma}$ and relative firm entry $J_g/J_1 = (Q_g/Q_1)/(\bar{q}_g/\bar{q}_1)$.
6. Pin down entry of firm type 1 (and thus all others) with entrepreneurial talent clearing: $J_1 = T/(\sum_g F_g J_g/J_1)$.
7. Obtain $\omega_h(\varepsilon)$ using expression 9.
8. Calculate the error in the system of equations, which has two components:
 - (a) The deviation between \bar{l}_{hg} found in step 2 and that implied by the labor supply curve (5).
 - (b) The deviation between profits and entry costs in Equation 11.

That system of equations can be solved using standard numerical procedures, with the restrictions that $\bar{q}_g > 0$, $p_g > 0$, and $0 \leq \bar{x}_{1g} \leq \bar{x}_{2g} \leq \dots \leq \bar{x}_{Hg} \forall g$. These restrictions can be imposed through transformations of the choice variables: log prices, log quantities, log of the lowest task thresholds \bar{x}_{1g} , and log of differences between consecutive thresholds $\bar{x}_{hg} - \bar{x}_{h-1,g}$ for $h = 2, \dots, H - 1$.

D Appendix to the quantitative exercise

D.1 Summary statistics

Descriptive statistics for the RAIS dataset are presented in Table D1. Statistics are presented for the whole sample and separately by schooling group.

D.2 Wage inequality and schooling trends using PNAD data

In this Appendix, I analyze the robustness of the main facts presented in Section 4.1 using an alternative data source, the PNAD survey. I proceed in three steps. First, I compare wage inequality and schooling trends for formal sector workers in the two datasets. Second, I expand the sample to include both formal and informal workers to check whether these trends are restricted to the formal sector. Third, I look at schooling achievement for Brazilian adults regardless of their workforce participation status, as a way of investigating whether increased schooling achievement among employed workers reflects changes in selection patterns into employment or fundamental changes in access to schooling for the whole population.

The PNAD is a household survey with national coverage administered by the by the Brazilian Statistical Bureau (IBGE). Jointly with the Census, it is one of the primary sources of nationally representative data on a series of topics that include labor market participation, earnings, and education. It contains thorough information on employment status, including whether workers had a signed "labor card" — that is, whether the employment relationship is formally registered.

This Appendix analyzes PNAD data from 1998 through 2012. The sample I use includes adults 18 through 54 years old that are not in school, the same criterion imposed on RAIS data. I use public use software developed by PUC-Rio's Datazoom project to read the data, make it compatible across years, and deflate income variables. More information about the resulting dataset is available at Datazoom's website.²⁸

D.2.1 Comparing RAIS data and PNAD data for formal sector workers

Figure D1 replicates Figure 5 using PNAD data instead of RAIS data. The PNAD sample is constructed to match the RAIS sample, including only formal employees. Overall, the patterns are broadly similar: they show decreased wage inequality along different dimensions.

²⁸Currently located at <http://www.econ.puc-rio.br/datazoom/english/index.html>.

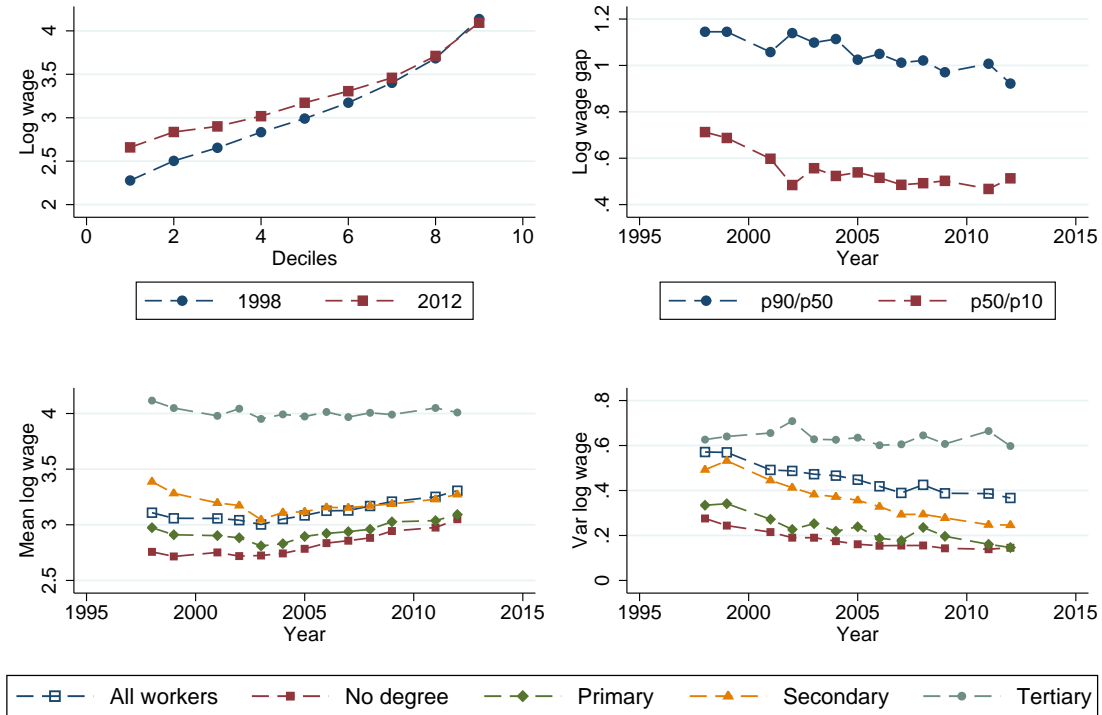
Table D1: Summary statistics, RAIS data.

	All Workers	No degree	Primary	Secondary	Tertiary
<i>Panel A: 1998</i>					
Age	33.867 (9.338)	34.838 (9.601)	32.164 (9.296)	32.054 (8.615)	38.716 (7.700)
Female	0.411 (0.492)	0.303 (0.459)	0.365 (0.482)	0.536 (0.499)	0.626 (0.484)
Log wage	1.759 (0.828)	1.447 (0.593)	1.596 (0.688)	2.006 (0.840)	2.725 (0.938)
Public sector	0.274 (0.446)	0.195 (0.396)	0.190 (0.392)	0.335 (0.472)	0.622 (0.485)
Monthly hours	179.374 (26.578)	185.830 (19.033)	183.650 (20.340)	173.885 (29.749)	158.111 (39.292)
Number of workers	1,494,186	574,904	394,990	364,376	159,916
<i>Panel B: 2012</i>					
Age	34.501 (9.890)	38.682 (9.865)	34.015 (10.513)	32.554 (9.329)	37.727 (8.642)
Female	0.452 (0.498)	0.327 (0.469)	0.361 (0.480)	0.476 (0.499)	0.636 (0.481)
Log wage	1.978 (0.701)	1.692 (0.434)	1.732 (0.487)	1.903 (0.597)	2.909 (0.776)
Public sector	0.192 (0.393)	0.138 (0.344)	0.109 (0.311)	0.152 (0.359)	0.512 (0.500)
Monthly hours	179.376 (27.319)	186.569 (17.788)	185.134 (19.368)	182.107 (22.342)	153.702 (42.728)
Number of workers	2,398,391	350,704	517,748	1,189,063	340,876

This table presents summary statistics (means and standard deviations, in parenthesis) for the RAIS data. The sample includes adults in Rio Grande do Sul state from 18 to 54 years of age who are not in school and who are employed in December, having been hired in November or earlier. Wages are in 2010 Brazilian Reais and are winsorized at the top and bottom 1 percent of the wage distribution in each year.

Figure D1: Measures of wage dispersion, PNAD data, formal sector

Wage inequality, PNAD data, formal workers

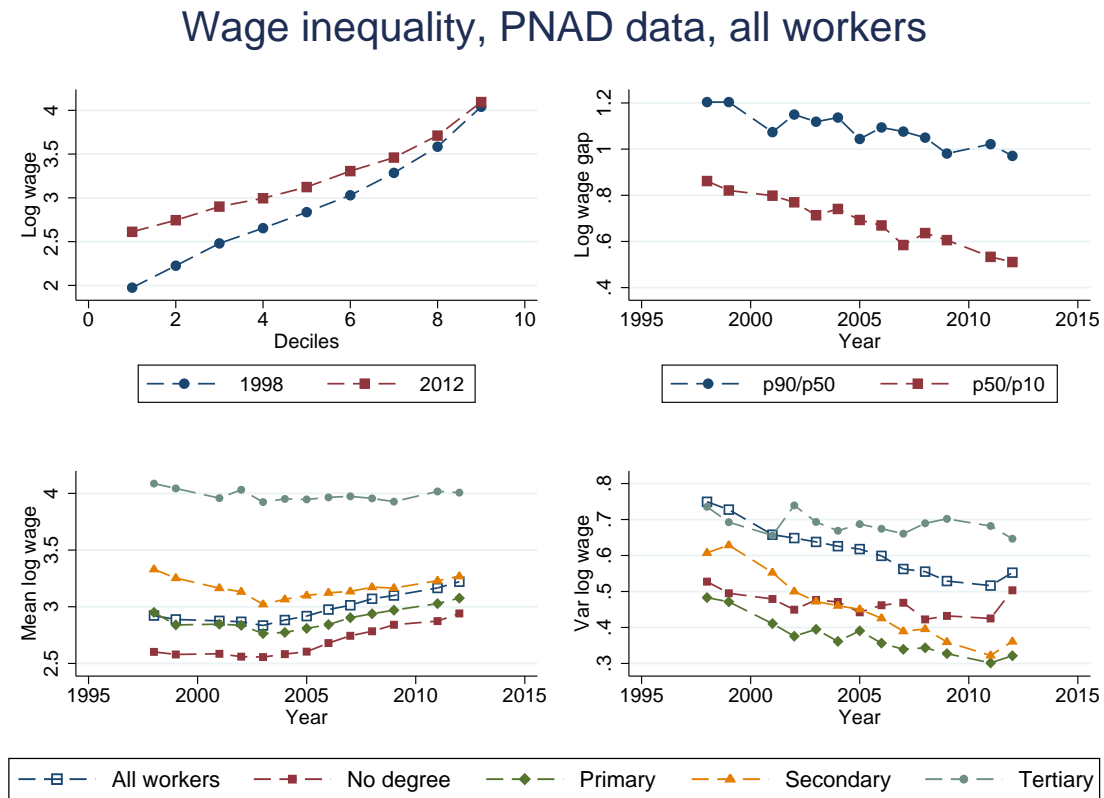


Notes: PNAD data, Rio Grande do Sul, Brazil. Formal sector employees only (including public sector). Observations are weighted by sampling weights multiplied by hours worked.

There are two significant differences. First, the mean log wage gap between college and high school workers is stable from 1998 to 2012 in PNAD, but increasing in RAIS. Second, variances of log wages within groups and for the whole sample are larger with the RAIS data for 1998, but not for 2012. Thus, RAIS shows larger reductions in inequality using this measure.

The first panel in Figure D3 replicates the evolution of schooling achievement of formal employees, shown in Figure 6. Again, the overall patterns are broadly similar: there is a substantial decline in the share of hours supplied by workers without any educational degree, accompanied by a similarly large increase in the percentage of hours supplied by workers with complete high school (secondary). There is also an increase in hours supplied by workers with college degrees. There are small changes in the shares; in particular, the PNAD shows a higher fraction of college-educated workers.

Figure D2: Measures of wage dispersion, PNAD data, all workers



Notes: PNAD data, Rio Grande do Sul, Brazil. All employees (including public sector and informal sector). Observations are weighted by sampling weights multiplied by hours worked.

There are three reasons for differences between the PNAD and RAIS. First, the RAIS is a census of formal employees, while PNAD is a small sample of that population. While the latter is designed to be representative, it might under-sample some workers with very high or very low earnings. Second, RAIS data are reported by firms, while PNAD data are reported by workers. That might lead to differences if, e.g., workers with high wages under-report in the PNAD or firms misreport the education of workers. Third, there are differences in the primitive questions used to construct wages and years of schooling in each dataset. [De Negri et al. \(2001\)](#) compares PNAD data and RAIS data and provides a detailed account of those differences. The first two reasons suggest that, when assessing inequality trends in the formal sector, RAIS data are probably more reliable than PNAD data.

D.2.2 Inequality trends for the whole workforce

Figure D2 is constructed similarly to Figure D1 above, but the data includes both formal and informal workers. I use a broad definition of the informal sector that includes domestic and self-employed workers. There are no substantial changes in qualitative patterns once informal workers are taken into account. The amount of wage dispersion is higher for the whole sample than for the restricted sample, especially in the lower tail of the wage distribution. One possible candidate for these differences is the presence of the binding minimum wage.

Differences in schooling achievement between the formal sample and the full sample can be observed by comparing the first two panels in Figure D3. Formal sector workers are a selected subsample with higher education levels. However, trends for the whole sample are, again, similar to those obtained from the formal sample.

D.2.3 Changes in relative labor supply

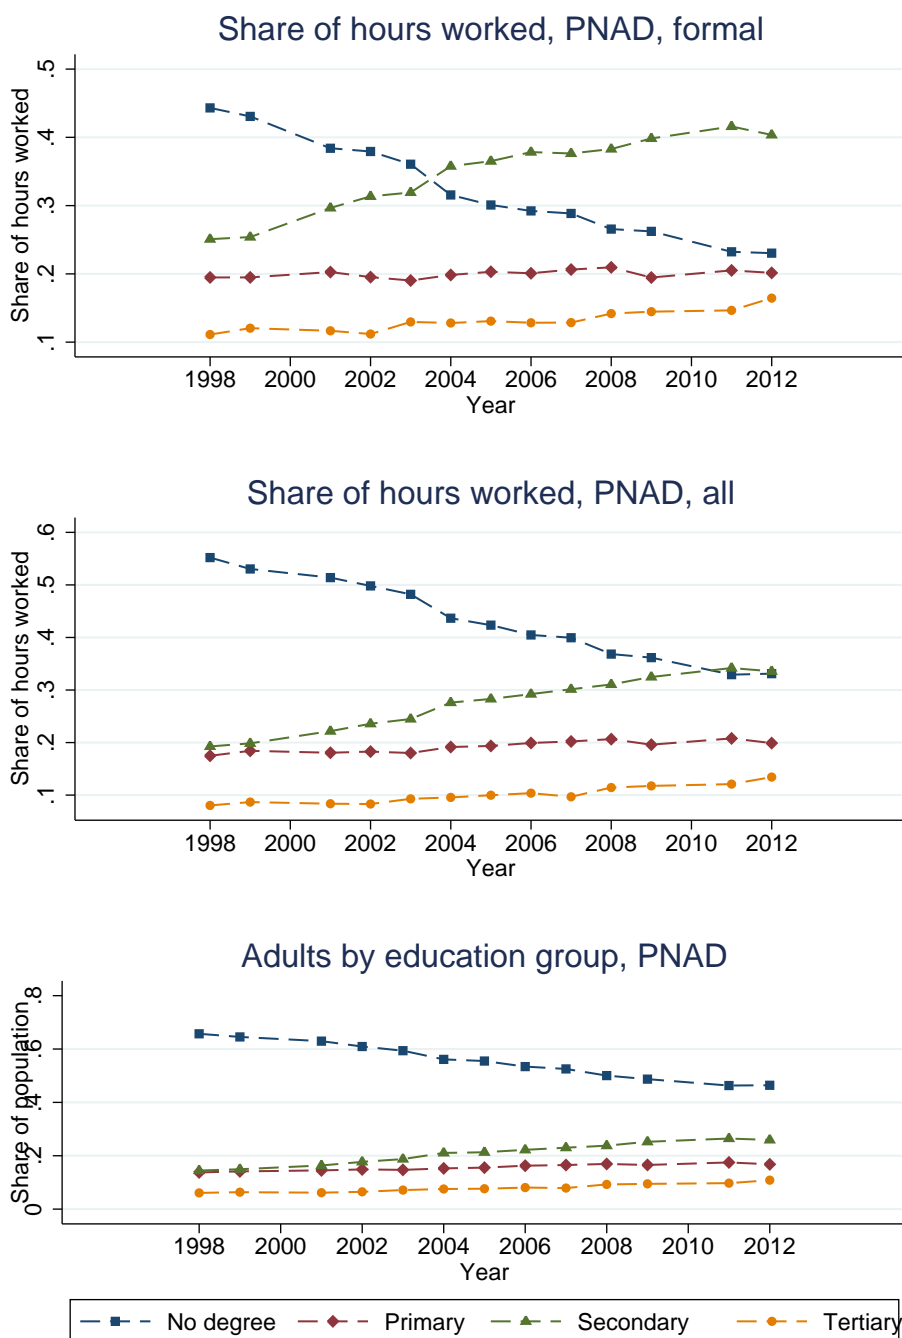
The first two panels of Figure D3, along with Figure 6 in the main text, show shares of hours worked supplied by each schooling group. One might wonder whether these could reflect changes in selection patterns into employment over time (coming, e.g., from business cycle fluctuations) instead of changes in labor supply. The third panel in Figure D3 shows that this is not the case. That graph shows the share of adults out of school, aged 18 through 54, in each educational group — regardless of whether they are employed, looking for jobs, or not in the labor force. The changes in educational achievement from that figure are similar in magnitude to those in the second and first panels. The levels are different, though, suggesting selection into employment by education.

D.3 Variance decomposition using Kline, Saggio and Sølvssten (2018)

The estimation of variance components follows the methodology proposed in Kline, Saggio and Sølvssten (2018), henceforth KSS. For the 1998 period, I use data for two years: 1997 and 1999. I use non-consecutive years to increase the number of firm-to-firm transitions.

The sample used for estimation is the largest leave-one-out connected set. This concept differs from the usual connected set in matched employer-employee datasets because it requires that firms need to be connected by at least two movers, such that removing any worker from the sample does not disconnect this set. Table D2 presents the size of that largest connected set in each period.

Figure D3: Changes in educational achievement, PNAD data



Notes: PNAD data, Rio Grande do Sul, Brazil. In the first two panels, the sample includes employed workers and observations are weighted by sampling weights multiplied by hours worked. In the third panel, the sample is composed of all adults 18-54 who are not in school, weighted by sampling weights.

Table D2: Sample sizes in variance decomposition exercise

	1998	2012
Person-year observations	1,618,478	2,570,016
Workers	809,239	1,285,008
Firms	31,107	65,466
Movers	174,299	334,206

This table presents the number of observations in the largest leave-one-out connected set.

I implement the variance decomposition using the code provided by KSS.²⁹ There are some implementation choices required in this estimation, stated below:

- Dealing with controls (year fixed effects): "Partialled out" prior to estimation (option 1 in the `resid_controls` argument).
- Computation of local linear regressions: stratified by grids, separate for movers and stayers (option 2 in the `subsample_llr_fit` argument).
- Sample selection: includes both movers and stayers (option 0 in the `restrict_movers` argument).
- Algorithm: Random projection method (option "JLL" in `type_of_algorithm` option, with `epsilon=0.005`).

D.4 Minimum Distance Estimation

D.4.1 Covariance matrix of estimated moments

I use the bootstrap to obtain the covariance matrix of all moments except those from the AKM decomposition. I re-sample workers in each period (keeping their histories in the two-period panel intact) and use a total of 500 replications. The KSS procedure provides standard errors for the variance of firm effects and the covariance of worker and firm effects. I assume that these estimates are uncorrelated between each other and with the other moments, filling the corresponding rows and columns in the covariance matrix with zeros.

²⁹Currently available at <https://github.com/rsaggio87/LeaveOutTwoWay>.

D.4.2 Simulating measures of wage inequality

After solving for equilibrium for a given set of parameters, I use the following procedure to simulate data which will be used to compute measures of inequality and the AKM decomposition:

1. For each $h = 1, \dots, H$:
 - (a) Create a sorted list of $2 \times G + 1$ values of log efficiency units, corresponding to (i) log hiring thresholds for all g ; (ii) points where the minimum wage ceases to bind, $\log(\underline{y}/w_{hg})$ for all g ; and (iii) quantile 0.999 of the $\log \varepsilon$ distribution.
 - (b) For each interval in $\log \varepsilon$ space, corresponding to consecutive values in that list:
 - i. If no firms pay the minimum wage in that interval, split it into 200 equal segments. Else, if some firms pay more than the minimum wage, split it into 40 segments. Else (all firms pay the minimum wage), use a single segment corresponding to the whole interval.
 - ii. For each $\log \varepsilon$ corresponding to a segment edge, calculate log earnings at each firm type and share of workers of that (h, ε) employed by firms of each type.
 - iii. Each segment/firm type combination corresponds to discretized group of workers, with the employment share for each g defined as the mean of the two shares at the edges of the segment, log earnings as the mean of the log earnings values at the edges, and quantity of workers defined by the mean employment share across g multiplied by N_h and by the share of workers of that h in that particular ε interval.

Using that simulated data to calculate means and variances of log earnings by group (and for the whole workforce) is straightforward.

D.4.3 Simulating the AKM decomposition

To reason about AKM decompositions in the theory, I need a two-period version of the model, from which panel data could be simulated if needed. I assume that, with some probability $R > 0$, workers re-draw their full vector of idiosyncratic preferences η_i from period one to period two. I also assume that only part of the efficiency units of labor of a worker is transferable $\log \varepsilon_{t=2} = A \log \varepsilon_{t=1} + (1 - A^2)^{0.5} \log \varepsilon'$, where ε' is a new i.i.d. draw from the

same distribution of efficiency units (given h). After the re-draws, the labor market clears in the same way as in period 1.

Because the cross-sectional distribution of (h, ε, η) remains the same as before, firm choices and the equilibrium allocation remain the same, except for the identity of workers employed by each firm. That model of job-to-job transitions implies that, whenever a given worker type (h, ε) is employed in equilibrium by the two firm types, there is a positive probability that some of those workers moved from a firm of type $g = 1$ to another of type $g = 2$ (and vice-versa).

Furthermore, I assume that firms are large, in the sense that there are many movers and firm fixed effects in the AKM regression are precisely estimated. Together with Lemma 2, that assumption implies that all firms producing the same good will have the same estimated fixed effect.

Given these assumptions, the results of an AKM decomposition of log wages using simulated panel data are identical to the alternative regression proposed in the main text. Each observation is a (h, ε, g) cell, using the discretized distribution of ε discussed above. There is a worker type dummy for each (h, ε) pair. The regression of log wages on worker type and firm type dummies is weighted by the share of the employed population in the corresponding cell. Finally, the estimated worker fixed effects are shrunk by the factor A , since they correspond only to the portable portion of productivity. The persistence parameter A is calibrated such that the R^2 of the simulated AKM regression is 0.9, about the same as the empirical regressions.³⁰

This approach ignores granularity issues in the simulation of AKM moments. That is conceptually consistent with the way the corresponding moments are estimated from the data, since the KSS estimator is not subject to limited mobility bias.

D.4.4 Numerical implementation and starting points

The estimation procedure is coded in Matlab. The objective function is defined in terms of transformed variables (e.g., log or logit-like functions) that ensure parameters are valid without the need for constrained optimization. When evaluating the objective function, the tolerance for the computation of equilibria is set to 10^{-14} . For minimization, I use a standard gradient descent approach, with the gradient calculated using forward differences shifts of

³⁰The persistence parameter is allowed to change between 1998 and 2012.

10^{-6} in the transformed variables. I used 126 starting points with random shifts of plus or minus 0.5 in all transformed variables, relative to a reasonable starting point. Only 2 did not converge to the optimal point.

D.5 Histograms of log wages, model and data

See Figure D4.

D.6 Minimum wage spillovers

The empirical model of minimum wage spillovers is:

$$\begin{aligned} \log y_{st}(p) - \log y_{st}(50) = & \beta_1(p) [\log \underline{y}_t - \log y_{st}(50)] + \beta_2(p) [\log \underline{y}_t - \log y_{st}(50)]^2 \\ & + \zeta_{0s}(p) + \zeta_{1s}(p) \times time_t + \zeta_2(p) \times (time_t)^2 + u_{st}(p) \end{aligned} \quad (16)$$

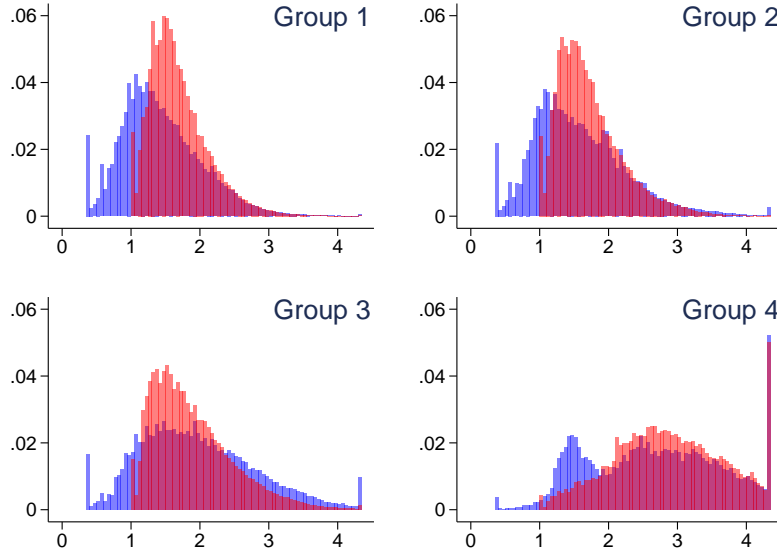
where $y_{st}(p)$ is the p -th percentile of the real wage distribution in state s at time t ; \underline{y}_t is the national minimum wage at time t ; $\zeta_{0s}(p)$ and $\zeta_{1s}(p)$ are state-quantile fixed effects and linear trends, respectively; $\zeta_2(p)$ is a national quadratic trend; and $u_{st}(p)$ is the residual.

This expression parameterizes the impact of the "effective minimum wage" $\underline{y}_t - \log y_{st}(50)$ — the minimum wage relative to the median wage in any given state and year — on any quantile p of the wage distribution, again relative to the median. The quadratic specification accounts for possibly non-linear effects of the effective minimum wage. The regression includes state-percentile fixed effects and linear trends to account for state-level changes in the shape of the wage distribution that are unrelated to the minimum wage. It also includes a national quadratic trend for each percentile, accounting for flexible changes in the shape of the wage distribution that are common across states. I use this trend instead of year effects, as in [Autor, Manning and Smith \(2016\)](#), because the statutory minimum wage is set at the federal level in Brazil.

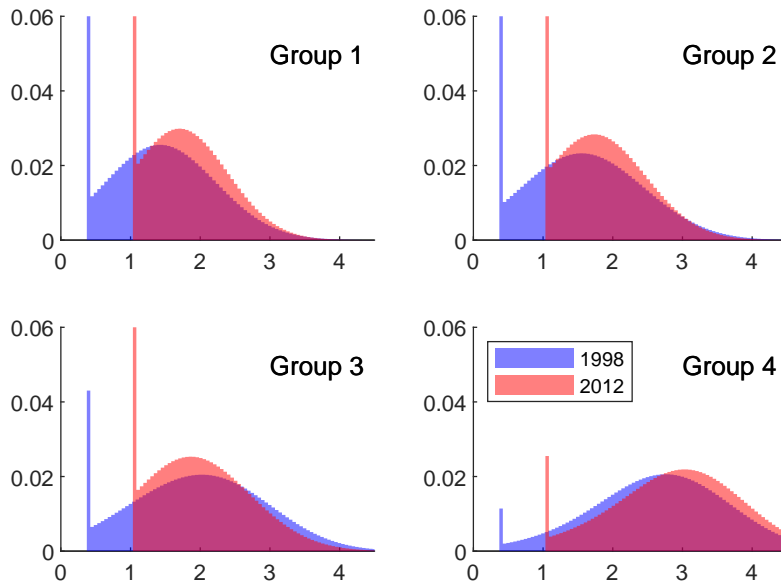
[Autor, Manning and Smith \(2016\)](#) argue that the effective minimum might correlate with the residual term because median wages are used to construct both the independent and the dependent variables. I follow their approach to solve this problem. Specifically, I use an instrument set composed of the log real minimum wage, the square of the log real minimum wage, and an interaction of the log real minimum wage with the average median real wage in state s for the whole period.

Figure D4: Distribution of log wages, data and model

(a) Data



(b) Simulation from estimated model



Notes: This figure shows histograms of log wages using 0.05-sized bins, separately by educational group (No degree, Primary, Secondary, and Tertiary) and time (1998 in blue, 2012 in red). Panel (a) shows data from RAIS, Rio Grande do Sul, Brazil, hours-weighted. Panel (b) shows histograms predicted by the estimated model.

Table D3 shows ordinary least squares and instrumental variables estimates of the marginal effect of minimum wages over different quantiles of the wage distribution. I estimate specifications in levels and in differences. The specification in differences presents much stronger first stages (measured by the Cragg-Donald (1993) F statistic). In addition, it shows no spillovers in the upper tail, a criterion that has been used for model selection when studying the impact of minimum wages on the wage distribution (e.g. Autor, Katz and Kearney (2008) and Cengiz et al. (2019)). For these reasons, it is my preferred specification.

These estimates, which are plotted in Figure 8, show spillovers that are economically and statistically significant up to percentile 40. Spillovers on the upper tail are small and indistinguishable from zero. These estimates are larger than what Autor, Manning and Smith (2016) found for the US, consistent with the fact that the minimum wage is more binding in Brazil and that only a small fraction of the workforce is in possession of a tertiary education.³¹

Model-based spillovers are calculated by making changes to the minimum wage, solving for equilibrium, and dividing the change in the desired quantile (relative to the median) by the change in effective minimum wage. The value shown in the graph is the average of 50 estimated spillovers, with parameters changing smoothly from the 1998 levels to the 2012 levels. The 50 step-wise changes in the minimum wage are the same ones used in the decomposition of inequality trends and sorting.

³¹Engbom and Moser (2018) also estimate reduced-form estimates of minimum wage spillovers for Brazil and compare them to the predictions of a structural model.

Table D3: Reduced form estimates of minimum wage spillovers

Quantile	Levels		Differences	
	OLS	IV	OLS	IV
10	0.584 (0.062)	0.427 (0.068)	0.641 (0.050)	0.540 (0.052)
20	0.369 (0.049)	0.246 (0.043)	0.389 (0.036)	0.321 (0.029)
30	0.204 (0.073)	0.158 (0.054)	0.241 (0.034)	0.167 (0.022)
40	0.106 (0.032)	0.025 (0.031)	0.119 (0.029)	0.052 (0.023)
60	-0.051 (0.037)	0.044 (0.029)	-0.084 (0.041)	-0.019 (0.024)
70	0.091 (0.095)	0.259 (0.060)	-0.037 (0.059)	0.067 (0.053)
80	0.113 (0.108)	0.281 (0.088)	0.015 (0.079)	0.028 (0.063)
90	0.230 (0.085)	0.282 (0.093)	0.113 (0.073)	-0.011 (0.068)
N	378	378	351	351
Cragg-Donald F		11.50		43.24

Notes: Each cell in this table corresponds to the marginal effects of the "effective minimum wage" (log statutory minimum wage minus median log wage) on quantiles of the wage distribution relative to the median log wage, coming from separate (quantile-specific) regressions. Each observation is a state-year and the regression is weighted by total hours worked. All years from 1996 through 2013 are included except 2002, 2003, 2004 and 2010, years in which data is not available for some states. Marginal effects are calculated at the median effective minimum wage for the whole sample (hours weighted). Regressions in levels include state fixed effects, state linear trends, and a national quadratic trend. Regressions in differences include state fixed effects and a national linear trend. Standard errors are clustered by state (27 clusters).