

Changes in Between-group Inequality: Computers, Occupations, and International Trade*

Ariel Burstein
UCLA

Eduardo Morales
Princeton University

Jonathan Vogel
Columbia University

June 6, 2018

Abstract

We provide a unifying framework to quantify the impact of several determinants of changes in U.S. between-group inequality. We use an assignment framework with many labor groups, equipment types, and occupations in which changes in inequality are driven by changes in workforce composition, occupation demand, computerization, and labor productivity. We parameterize the model using direct measures of computer usage within labor group-occupation pairs and quantify the impact of each shock for various dimensions of between-group inequality between 1984 and 2003. We find, for example, that computerization and shifts in occupation demand jointly account for roughly 80 percent of the rise in the skill premium, with computerization alone accounting for roughly 60 percent. In an open economy extension of the model, we show how computerization and changes in occupation demand can be caused by changes in the extent of international trade and perform counterfactual exercises to quantify these effects.

*We thank Treb Allen, Costas Arkolakis, David Autor, Gadi Barlevy, Lorenzo Caliendo, Davin Chor, Arnaud Costinot, Jonathan Eaton, Pablo Fajgelbaum, Gene Grossman, David Lagakos, Bernard Salanié, Nancy Stokey, Mike Waugh, and the referees for helpful comments. Vogel thanks the Princeton International Economics Section for their support. A previous version of this paper circulated under the name "Accounting for Changes in Between-group Inequality."

1 Introduction

The last few decades have witnessed pronounced changes in relative average wages across groups of workers with different observable characteristics (*between-group inequality*). Most notably, the wages of workers with more education relative to those with less and of women relative to men have increased substantially in the United States.

A large literature has emerged studying how changes in relative supply and demand for labor groups shape their relative wages. Changes in relative demand across labor groups have been linked prominently to computerization (or a reduction in the price of equipment more generally)—see e.g. [Krusell et al. \(2000\)](#), [Autor and Dorn \(2013\)](#), and [Beaudry and Lewis \(2014\)](#)—and to changes in relative demand across occupations and sectors, driven by structural transformation, offshoring, and international trade—see e.g. [Autor et al. \(2003\)](#), [Buera et al. \(2015\)](#), and [Galle et al. \(2015\)](#). Suggestive of the first hypothesis, Panel A of Table 1 shows that computer use in the United States rose dramatically between 1984 and 2003 and that computers are used more intensively by educated workers and women. Suggestive of the second hypothesis, Figure 1 shows that education- and female-intensive occupations grew relatively quickly over the same time period.¹

			Computer use over time				
			1984	1989	1993	1997	2003
Panel A.	All		27.2	39.7	49.1	53.0	58.0
	Gender	Female	32.7	47.4	56.8	60.9	65.1
		Male	23.3	34.0	43.2	46.8	52.3
	Education	College degree	45.8	62.8	74.6	79.9	85.8
		No college degree	21.9	32.3	40.4	42.8	45.4
Panel B.	Gender	Difference	9.4	13.5	13.6	14.1	12.9
		Contribution within	27.6	18.6	15.6	20.6	34.2
		Contribution between	72.4	81.4	84.4	79.4	65.8
	Education	Difference	23.9	30.4	34.2	37.1	40.4
		Contribution within	51.3	48.3	48.6	43.4	54.3
		Contribution between	48.7	51.7	51.4	56.6	45.7

Table 1: Panel A: Share of hours worked with computers by all workers, gender, and education. Panel B: The difference across gender and education in computer use, and the share of these differences accounted for by between- and within-occupation differences; see Appendix F.1 for details of the decomposition.

In this paper, we provide a model-based quantification of the impacts of various shocks on between-group inequality in the U.S. To motivate our approach, consider first

¹We describe our data sources in Section 4.1 and Appendix B.




Figure 1: Growth (1984-2003) of the occupation share of labor payments and the average (1984 & 2003) of the share of workers in the occupation who have a college degree (left) and are female (right)

the facts presented in Panel A of Table 1—that educated and female workers use computers more and that computer usage has risen over time—and what they imply for the college- and female-bias of computerization. These patterns may be caused by (i) variation in computer usage across worker groups within occupations or (ii) variation in the allocation of worker groups across occupations that differ in computer intensity. In Panel B, we decompose the difference in computer usage between women and men and between college and non-college workers into a within-occupation component and a between-occupation component. Women use computers more largely because they are disproportionately employed in computer-intensive occupations whereas more-educated workers use computers more both because of the within- and the between-occupation components (roughly in equal parts). Computerization raises the marginal product of worker groups that use computers intensively but reduces the price of output in occupations that are intensive in computers. Since wages equal the value marginal product of labor, the impact of computerization on wage inequality depends crucially both on variation in aggregate computer usage across worker groups—Panel A—and on the source of this variation—Panel B. Consequently, in order to quantify the impact of computerization on inequality, we must take into consideration the possibility that workers differ in their computer usage both within and across occupations.

Next, consider the facts presented in Figure 1—that college- and female-intensive occupations grew over time—and what they imply for the college- and female-bias of changes in occupation demand. Occupation growth is shaped not only by changes in relative demand shifters for different occupations, but also by changes in the wages and productivities of different groups of workers and changes in the prices and productivities of different types of capital equipment. For instance, if occupations are gross substitutes,

then computerization will tend to shift employment to computer-intensive occupations. Consequently, in order to quantify the impact of occupation demand shifters, we must take into consideration the possibility that computerization and changes in labor supply and productivity shape the size of occupations.

Motivated by these observations, we build a model with many labor groups (to study changes in between-group inequality along various dimensions) that features many occupations and equipment types (to incorporate heterogeneous computer usage within and across occupations as well shocks to equipment productivity and occupation demand). Given the difficulties in estimating a general production function with a large number of factors by occupation, we parameterize a Roy (1951)-type assignment model that builds on Lagakos and Waugh (2013) and Hsieh et al. (2016), extended to include equipment and international trade. This framework allows for potentially rich patterns of complementarity between labor groups, equipment types, and occupations and, in spite of its high dimensionality, remains tractable enough to perform aggregate counterfactuals in a parsimonious manner. Three-dimensional comparative advantage—between labor groups, equipment types, and occupations—shapes the allocation of workers in each labor group across equipment types and occupations and, thus, also shapes the impact of changes in the economic environment on between-group inequality. Our model’s aggregate implications for relative wages nest those of workhorse models of between-group inequality in a closed economy, e.g. Katz and Murphy (1992) and Krusell et al. (2000), and in an open economy, e.g. Heckscher-Ohlin.

We use our model to assess the impact of computerization and changes in occupation demand, labor composition, and labor productivity on between-group inequality in the United States. We additionally provide simple counterfactual exercises in which we quantify the extent to which computerization and changes in occupation demand are driven by international trade in equipment goods and occupation output.

In our first exercise, we decompose observed changes in relative wages between labor groups in the US between 1984 and 2003 into changes in (price-adjusted) equipment productivity, occupation demand, labor supply, and labor productivity, in a closed economy. This time period featured large changes in between-group inequality. For instance, the skill premium rose by 15.1 (log) percent and rose by 7.2 percent more amongst the old than the middle aged; and the gender gap contracted by 13.3 percent; see tables in Section 4.4.² We infer changes in equipment productivity that result in computerization from

²The relative importance of between- and within-group inequality is an area of active research, and is beyond the scope of this paper. On the one hand based on U.S. data, Autor (2014) writes that “about two-thirds of the overall rise of earnings dispersion between 1980 and 2005 is proximately accounted for by the increased premium associated with schooling in general and postsecondary education in particular,” and

changes in the allocation of workers to computers within labor group-occupation pairs; focusing on changes within labor group-occupation pairs is important because aggregate computer usage will rise even in the absence of changes in equipment productivity if labor groups that have a comparative advantage using computers or occupations that have a comparative advantage with computers grow. We infer occupation demand shifters from changes in the allocation of workers to occupations within labor group-equipment type pairs and in labor income shares across occupations. Changes in labor composition are directly observed in the data. We measure labor productivity as a residual to match changes in the average wage of each labor group. Finally, we leverage plausibly exogenous variation and the structure of our model to identify key elasticities.

We find that computerization alone accounts for roughly 60 percent of all shocks that have had a positive impact on the skill premium (i.e. the relative wage of workers with a college degree to those without) between 1984 and 2003 and plays a similar role in explaining disaggregated measures of between-education-group inequality (e.g., the wage of workers with graduate training relative to the average wage). Our model's prediction is driven by the following three facts observed in the data. First, there has been a large rise in the share of workers using computers within labor group-occupation pairs, which our model interprets as a large increase in computer productivity (i.e. computerization).³ Second, more educated workers use computers within occupations relatively more than less educated workers, which our model interprets as educated workers having a comparative advantage with computers. This pattern of comparative advantage, together with computerization, yields a rise in the relative wages of educated workers according to our model. Third, more educated workers are also disproportionately employed in occupations in which all workers use computers relatively intensively, which our model interprets as educated workers having a comparative advantage in occupations in which computers have a comparative advantage. This pattern of sorting across occupations, together with computerization and an estimated elasticity of substitution between occupations greater than one, also yields a rise in the relative wages of educated workers.

The combination of computerization and changes in occupation demand accounts for roughly 80 percent of the rise in the skill premium, leaving only 20 percent to be explained by labor productivity. This is remarkable, given that we measure changes in labor pro-

Lemieux (2006) writes "that residual wage inequality accounts for only a modest share of the growth in overall inequality between 1973 and 2003." On the other hand, based on 1986-1995 Brazilian data Helpman et al. (2017) concludes that "residual wage inequality is at least as important as worker observables in explaining the overall level and growth of wage inequality."

³This observation is consistent with ample direct evidence showing a rapid decline in the price of computers relative to all other equipment types and structures (summarized in footnote 25 below) which we do not directly use in our estimation.

ductivity as a residual that allows our framework to match exactly observed changes in relative wages. We find that computerization, occupation shifters, and labor productivity all play important roles in accounting for the reduction in the gender gap (i.e. the relative wage of male to female workers). Computerization decreases the gender gap because women are disproportionately employed in occupations in which all workers use computers intensively and our estimate of the elasticity of substitution across occupations is larger than one. We show that computerization plays the central role in explaining the distinct evolution of the skill premium across age groups, while labor composition—the focus of previous literature on this issue, e.g. [Card and Lemieux \(2001\)](#)— plays a much smaller role. We obtain this result because computer use among college-educated relative to non-college-educated workers is substantially higher among young workers than old workers in every year of our sample. More generally, we show that computerization explains over 50 percent of the variance in the change in relative wages (implied by the combination of the three demand-side shocks) across the thirty labor groups considered in our analysis.

Whereas in our closed economy model we treat computerization and occupation demand shifters as primitives, in sections 5 and 6 we study the extent to which these changes are a consequence of international trade. Given estimates of the relevant elasticity parameters, we show theoretically that the procedure to quantify the impact on relative wages of moving to autarky in equipment and occupation trade is equivalent to the procedure we follow in our closed economy model to calculate the impact of domestic shocks on relative wages, with the only difference that the computerization and occupation demand shocks are now measured as functions of import shares of absorption and export shares of output of different equipment types and occupations. For example, if an occupation ω has a high import share relative to another occupation ω' , then moving to autarky has an equivalent impact on relative wages in a closed economy as increasing domestic occupation demand for ω relative to ω' . Given the lack of data on the occupation content of exports and imports, measuring occupation trade shares is a challenge; for a full discussion of the difficulties, see [Grossman and Rossi-Hansberg \(2007\)](#). Given these challenges, we consider alternative simple, albeit imperfect, approaches to measure the occupation content of exports and imports. Using our preferred approach, we find that moving from 2003 trade shares to equipment (occupation) autarky would generate a 2.2 (6.3) percent reduction in the skill premium. We also provide a simple procedure to quantify the differential effects on wages in a given country of changes in primitives (i.e. worldwide technologies, labor compositions, and trade costs) between two time periods relative to the effects of the same changes in primitives if that country were a closed economy. Our results indicate that

changes in equipment (occupation) trade between 1984 and 2003 accounted for roughly 15 percent (25 percent) of the impact of changes in equipment productivity (occupation shifters) on the skill premium that we obtained in our closed economy calculations.

Our paper is organized as follows. In Section 2, we discuss the related literature. We describe our closed economy framework, characterize its equilibrium, and discuss its mechanisms in Section 3. We parameterize the model and present our closed economy results in Section 4. We extend our model to incorporate international trade in equipment goods and occupation services in Section 5 and quantify the impact of international trade in Section 6. We conclude in Section 7. Additional details and robustness exercises are relegated to appendices.

2 Literature

In trying to explain the evolution of between-group inequality as a function of changes in observables, our paper's objective is most similar to [Krusell et al. \(2000\)](#) and [Lee and Wolpin \(2010\)](#). [Krusell et al. \(2000\)](#) estimate an aggregate production function that permits capital-skill complementarity and show that changes in aggregate stocks of equipment, skilled labor, and unskilled labor can account for much of the variation in the U.S. skill premium. Whereas [Krusell et al. \(2000\)](#) identify the degree of capital-skill complementarity using aggregate time series data, our approach leverages information on the allocation of workers to computers and occupations and, consequently, yields parameter estimates shaping the degree of equipment-labor group complementarities that are robust to allowing for time trends in the relative productivity of each labor group; see [Acemoglu \(2002\)](#) for a discussion of the relevance of allowing for these time trends in this context. Our decomposition corroborates the findings in [Krusell et al. \(2000\)](#) and extends them by additionally considering the impact of equipment productivity growth on the gender gap and other measures of between-group inequality.

[Lee and Wolpin \(2010\)](#) use a dynamic model of endogenous human capital accumulation and find that labor group productivity (also treated as a residual in their analysis) plays the central role in explaining changes in the skill premium. By considering a greater degree of disaggregation in occupations and labor groups (which would require estimating hundreds of additional elasticities using their aggregate-production-function approach, but which is straightforward in our assignment approach), and by introducing three-dimensional comparative advantage (in order to account for the patterns in Panel B of Table 1, which cannot be incorporated into an aggregate production function), our results largely reduce the role of changes in the residual in shaping changes in the skill

premium and allow us to study various additional dimensions of between-group inequality. On the other hand, in contrast to [Lee and Wolpin \(2010\)](#), we treat labor composition as exogenous.⁴

In modeling international trade, we operationalize in a quantitative setting the theoretical insights of [Costinot and Vogel \(2010\)](#), [Sampson \(2014\)](#), and [Costinot and Vogel \(2015\)](#) regarding the impact of international trade on inequality in high-dimensional environments. We show that one can use a similar approach to that introduced by [Dekle et al. \(2008\)](#) in a single-factor trade model—i.e. replacing a large number of unknown parameters with observable allocations in an initial equilibrium—in a many-factor assignment model. In this respect, our paper is complementary to concurrent work quantifying the impact of international trade on between-group inequality; see e.g. [Adao \(2015\)](#), [Dix-Carneiro and Kovak \(2015\)](#), [Galle et al. \(2015\)](#), and [Lee \(2016\)](#). Relative to this concurrent work, we introduce equipment in the framework and quantify the impact of trade both in equipment and in occupations on inequality; however, relative to this work, we only use data for one region: the US. Whereas a large literature has emerged to argue that trade in occupation (or even task) output is a potentially important force generating changes in inequality—see e.g. [Grossman and Rossi-Hansberg \(2008\)](#)—there has been much less work conducting model-based counterfactuals to quantify the importance of this phenomenon.⁵ Given the crudeness of our measures of occupation trade, our quantification of the impact of trade in occupations on between-group inequality should be viewed as a first step. Our modeling of international trade in equipment extends the quantitative analyses of [Burstein et al. \(2013\)](#) and [Parro \(2013\)](#), who study the impact of trade in capital

⁴Extending our model to endogenize education and labor participation—maintaining a static environment—would give rise to the same equilibrium equations determining factor allocations and wages conditional on labor composition. Our measures of shocks—to occupation demand shifters, equipment productivity, and labor productivity—and our estimates of model parameters are therefore robust to extending our model to endogenize the supply of each labor group in a static environment. Through the lens of this model, our counterfactual results would be interpreted as the direct effect of shocks on labor demand and wages, taking labor composition as given. To take into consideration the accumulation of occupation-specific human capital as in, e.g., [Kambourov and Manovskii \(2009a,b\)](#), we would have to include occupational experience as a worker characteristic when defining labor groups in the data (unfortunately, the October CPS does not contain this information) and model the corresponding dynamic optimization problem that workers would solve when deciding which occupation to sort into.

⁵[Traiberman \(2016\)](#) quantifies the impact of import competition on occupation demand and wages. He constructs a measure of import exposure by occupation by allocating observed sectoral imports to occupations according to the share of labor payments going to each occupation in each sector. As we argue below, this measure is likely to be biased if, within industries, certain occupations are more traded than others. [Goos et al. \(2014\)](#) measure the impact of offshoring on the growth of occupations imposing a common effect of occupation offshorability on changes in occupation size, independently of whether a country is a net exporter or a net importer in this occupation. Our model however emphasizes the importance of distinguishing whether a country is a net exporter or importer of each occupation when looking at the impact of international trade on occupation growth.

equipment on the skill premium using the model of [Krusell et al. \(2000\)](#).

A large number of papers study the impact of changes in demand across sectors or occupations on relative wages using a shift-share analysis; see e.g. [Katz and Murphy \(1992\)](#) and [Autor et al. \(2003\)](#). This approach directly interprets [Figure 1](#) as evidence of the extent to which changes in sector or occupation demand have affected relative wages. We contribute to this literature by estimating the elasticity of substitution across occupations (a necessary condition for a causal interpretation of a shift-share analysis is that this elasticity is one) and incorporating a range of additional possible causes of changes in relative wages and occupation employment and income shares, which becomes necessary to quantify the impact of changes in sector or occupation demand when this elasticity differs from one.⁶

Two related papers use information on regional exposure to computerization to study the differential effect across regions of technical change on the polarization of U.S. employment and wages, [Autor and Dorn \(2013\)](#), and on the gender gap and skill premium, [Beaudry and Lewis \(2014\)](#). Our approach complements these papers, embedding computerization into a general equilibrium model that allows us to quantify at the aggregate level the effect of computerization (as well as other shocks) on changes in between-group inequality (as opposed to the differential effect across regions). By incorporating three-dimensional comparative advantage, our framework introduces heterogeneous treatment effects of computerization across labor groups that use computers at a similar rate.

In exploiting data on workers' computer usage, our paper is related to an earlier literature studying the impact of computer use on wages; see e.g. [Krueger \(1993\)](#) and [Entorf et al. \(1999\)](#). This literature identifies the impact of computer usage on wages by regressing wages of different workers on a dummy for computer usage, an identification approach that [DiNardo and Pischke \(1997\)](#) criticize. Our approach to estimate key model parameters does not rely on such a regression.

3 Closed Economy Model

In this section, we introduce the closed economy version of our model, characterize its equilibrium, and provide intuition for how different changes in the economic environment affect relative wages.

⁶The approach that we use to bring our model to the data does not require mapping occupations into observable characteristics such as those introduced in [Autor et al. \(2003\)](#). Instead, we estimate measures of comparative advantage and demand shocks that vary flexibly across occupations, independently of the similarity in the task composition of these occupations.

3.1 Environment

At time t there is a continuum of workers indexed by $z \in \mathcal{Z}_t$, each of whom inelastically supplies one unit of labor. We divide workers into a finite number of labor groups, indexed by λ . The set of workers in group λ is given by $\mathcal{Z}_t(\lambda) \subseteq \mathcal{Z}_t$, which has mass $L_t(\lambda)$. There is a finite number of equipment types, indexed by κ . Workers and equipment are employed by production units to produce a finite number of occupations, indexed by ω .

A final good is produced combining the services of occupations according to a constant elasticity of substitution (CES) production function

$$Y_t = \left(\sum_{\omega} \mu_t(\omega)^{1/\rho} Y_t(\omega)^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)}, \quad (1)$$

where $\rho > 0$ is the elasticity of substitution across occupations, $Y_t(\omega) \geq 0$ is the output of occupation ω , and $\mu_t(\omega) \geq 0$ is an exogenous demand shifter for occupation ω .⁷ The final good is used to produce consumption, C_t , and equipment, $Y_t(\kappa)$, according to the resource constraint

$$Y_t = C_t + \sum_{\kappa} \tilde{p}_t(\kappa) Y_t(\kappa), \quad (2)$$

where $\tilde{p}_t(\kappa)$ denotes the cost of a unit of equipment κ in terms of units of the final good.⁸

Occupation services are produced by perfectly competitive production units. A unit hiring k units of equipment type κ and l efficiency units of labor group λ produces $k^\alpha [T_t(\lambda, \kappa, \omega) l]^{1-\alpha}$ units of output, where α denotes the output elasticity of equipment and $T_t(\lambda, \kappa, \omega)$ denotes the productivity of an efficiency unit of group λ 's labor in occupation ω when using equipment κ .⁹ Comparative advantage between labor and equipment is defined as follows: λ' has a comparative advantage (relative to λ) using equipment κ' (relative to κ) in occupation ω if $T_t(\lambda', \kappa', \omega) / T_t(\lambda', \kappa, \omega) \geq T_t(\lambda, \kappa', \omega) / T_t(\lambda, \kappa, \omega)$.

⁷We show in sections 5 and 6 that changes in the extent of international trade in occupation services generate changes in $\mu_t(\omega)$.

⁸We show in sections 5 and 6 that changes in the extent of international trade in equipment generate changes in equipment prices $\tilde{p}_t(\kappa)$. Our assumption that equipment must be produced every period (five-years in our quantitative analysis) is equivalent to assuming that equipment fully depreciates every period. Alternatively, we could have assumed that $Y_t(\kappa)$ denotes investment in capital equipment κ , which depreciates at a given finite rate. All our counterfactual exercises are consistent with this alternative model with capital accumulation: they would correspond to comparisons across balanced growth paths in which the real interest rate and the growth rate of relative productivity across equipment types are constant over time.

⁹We can extend the model to incorporate other inputs such as structures or intermediate inputs s ; however, they would not affect any of our results as long as they are produced linearly using the final good and enter the production function multiplicatively as $s^{1-\eta} (k^\alpha [T_t(\lambda, \kappa, \omega) l]^{1-\alpha})^\eta$. Notice that, in either case, α is the share of equipment relative to the combination of equipment and labor. We restrict α to be common across occupations because we do not have the data to estimate $\alpha(\omega)$ by occupation.

Labor-occupation and equipment-occupation comparative advantage are defined analogously.

A worker $z \in \mathcal{Z}_t(\lambda)$ supplies $\epsilon(z) \times \varepsilon(z, \kappa, \omega)$ efficiency units of labor if teamed with equipment κ in occupation ω . Each worker is associated with a unique $\epsilon(z)$, allowing some workers within $\mathcal{Z}_t(\lambda)$ to be more productive than others across all possible (κ, ω) ; we normalize the average value of $\epsilon(z)$ across workers within each λ to be one, which is without loss of generality (see Appendix A). Each worker is also associated with a vector of $\varepsilon(z, \kappa, \omega)$, one for each (κ, ω) pair, allowing workers within $\mathcal{Z}_t(\lambda)$ to vary in their relative productivities across (κ, ω) pairs. We impose two restrictions. First, the distribution of $\epsilon(z)$ has finite support and is independent of the distribution of $\varepsilon(z, \kappa, \omega)$ for each (κ, ω) . Second, each $\varepsilon(z, \kappa, \omega)$ is drawn independently from a Fréchet distribution with cumulative distribution function $G(\varepsilon) = \exp(-\varepsilon^{-\theta})$, where a higher value of $\theta > 1$ implies lower within-worker dispersion of efficiency units across (κ, ω) pairs.¹⁰

Total output of occupation ω , $Y_t(\omega)$, is the sum of output across all units producing occupation ω . All markets are perfectly competitive and all factors are freely mobile across occupations and equipment types.

Relation to alternative frameworks. Whereas our framework imposes strong restrictions on occupation production functions, its aggregate implications for wages nest those in [Katz and Murphy \(1992\)](#) and [Krusell et al. \(2000\)](#). Specifically, the aggregate implications of our model for relative wages are equivalent to those in [Katz and Murphy \(1992\)](#) if we assume no equipment (i.e. $\alpha = 0$) and two labor groups, each of which has a positive productivity in only one occupation. Similarly, the aggregate implications of our model for relative wages are equivalent to those in [Krusell et al. \(2000\)](#) if we allow for only two labor groups and one type of equipment, each labor group has positive productivity in only one occupation and the equipment share is positive in only one occupation.

3.2 Equilibrium

We characterize the competitive equilibrium, first taking occupation prices as given and then in general equilibrium. Additional derivations are provided in Appendix A.

Partial equilibrium. With perfect competition, equation (2) implies that the price of

¹⁰The assumption that $\varepsilon(z, \kappa, \omega)$ is distributed Fréchet is made for tractability, as it implies that the average wage and the allocation of workers (in each labor group) across occupations and equipment types are both globally constant elasticity functions of occupation prices and equipment productivities. The wage distribution implied by the Fréchet assumption is a good approximation to the observed distribution of individual wages; see e.g. [Saez \(2001\)](#) and Figure 3 in Appendix D.3. In Appendix E.4 we show that our results are robust to allowing for variation across labor groups in the dispersion parameter θ , restricting the statistical dependence of $\varepsilon(z, \kappa, \omega)$ across (κ, ω) pairs by using a multivariate Fréchet distribution.

equipment κ is simply $p_t(\kappa) = \tilde{p}_t(\kappa) P_t$, where P_t is the price of the final good, which we normalize to one so that $p_t(\kappa) = \tilde{p}_t(\kappa)$. An occupation production unit hiring k units of equipment κ and l efficiency units of labor λ earns revenues $p_t(\omega) k^\alpha [T_t(\lambda, \kappa, \omega) l]^{1-\alpha}$ and incurs costs $p_t(\kappa) k + v_t(\lambda, \kappa, \omega) l$, where $v_t(\lambda, \kappa, \omega)$ is the wage per efficiency unit of labor λ when teamed with equipment κ in occupation ω and $p_t(\omega)$ is the price of occupation ω output. The profit maximizing choice of equipment quantity and the zero profit condition—due to costless entry of production units—yield

$$v_t(\lambda, \kappa, \omega) = \bar{\alpha} p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} p_t(\omega)^{\frac{1}{1-\alpha}} T_t(\lambda, \kappa, \omega)$$

if there is positive entry in $(\lambda, \kappa, \omega)$, where $\bar{\alpha} \equiv (1 - \alpha) \alpha^{\alpha/(1-\alpha)}$. Facing the wage profile $v_t(\lambda, \kappa, \omega)$, each worker $z \in \mathcal{Z}_t(\lambda)$ chooses the equipment-occupation pair (κ, ω) that maximizes her wage, $\epsilon(z) \varepsilon(z, \kappa, \omega) v_t(\lambda, \kappa, \omega)$. The assumption that $\varepsilon(z, \kappa, \omega)$ is distributed Fréchet and independent of $\epsilon(z)$ implies that the probability that a randomly sampled worker, $z \in \mathcal{Z}_t(\lambda)$, uses equipment κ in occupation ω is

$$\pi_t(\lambda, \kappa, \omega) = \frac{\left[T_t(\lambda, \kappa, \omega) p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} p_t(\omega)^{\frac{1}{1-\alpha}} \right]^\theta}{\sum_{\kappa', \omega'} \left[T_t(\lambda, \kappa', \omega') p_t(\kappa')^{\frac{-\alpha}{1-\alpha}} p_t(\omega')^{\frac{1}{1-\alpha}} \right]^\theta}. \quad (3)$$

The higher is θ , the more responsive are factor allocations to changes in prices or productivities. According to equation (3), comparative advantage shapes factor allocations. As an example, the assignment of workers across equipment types within any given occupation satisfies

$$\frac{T_t(\lambda', \kappa', \omega)}{T_t(\lambda', \kappa, \omega)} \bigg/ \frac{T_t(\lambda, \kappa', \omega)}{T_t(\lambda, \kappa, \omega)} = \left(\frac{\pi_t(\lambda', \kappa', \omega)}{\pi_t(\lambda', \kappa, \omega)} \bigg/ \frac{\pi_t(\lambda, \kappa', \omega)}{\pi_t(\lambda, \kappa, \omega)} \right)^{1/\theta},$$

so that, if λ' workers (relative to λ) have a comparative advantage using κ' (relative to κ) in occupation ω , then they are relatively more likely to be allocated to κ' in occupation ω .

The wage per efficiency unit of labor λ when teamed with equipment κ in occupation ω , $v_t(\lambda, \kappa, \omega)$, differs from the *average wage* of workers in group λ teamed with equipment κ in occupation ω , denoted by $w_t(\lambda, \kappa, \omega)$, which is the integral of $\epsilon(z) \varepsilon(z, \kappa, \omega) v_t(\lambda, \kappa, \omega)$ across workers teamed with κ in occupation ω , divided by the mass of these workers. Given our assumptions on $\epsilon(z)$ and $\varepsilon(z, \kappa, \omega)$, we obtain

$$w_t(\lambda, \kappa, \omega) = \bar{\alpha} \gamma T_t(\lambda, \kappa, \omega) p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} p_t(\omega)^{\frac{1}{1-\alpha}} \pi_t(\lambda, \kappa, \omega)^{-1/\theta},$$

where $\gamma \equiv \Gamma(1 - 1/\theta)$ and $\Gamma(\cdot)$ is the Gamma function. An increase in productivity or occupation price, $T_t(\lambda, \kappa, \omega)$ or $p_t(\omega)$, or a decrease in equipment price, $p_t(\kappa)$, raises the wages of infra-marginal λ workers allocated to (κ, ω) . However, the average wage across all λ workers in (κ, ω) increases less than that of infra-marginal workers due to self-selection, i.e. $\pi_t(\lambda, \kappa, \omega)$ increases, which lowers the average efficiency units of the λ workers who choose to use equipment κ in occupation ω .

Denoting by $w_t(\lambda)$ the average wage of workers in group λ (i.e. total income of the workers in group λ divided by their mass), the previous expression and equation (3) imply $w_t(\lambda) = w_t(\lambda, \kappa, \omega)$ for all (κ, ω) , where¹¹

$$w_t(\lambda) = \bar{\alpha} \gamma \left(\sum_{\kappa, \omega} \left(T_t(\lambda, \kappa, \omega) p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} p_t(\omega)^{\frac{1}{1-\alpha}} \right)^\theta \right)^{1/\theta}. \quad (4)$$

General equilibrium. In any period, occupation prices $p_t(\omega)$ must be such that total expenditure in occupation ω is equal to total revenue earned by all factors employed in ω ,

$$\mu_t(\omega) p_t(\omega)^{1-\rho} E_t = \frac{1}{1-\alpha} \zeta_t(\omega), \quad (5)$$

where $\zeta_t(\omega) \equiv \sum_{\lambda, \kappa} w_t(\lambda) L_t(\lambda) \pi_t(\lambda, \kappa, \omega)$ is total labor income in occupation ω , and $E_t \equiv \sum_{\omega} \zeta_t(\omega)$ is total expenditures. The aggregate quantity of the final good is

$$Y_t = E_t = (1 - \alpha)^{-1} \sum_{\lambda} w_t(\lambda) L_t(\lambda),$$

the aggregate quantity of equipment κ is

$$Y_t(\kappa) = \frac{1}{p_t(\kappa)} \frac{\alpha}{1-\alpha} \sum_{\lambda, \omega} \pi_t(\lambda, \kappa, \omega) w_t(\lambda) L_t(\lambda),$$

and aggregate consumption is given by equation (2).

¹¹The model implies that changes in wages are common across occupations within a labor group λ . In Appendix G we decompose the observed *changes* in average wages for each labor group into a between occupation component (which is zero in our model) and a within occupation component. We show that the between component is relatively small in the data (accounting for approximately 15% of the variation in wage changes). In the presence of compensating differentials, average wages would vary across (κ, ω) pairs within each labor group λ .

3.3 System in Changes

To solve for changes in equilibrium wages in response to changes in the economic environment, we express the system of equilibrium equations described in Section 3.2 in changes, denoting changes in any variable x between any two periods t_0 and t_1 by $\hat{x} \equiv x_{t_1}/x_{t_0}$. Changes in average wages are given by

$$\hat{w}(\lambda) = \left(\sum_{\kappa, \omega} \left(\hat{T}(\lambda, \kappa, \omega) \hat{p}(\kappa)^{\frac{-\alpha}{1-\alpha}} \hat{p}(\omega)^{\frac{1}{1-\alpha}} \right)^\theta \pi_{t_0}(\lambda, \kappa, \omega) \right)^{1/\theta} \quad (6)$$

and changes in occupation prices are determined by the following system of equations

$$\hat{\pi}(\lambda, \kappa, \omega) = \frac{\left[\hat{T}(\lambda, \kappa, \omega) \hat{p}(\kappa)^{\frac{-\alpha}{1-\alpha}} \hat{p}(\omega)^{\frac{1}{1-\alpha}} \right]^\theta}{\sum_{\kappa', \omega'} \left[\hat{T}(\lambda, \kappa', \omega') \hat{p}(\kappa')^{\frac{-\alpha}{1-\alpha}} \hat{p}(\omega')^{\frac{1}{1-\alpha}} \right]^\theta \pi_{t_0}(\lambda, \kappa', \omega')} \quad (7)$$

$$\hat{\mu}(\omega) \hat{p}(\omega)^{1-\rho} \hat{E} = \frac{1}{\zeta_{t_0}(\omega)} \sum_{\lambda, \kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega), \quad (8)$$

where $\hat{E} = \sum_{\lambda} \frac{w_{t_0}(\lambda) L_{t_0}(\lambda)}{(1-\alpha) E_{t_0}} \hat{w}(\lambda) \hat{L}(\lambda)$. In this system, the forcing variables are: $\hat{L}(\lambda)$, $\hat{T}(\lambda, \kappa, \omega)$, $\hat{\mu}(\omega)$, and $\hat{p}(\kappa)$. Given these variables and a choice of a numeraire, equations (6)-(8) yield the model's implied values of changes in average wages, $\hat{w}(\lambda)$, allocations, $\hat{\pi}(\lambda, \kappa, \omega)$, occupation prices, $\hat{p}(\omega)$, and total expenditure, \hat{E} .

3.4 Intuition

The impact of shocks on wages, and the role of ρ . According to equation (6), changes in productivities $\hat{T}(\lambda, \kappa, \omega)$ and in the price of equipment $\hat{p}(\kappa)$ have a direct effect (i.e. holding occupation prices constant) on wages. In addition, both these shocks as well as changes in occupation shifters and in labor composition affect wages indirectly through their impact on occupation prices $\hat{p}(\omega)$. The importance of changes in productivities, equipment prices, and occupation prices for changes in group λ 's average wage depends on factor allocations in the initial period, $\pi_{t_0}(\lambda, \kappa, \omega)$. For example, an increase in occupation ω 's price or a reduction in the price of equipment κ between any two periods t_0 and t_1 raises the relative average wage of labor groups that disproportionately work in occupation ω or use equipment κ in period t_0 .

We now discuss in more detail the intuition for the impact of each of the shocks (the forcing variables) we consider in our quantitative analysis. Consider an increase

in occupation- ω demand, i.e. $\hat{\mu}(\omega) > 1$. This shock raises the price of occupation ω and, therefore, the relative wage of labor groups that are disproportionately employed in ω . Similarly, a decrease in the labor supply of group λ , i.e. $\hat{L}(\lambda) < 1$, raises the relative price of occupations in which group λ is disproportionately employed. This raises the relative wage not only of group λ , but also of other labor groups employed in similar occupations as λ . A proportional increase across all (κ, ω) in labor productivity for group λ —i.e. $\hat{T}(\lambda, \kappa, \omega) = \hat{T}(\lambda) > 1$ —directly raises the relative wage of group λ and reduces the relative price of occupations in which this group is disproportionately employed, thus reducing the relative wage of labor groups employed in similar occupations as λ .¹²

Consider the impact of a decrease in the price of equipment κ , i.e. $\hat{p}(\kappa) < 1$, or equivalently a proportional increase in its productivity across all (λ, ω) , i.e. $\hat{T}(\lambda, \kappa, \omega) = \hat{T}(\kappa) > 1$. The partial equilibrium impact of this shock is to raise the relative wage of labor groups that use κ intensively at the aggregate level. In general equilibrium this shock also reduces the relative price of occupations in which κ is used intensively, lowering the relative wage of labor groups that are disproportionately employed in these occupations. Overall, the impact on a group's average wage of changes in equipment price depends on the value of ρ and the extent to which the group's aggregate computer use is shaped by variation in computer usage within or across occupations.

If the only form of comparative advantage is between workers and equipment—in which case variation in aggregate computer use across labor groups is fully explained by within-occupation variation in computer use—then a decrease in the price of equipment κ does not affect relative occupation prices. Hence, a decrease in the price of κ will increase the relative wages of worker groups that use equipment κ more intensively in the initial period.

On the other hand, if the only forms of comparative advantage are between workers and occupations and between equipment and occupations—in which case variation in aggregate computer use is fully explained by between-occupation allocations—then a decrease in the price of equipment κ has two opposite effects on the relative wages of worker groups that use equipment κ intensively at the aggregate level: a positive effect conditional on occupation prices and a negative effect indirectly through its effect on occupation prices. The relative strength of the direct and indirect channels depends on ρ . The direct effect dominates if and only if $\rho > 1$. Intuitively, a decrease in the price of κ

¹²Costinot and Vogel (2010) provide analytic results on the implications for relative wages of changes in labor composition, $L_t(\lambda)$, and occupation demand, $\mu_t(\omega)$, in a restricted version of our model in which there are no differences in efficiency units across workers in the same labor group (i.e. $\theta = \infty$), there is no capital equipment (i.e. $\alpha = 0$), and $T(\lambda, \omega)$ —i.e. our $T(\lambda, \kappa, \omega)$ in the absence of equipment—is log-supermodular.

acts like a positive productivity shock to the occupations in which κ has a comparative advantage. If $\rho > 1$, this increases employment in these occupations as well as the relative wages of labor groups disproportionately employed in these occupations. On the other hand, if $\rho < 1$, employment falls in κ -intensive occupations, and relative wages fall disproportionately for labor groups employed in these occupations.

More generally—outside of these two extreme cases—a decrease in the price of computers can simultaneously raise the wage (relative to the average wage) of one labor group (λ_1) that uses computers intensively while reducing the wage (relative to the average wage) of another labor group (λ_2) that also uses computers intensively. This would be the case, for example, if λ_1 has a comparative advantage using computers, λ_2 has a comparative advantage in occupations in which computers have a comparative advantage, and $\rho < 1$. Moreover, our framework is flexible enough that computerization can either contract or expand employment in occupations that are computer intensive.

The role of θ . The parameter θ , which governs the degree of within-worker productivity dispersion across occupation-equipment type pairs, determines the extent of worker reallocation in response to changes in occupation prices, equipment prices, and productivities: a higher dispersion of idiosyncratic draws, as given by a lower value of θ , results in less worker reallocation. In order to gain intuition on the role of θ for changes in labor group average wages, we take a first-order approximation of equation (6) at period t_0 allocations:

$$\log \hat{w}(\lambda) = \sum_{\kappa, \omega} \pi_{t_0}(\lambda, \kappa, \omega) \left(\log \hat{T}(\lambda, \kappa, \omega) + \frac{1}{1-\alpha} \log \hat{p}(\omega) - \frac{\alpha}{1-\alpha} \log \hat{p}(\kappa) \right). \quad (9)$$

This expression shows that, given changes in occupation and equipment prices as well as productivities, the change in average wages does not depend on θ up to a first-order approximation.¹³ However, a lower value of θ results in less worker reallocation across occupations in response to a shock and, therefore, larger changes in occupation prices. Hence, to a first-order approximation, the value of θ affects the response of average wages to shocks by affecting occupation price changes.

¹³More generally, given changes in occupation prices, the shape of the distribution of $\varepsilon(z, \kappa, \omega)$ —which we have assumed to be Fréchet—does not matter for the first-order effect of any shock on average wage changes. The reason is that, for any worker group λ , the marginal worker’s wage is equal across occupations and equipment types.

4 Decomposition

In this section we use our closed economy model to quantify the impact of computerization, changes in labor composition, changes in occupation demand, and changes in labor-group specific productivities on observed changes in relative wages in the US. Throughout this section we assume that $T_t(\lambda, \kappa, \omega)$ can be expressed as

$$T_t(\lambda, \kappa, \omega) \equiv T_t(\lambda) T_t(\kappa) T_t(\omega) T(\lambda, \kappa, \omega). \quad (10)$$

Accordingly, whereas we allow labor group, $T_t(\lambda)$, equipment type, $T_t(\kappa)$, and occupation, $T_t(\omega)$, productivity to vary over time, we impose that the interaction between labor group, equipment type, and occupation productivity, $T(\lambda, \kappa, \omega)$, is constant over time. We therefore assume that comparative advantage is fixed over time. This restriction allows us to separate cleanly the impact on relative wages of λ -, κ -, and ω -specific productivity shocks.

How restrictive is the assumption in equation (10) in practice? Note that, under this assumption, if we use data generated by our model and regress (i) the log of allocations in 2003 on the log of allocations in 1984 as well as labor-group, equipment-type, and occupation fixed effects or (ii) the log change in allocations between 1984 and 2003 on the same set of fixed effects, we would obtain $R^2 = 1$ in both cases. Using the data described in Section 4.1 instead, we run the two regressions described in (i) and (ii), weighing observations by the employment within each $(\lambda, \kappa, \omega)$ cell in 1984, and we obtain $R^2 = 0.85$ and $R^2 = 0.71$, respectively. In Appendix E.3 we show that our main conclusions are robust to a set of decompositions in which we allow two-dimensional comparative advantage to vary over time; each of these specifications is less restrictive than that in equation (10), which holds all comparative advantage fixed over time.¹⁴ We drop the restriction in equation (10) when computing the international trade counterfactuals in Section 6.

Given the restriction in equation (10), the system in equations (6)-(8) simplifies to

$$\hat{w}(\lambda) = \hat{T}(\lambda) \left[\sum_{\kappa, \omega} (\hat{q}(\omega) \hat{q}(\kappa))^\theta \pi_{t_0}(\lambda, \kappa, \omega) \right]^{1/\theta} \quad (11)$$

$$\hat{\pi}(\lambda, \kappa, \omega) = \frac{(\hat{q}(\omega) \hat{q}(\kappa))^\theta}{\sum_{\kappa', \omega'} (\hat{q}(\omega') \hat{q}(\kappa'))^\theta \pi_{t_0}(\lambda, \kappa', \omega')} \quad (12)$$

¹⁴In the most general case, if we let $T_t(\lambda, \kappa, \omega)$ vary freely over time in the closed economy, we would only be able to report the joint effect of the combination of all λ -, κ -, and ω -specific shocks on relative wages (while the impact of labor composition would be exactly the same as in our baseline specification).

$$\hat{a}(\omega)\hat{q}(\omega)^{(1-\alpha)(1-\rho)}\hat{E} = \frac{\sum_{\lambda,\kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega)}{\zeta_{t_0}(\omega)}. \quad (13)$$

In this simplified system, we have combined equipment productivity, $T_t(\kappa)$, with equipment prices, $p_t(\kappa)$, to form a composite equipment productivity measure: $q_t(\kappa) \equiv p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} T_t(\kappa)$. Similarly, we have combined occupation productivity, $T_t(\omega)$, with occupation demand, $\mu_t(\omega)$, to form a composite occupation shifter: $a_t(\omega) \equiv \mu_t(\omega) T_t(\omega)^{(1-\alpha)(\rho-1)}$. Finally, we have also combined occupation productivity, $T_t(\omega)$, with occupation prices, $p_t(\omega)$, to form transformed occupation prices: $q_t(\omega) \equiv p_t(\omega)^{1/(1-\alpha)} T_t(\omega)$. Thus, the forcing variables (shocks) in this system are: (i) $\hat{L}(\lambda)$, which we refer to as *labor composition*; (ii) $\hat{q}(\kappa)$, which we refer to as *equipment productivity*; (iii) $\hat{a}(\omega)$, which we refer to as *occupation shifters*; and (iv) $\hat{T}(\lambda)$, which we refer to as *labor productivity*. Given these shocks and a choice of a numeraire, equations (11)-(13) yield the model's implied values of changes in average wages, $\hat{w}(\lambda)$, allocations, $\hat{\pi}(\lambda, \kappa, \omega)$, and transformed occupation prices, $\hat{q}(\omega)$.¹⁵

According to equations (11)-(13), we need the following ingredients to perform our decomposition. First, we require measures of factor allocations, $\pi_{t_0}(\lambda, \kappa, \omega)$, average wages, $w_{t_0}(\lambda)$, labor composition, $L_{t_0}(\lambda)$, and labor payments by occupation, $\zeta_{t_0}(\omega)$ in the initial period. We describe the data we employ to obtain each of these measures in Section 4.1. Second, we require measures of relative shocks to labor composition, $\hat{L}(\lambda)/\hat{L}(\lambda_1)$, occupation shifters, $\hat{a}(\omega)/\hat{a}(\omega_1)$, equipment productivity to the power θ , $\hat{q}(\kappa)^\theta/\hat{q}(\kappa_1)^\theta$, and labor productivity, $\hat{T}(\lambda)/\hat{T}(\lambda_1)$. We describe our approach to measure these shocks in Section 4.2. Finally, we require estimates of α , ρ , and θ . We describe our estimation strategy in Section 4.3.

4.1 Data

We use data from the CPS's Merged Outgoing Rotation Group (MORG CPS) and the October CPS Supplement (October Supplement) for the years 1984, 1989, 1993, 1997, and 2003. We restrict our sample by dropping workers who are younger than 17 or older than 64 years old, do not report positive paid hours worked, or are self employed. After cleaning, the MORG CPS and October Supplement contain data for roughly 115,000 and

¹⁵The two components of equipment productivity, $T_t(\kappa)$ and $p_t(\kappa)$, could be measured separately using information on equipment prices, which are subject to known quality-adjustment issues raised by, e.g., Gordon (1990). Similarly, the two components of the occupation shifter, $T_t(\omega)$ and $\mu_t(\omega)$, could be measured separately using information on occupation prices, which are hard to measure in practice. If the aim of our exercise were not to determine the impact of observed changes in the economic environment on observed wages but rather to predict the impact of hypothetical changes in the economic environment, then it would not be necessary to combine changes in equipment productivity and equipment production costs into a single shock or to combine changes in occupation demand and occupation productivity into a single composite shock.

50,000 individuals, respectively, in each year. We provide additional details in Appendix B.

We divide workers into thirty labor groups by gender, education (high school dropouts, high school graduates, some college completed, college completed, and graduate training), and age (17-30, 31-43, and 44-64). We consider two types of equipment: computers and other equipment. We use thirty occupations, which we list, together with summary statistics, in Table 10 in Appendix B.

We use the MORG CPS to construct total hours worked and average hourly wages by labor group by year.¹⁶ We use the October Supplement to construct a measure of the share of total hours worked by labor group λ that is spent using equipment type κ in occupation ω in year t , $\pi_t(\lambda, \kappa, \omega)$. In 1984, 1989, 1993, 1997, and 2003, the October Supplement asked respondents whether they “have direct or hands on use of computers at work,” “directly use a computer at work,” or “use a computer at/for his/her/your main job.” This question defines a computer as a machine with typewriter like keyboards, whether a personal computer, laptop, mini computer, or mainframe. Identifying computers with the index κ_c and the non-computer equipment with the index κ_n , we construct $\pi_t(\lambda, \kappa_c, \omega)$ as the hours worked in occupation ω by λ workers who report that they use a computer at work relative to the total hours worked by labor group λ in year t . We construct $\pi_t(\lambda, \kappa_n, \omega)$ as the hours worked in occupation ω by λ workers who report that they do not use a computer at work relative to the total hours worked by labor group λ in year t .¹⁷

When interpreting our measures of factor allocations, $\pi_t(\lambda, \kappa, \omega)$, one should bear in mind four limitations. First, our measure of computerization is narrow, as it does not capture, e.g., automation of assembly lines. Second, at the individual level our computer-use variable takes only two values, zero or one, without detailing the share of each worker’s time at work spent using computers. Third, we are not using any information on the allocation of non-computer equipment. Finally, the computer use question was discontinued after 2003.¹⁸

¹⁶We measure wages using the MORG CPS rather than the March CPS. Both datasets imply similar changes in average wages for each of our thirty labor groups. However, the March CPS does not directly measure hourly wages of workers paid by the hour and, therefore, introduces substantial measurement error in individual wages (see e.g. Lemieux 2006). For most of our analysis, we do not use individual wages; however, we do so in sensitivity analyses in Appendix D.3.

¹⁷On average across the five years considered in the analysis, we measure $\pi_t(\lambda, \kappa, \omega) = 0$ for 26.3% of the $(\lambda, \kappa, \omega)$ triplets. As a robustness check, in Appendix F.3 we drop age as a characteristic defining a labor group and redo our analysis with the resulting ten labor groups. With only ten labor groups, we measure $\pi_t(\lambda, \kappa, \omega) = 0$ for 11.7% of the $(\lambda, \kappa, \omega)$ triplets. Our main decomposition results are largely robust to decreasing the number of labor groups from 30 to 10.

¹⁸The period 1984-2003, however, accounts for a substantial share of the increase in the skill premium and reduction in the gender gap observed in the US since the late 1960s. In unreported results, we show that the German *Qualification and Working Conditions* survey, which alleviates some of the limitations of the

Factor allocation. Aggregating $\pi_t(\lambda, \kappa, \omega)$ across ω and λ , Panel A of Table 1 shows that women and more educated workers use computers more intensively in the aggregate than men and less educated workers, respectively. Panel B shows that this aggregate difference between women and men is largely driven by differences in female and male allocations across occupations whereas this aggregate difference between college- and non-college-educated workers is driven in equal shares by differences in allocations between occupations and differences in computer usage within occupations. The disaggregated $\pi_t(\lambda, \kappa, \omega)$ data allows us also to identify sorting patterns across equipment types of worker groups differing along a particular dimension of heterogeneity (e.g. education) while holding fixed the other dimensions of heterogeneity (e.g. age, gender, and year) and the occupation. Specifically, to determine the extent to which college educated workers (λ') compared to workers with high school degrees in the same gender-age group (λ) use computers (κ_c) relatively more than non-computer equipment (κ_n) within occupations (ω), the left panel of Figure 2 plots the histogram of

$$\log \frac{\pi_t(\lambda', \kappa_c, \omega)}{\pi_t(\lambda', \kappa_n, \omega)} - \log \frac{\pi_t(\lambda, \kappa_c, \omega)}{\pi_t(\lambda, \kappa_n, \omega)}$$

across all five years, thirty occupations, and six gender-age groups described above. Figure 2 shows that college-educated workers are relatively more likely to use computers within occupations compared to high-school-educated workers; hence, according to our model, college-educated workers have a comparative advantage using computers within occupations relative to high-school-educated workers. A similar conclusion holds comparing across other education groups: more educated workers systematically have a comparative advantage using computers.

The right panel of Figure 2 plots a similar histogram, where λ' and λ now denote female and male workers in the same education-age group. This figure shows that there is no clear difference on average in computer usage across genders within occupations (i.e. the histogram is roughly centered around zero). Hence, our model rationalizes the fact that women use computers more than men at the aggregate level—see Panel A of Table 1—by concluding that women have a comparative advantage in occupations in which computers have a comparative advantage. For instance, we observe that women are much more likely than men to work in administrative support relative to construction occupations, conditional on the type of equipment used; and that each labor group is much more likely to use computers in administrative support than in construction oc-

CPS data, reveals similar patterns of comparative advantage in Germany as what the CPS data reveals for the US.




Figure 2: The distribution of computer relative to non-computer usage for college degree relative to high school degree workers (female relative to male workers) across all gender-age (education-age) cells in the left (right) panel

cupations.

Three-way comparative advantage—between labor groups, equipment types, and occupations, as summarized in the function $T(\lambda, \kappa, \omega)$ —is essential for our model to account for the detailed factor allocations that we summarize here and in Table 1.

4.2 Measuring Shocks

Here we describe our baseline procedure to measure the four shocks into which we decompose the 1984-2003 observed changes in relative average wages: labor composition, $\hat{L}(\lambda)/\hat{L}(\lambda_1)$, equipment productivity (to the power of θ), $\hat{q}(\kappa)^\theta/\hat{q}(\kappa_1)^\theta$, occupation shifters, $\hat{a}(\omega)/\hat{a}(\omega_1)$, and labor productivity, $\hat{T}(\lambda)/\hat{T}(\lambda_1)$. We measure changes in labor composition directly from the MORG CPS. We measure changes in equipment productivity using data only on changes in disaggregated factor allocations over time, $\hat{\pi}(\lambda, \kappa, \omega)$. We measure changes in occupation shifters using data on changes in disaggregated factor allocations and labor income shares across occupations as well as model parameters. Finally, we measure changes in labor productivity as a residual to match observed changes in relative wages. We provide details on the baseline procedure described here in Appendix C.1, and alternative procedures in Appendices C.2 and C.3. All these procedures yield very similar results.

First, consider the measurement of changes in equipment productivity to the power of θ . Equation (12) implies

$$\frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta} = \frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa_1, \omega)} \quad (14)$$

for any (λ, ω) pair. Hence, if computer productivity rises relative to that of other equipment between t_0 and t_1 , then the share of labor-group λ hours spent working with computers relative to other equipment in occupation ω will increase. This is true for any

(λ, ω) pair. It is important to condition on (λ, ω) pairs when identifying changes in equipment productivity because unconditional growth over time in computer usage, shown in Panel A of Table 1, may also reflect growth in the supply of labor groups who have a comparative advantage using computers and/or changes in occupation shifters that are biased towards occupations in which computers have a comparative advantage. Given the possibility of measurement error in $\hat{\pi}(\cdot, \cdot, \cdot)$, we construct an average of these (λ, ω) pair-specific measures to obtain a unique measure of changes in equipment productivity to the power of θ , $\hat{q}(\kappa)^\theta / \hat{q}(\kappa_1)^\theta$, as described in Appendix C.1.

Second, consider the measurement of changes in occupation shifters. Equation (13) implies

$$\frac{\hat{a}(\omega)}{\hat{a}(\omega_1)} = \frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega_1)} \left(\frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} \right)^{(1-\alpha)(\rho-1)}. \quad (15)$$

We construct the two components in the right-hand side of equation (15) as follows. First, equation (12) implies

$$\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta} = \frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa, \omega_1)} \quad (16)$$

for any (λ, κ) pair. For each ω , we construct an average of these (λ, κ) pair-specific measures to obtain a unique measure of changes in transformed occupation prices to the power of θ , $\hat{q}(\omega)^\theta / \hat{q}(\omega_1)^\theta$, as described in Appendix C.1.¹⁹ Given values of α , ρ , and θ (which we estimate as described below), we calculate $(\hat{q}(\omega) / \hat{q}(\omega_1))^{(1-\alpha)(\rho-1)}$. Second, given values of $\hat{q}(\kappa)^\theta / \hat{q}(\kappa_1)^\theta$ and $\hat{q}(\omega)^\theta / \hat{q}(\omega_1)^\theta$, we construct $\hat{\zeta}(\omega) / \hat{\zeta}(\omega_1)$ using the right-hand side of equation (13).

Finally, we measure changes in labor productivity as a residual to match changes in relative wages. Specifically, we rewrite equation (11) as

$$\frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} = \frac{\hat{T}(\lambda)}{\hat{T}(\lambda_1)} \left(\frac{\hat{s}(\lambda)}{\hat{s}(\lambda_1)} \right)^{1/\theta}, \quad (17)$$

where $\hat{s}(\lambda)$ is a labor-group-specific weighted average of changes in equipment productivity and transformed occupation prices, both raised to the power of θ ,

$$\hat{s}(\lambda) = \sum_{\kappa, \omega} \frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta} \frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta} \pi_{t_0}(\lambda, \kappa, \omega), \quad (18)$$

¹⁹We use data on observed allocations to measure these changes in transformed occupation prices. However, in calculating changes in relative wages in response to any subset of shocks in our decomposition, we solve for counterfactual changes in transformed occupation prices and allocations using equations (12) and (13).

which we can construct using observed allocations, changes in observed allocations, and equations (14) and (16). Given wage data and a value of θ , we recover $\hat{T}(\lambda) / \hat{T}(\lambda_1)$. Note that only our measures of changes in labor productivities are directly a function of the observed changes in relative wages that we aim to decompose: given parameters α , ρ , and θ , observed changes in wages do not affect our measures of changes in labor composition or equipment productivity to the power of θ , and they only affect our measures of occupation shifters indirectly through their impact on $\hat{\zeta}(\omega) / \hat{\zeta}(\omega_1)$.

4.3 Parameter Estimates

The Cobb-Douglas parameter α , which matters for our results only when ρ is different from one, determines payments to all equipment (computers and non-computer equipment) relative to the sum of payments to equipment and labor. We set $\alpha = 0.24$, consistent with estimates in [Burstein et al. \(2013\)](#). [Burstein et al. \(2013\)](#) disaggregate total capital payments (i.e the product of the capital stock and the rental rate) into structures and equipment using US data on the value of capital stocks and—since rental rates are not directly observable—setting the average rate of return over the period 1963-2000 of holding each type of capital (its rental rate plus price appreciation less the depreciation rate) equal to the real interest rate.

As a reminder, the parameter θ determines the within-worker dispersion of productivity across occupation-equipment type pairs and the parameter ρ is the elasticity of substitution across occupations in the production of the final good. We estimate these two parameters jointly using a method of moments approach. In order to derive the relevant moment conditions, we rewrite equation (17) as

$$\log \hat{w}(\lambda, t) = \zeta_\theta(t) + \frac{1}{\theta} \log \hat{s}(\lambda, t) + \iota_\theta(\lambda, t), \quad (19)$$

where $\hat{s}(\lambda, t)$ is defined in equation (18), $\zeta_\theta(t)$ is a time effect, and $\iota_\theta(\lambda, t) \equiv \log \hat{T}(\lambda, t)$ captures unobserved changes in labor productivity.²⁰ Similarly, equation (13) can be expressed as

$$\log \hat{\zeta}(\omega, t) = \zeta_\rho(t) + \frac{1}{\theta} (1 - \alpha) (1 - \rho) \log \frac{\hat{q}(\omega, t)^\theta}{\hat{q}(\omega_1, t)^\theta} + \iota_\rho(\omega, t). \quad (20)$$

where $\hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$ is defined in equation (16), $\zeta_\rho(t)$ is a time effect, and $\iota_\rho(\omega, t) \equiv$

²⁰Since our estimation uses the panel dimension of the data, here we must consider multiple changes (over different time periods) for each variable of interest. Hence, for the purposes of this section and in related appendices, for any variable x we use $\hat{x}(t)$ to denote the relative change in a variable x between any two consecutive periods t and $t' > t$.

$\log \hat{a}(\omega, t)$ captures unobserved changes in occupation shifters.

Equations (19) and (20) may be used to identify θ and ρ jointly. According to our model, however, the observed covariate in equation (19), $\log \hat{s}(\lambda, t)$, is predicted to be correlated with its error term, $\iota_\theta(\lambda, t)$, and the observed covariate in equation (20) is predicted to be correlated with its error term, $\iota_\rho(\omega, t)$. To address the endogeneity of these covariates, we construct instruments using different averages of observed changes in equipment productivity to the power of θ , $\hat{q}(\kappa, t)^\theta / \hat{q}(\kappa_1, t)^\theta$. Specifically, we instrument the covariate $\log \hat{s}(\lambda, t)$ using a labor-group-specific average,

$$\log \sum_{\kappa} \frac{\hat{q}(\kappa, t)^\theta}{\hat{q}(\kappa_1, t)^\theta} \sum_{\omega} \pi_{1984}(\lambda, \kappa, \omega),$$

and the covariate $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$ using an occupation-specific average

$$\log \sum_{\kappa} \frac{\hat{q}(\kappa, t)^\theta}{\hat{q}(\kappa_1, t)^\theta} \sum_{\lambda} \frac{L_{1984}(\lambda) \pi_{1984}(\lambda, \kappa, \omega)}{\sum_{\lambda', \kappa'} L_{1984}(\lambda') \pi_{1984}(\lambda', \kappa', \omega)}.$$

We use these two instruments and equations (19) and (20) to build moment conditions that identify the parameter vector (θ, ρ) .²¹

Our estimates exploit data on four time periods: 1984-1989, 1989-1993, 1993-1997, and 1997-2003. We report the point estimates and standard errors in the top row of Table 2. The resulting estimate of θ is 1.81 with a standard error of 0.28; the estimate of ρ is also 1.81 but with a slightly larger standard error of 0.36.²²

In reviewing [Krusell et al. \(2000\)](#), [Acemoglu \(2002\)](#) raises a concern that, if valid, will affect the estimates of θ and ρ computed using the expressions in equations (19) and (20) as estimating equations. Specifically, he points out that the presence of common trends in unobserved labor-group-specific productivity, explanatory variables, and instruments may bias the estimates of wage elasticities. In order to address this concern, we follow [Katz and Murphy \(1992\)](#) and [Acemoglu \(2002\)](#) and additionally control for a λ -specific time trend in equation (19). In addition, we also control for an occupation-specific time trend in equation (20). The estimates of θ and ρ that we obtain in this case are, respectively 1.26, with a standard error of 0.31, and 2.1, with a standard error of 0.75, as displayed in row 2 of Table 2.²³

²¹Additional details on these moment conditions are provided in Appendix D. This appendix also contains a detailed discussion of how, under the assumptions of our model, these two instruments are valid and expected to be correlated with the corresponding endogenous variables.

²²Estimating θ and ρ without instruments, we find that the direction of the bias in each parameter is consistent with the predictions of our theory (see Appendix D).

²³In Appendix D, we re-estimate θ separately for each labor group using moments of the unconditional distribution of observed wages within each labor group λ . The median value of $\theta(\lambda)$ is 2.62.

Parameter	Time Trend?	Estimate	(SE)
(θ, ρ)	NO	(1.81, 1.81)	(0.28, 0.36)
(θ, ρ)	YES	(1.26, 2.10)	(0.31, 0.75)

Table 2: Parameter estimates based on joint estimation

When computing the main results of our analysis, we verify in Appendix E.1 the robustness of our conclusions to the different estimates of θ and ρ reported in this section and in Appendix D.

4.4 Results

In this section we summarize the results of our decomposition of observed changes in relative wages in the US between 1984 and 2003. We construct various measures of changes in between-group inequality, each of them aggregating wage changes across our thirty labor groups in a different way. When doing so, both in the model and in the data, we construct *composition-adjusted* wage changes; that is, for each aggregated measure we report, we average disaggregated measures of wage changes across the corresponding labor groups using constant weights over time, as described in detail in Appendix B. For each measure of inequality, we report its cumulative log change between 1984 and 2003, calculated as the sum of the log change over all sub-periods in our data.²⁴ We also report log changes over certain sub-period in our data.

Skill premium. The first column in Table 3 reports the change observed in the data, which is also the change predicted by our model when all shocks—in labor composition, occupation shifters, equipment productivity, and labor productivity—are simultaneously considered. The skill premium increased by 15.1 (log) percent between 1984 and 2003. The subsequent four columns summarize the counterfactual percent change in the skill premium predicted by the model if only one of the four shocks is considered (i.e. holding the other exogenous parameters at their t_0 level).

Changes in labor composition decrease the skill premium in each sub-period. Specifically, the increase in hours worked by those with college degrees relative to those without of 47.4 percent between 1984 and 2003 decreases the skill premium by 11.3 percent. Changes in relative demand across labor groups must, therefore, compensate for the impact of changes in labor composition in order to generate the observed rise of the skill premium in the data.

²⁴We obtain very similar results if we directly compute changes in wages between 1984 and 2003 instead of adding changes in log relative wages over all sub-periods.

	Data	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
1984 - 1989	0.057	-0.031	0.024	0.052	0.010
1989 - 1993	0.064	-0.017	-0.004	0.045	0.040
1993 - 1997	0.037	-0.022	0.041	0.021	-0.002
1997 - 2003	-0.007	-0.042	-0.012	0.041	0.006
1984 - 2003	0.151	-0.113	0.049	0.159	0.055

Table 3: Decomposing changes in the log skill premium (the wage of workers with a college degree relative to those without)

Changes in equipment productivity, i.e. computerization, account for roughly 60 percent of the sum of the demand-side forces pushing the skill premium upwards: $0.60 \simeq 0.159 / (0.049 + 0.159 + 0.055)$. Furthermore, changes in equipment productivity are particularly important in generating increases in the skill premium in those sub-periods in which the skill premium rose most dramatically: 1984-1989 and 1989-1993. These are also precisely the years in which the overall share of workers using computers rose most rapidly; see Table 1. The intuition for why our model predicts that computerization had a large impact on the skill premium is the following. The procedure described in Section 4.2 to measure changes in computer productivity implies large growth in this variable.²⁵ This growth raises the skill premium for two reasons, as described in detail in Section 3.4. First, educated workers have a direct comparative advantage using computers within occupations, as shown in Panel B of Table 1 and, in more detail, in Figure 2. Second, educated workers have a comparative advantage in occupations in which computers have a comparative advantage and, given our estimate of ρ , this implies that computerization raises the wages of labor groups disproportionately employed in computer-intensive occupations.

Changes in occupation shifters account for roughly 19 percent of the sum of the forces pushing the skill premium upwards. Skill-intensive occupations grew disproportionately in our sample period, as documented in Figure 1 and in Table 10 in Appendix B. If $\rho = 1$, then our model would attribute this growth fully to occupation shifters and, thus, occupa-

²⁵This is consistent with ample direct evidence showing a rapid decline in the price of computers relative to all other equipment types and structures, which we do not directly use in our estimation. The decline over time in the U.S. price of equipment relative to structures—see e.g. Greenwood et al. (1997)—is mostly driven by a decline in computer prices. For example, between 1984 and 2003: (i) the prices of industrial equipment and of transportation equipment relative to the price of computers and peripheral equipment have risen by factors of 32 and 34, respectively (calculated using the BEA’s Price Indexes for Private Fixed Investment in Equipment and Software by Type), and (ii) the quantity of computers and peripheral equipment relative to the quantities of industrial equipment and of transportation equipment rose by a factor of 35 and 33, respectively (calculated using the BEA’s Chain-Type Quantity Indexes for Net Stock of Private Fixed Assets, Equipment and Software, and Structures by Type).

tions shifters would have accounted for a larger share of the growth of the skill premium (see Appendix E.1). However, because our estimate of ρ is different from 1, other shocks also play important roles in shaping income shares across occupations. In Appendix F.2, we document the importance of each shock for the growth of occupations with different characteristics (using O*NET constructed task measures), and show that changes in equipment productivity, labor composition, and labor productivity play a significant role in the systematic contraction of occupations that are intensive in routine manual as well as non-routine manual physical tasks and the systematic expansion of occupations that are intensive in non-routine cognitive analytical and non-routine cognitive interpersonal tasks.

Finally, labor productivity, which we estimate as a residual that matches observed changes in relative wages unexplained by the remaining forces included in our model, accounts for roughly 21 percent of the sum of the effects of all three demand-side mechanisms.

In Appendix E.2 we demonstrate the importance of accounting for all three forms of comparative advantage by performing similar exercises in versions of our model that omit some of them. For example, if we were to abstract from any comparative advantage at the level of occupations, we would incorrectly conclude that all of the rise in the skill premium has been driven by changes in relative equipment productivities. Similarly, if we were to abstract from any comparative advantage at the level of equipment, we would incorrectly magnify the importance of occupation shifters and labor productivity in explaining the rise of the skill premium.

Five education groups. Table 4 decomposes changes in between-education-group wage inequality at a higher level of disaggregation than what the skill premium captures. The

	Data	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
HS dropouts / Average	-0.083	0.085	-0.041	-0.181	0.052
HS grad / Average	-0.046	0.037	-0.019	-0.055	-0.007
Some college / Average	-0.009	-0.009	0.010	0.047	-0.058
College / Average	0.091	-0.074	0.024	0.112	0.027
Grad training / Average	0.149	-0.101	0.058	0.127	0.064

Table 4: Decomposing changes in log relative wages across education groups between 1984 and 2003

results reported in Table 3 are however qualitatively robust: computerization is the central force driving changes in between-education-group inequality, whereas labor productivity plays a relatively minor role.

Skill premium by age. The rise of the skill premium over our time period is substantially larger amongst middle aged and young workers than among older workers, as shown in column 1 of Table 5. This fact was first identified by [Katz and Murphy \(1992\)](#). In a framework featuring changes in labor composition and labor productivity, [Card and Lemieux \(2001\)](#) show that differential changes in labor composition can account for this pattern. While we find that changes in labor composition do push the skill premium of older workers below those of younger and middle-aged workers (column 2), changes in equipment productivity explain the vast majority of the lower growth of the skill premium amongst older workers (column 4). This result follows from the fact that, in our data, aggregate computer use of college-educated relative to non-college-educated workers is higher amongst young and middle-aged workers than amongst old workers.

Age group	Data	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Young	0.154	-0.110	0.040	0.204	0.020
Middle	0.179	-0.108	0.053	0.154	0.078
Old	0.107	-0.117	0.053	0.118	0.051

Table 5: Decomposing the change in the skill premium by age group over the period 1984-2003

Gender gap. The average wage of men relative to women, the gender gap, declined by 13.3 percent between 1984 and 2003. Table 6 decomposes changes in the gender gap over the full sample and over each sub-period. The increase in hours worked by women relative to men of roughly 12.6 percent between 1984 and 2003 increased the gender gap by 4.1 percent. Changes in relative demand across labor groups must, therefore, compensate for the impact of changes in labor composition in order to generate the observed fall in the gender gap.

	Data	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
1984 - 1989	-0.056	0.011	-0.012	-0.015	-0.041
1989 - 1993	-0.052	0.013	-0.036	-0.014	-0.015
1993 - 1997	-0.003	0.006	0.015	-0.005	-0.020
1997 - 2003	-0.021	0.011	-0.036	-0.012	0.017
1984 - 2003	-0.133	0.041	-0.069	-0.046	-0.058

Table 6: Decomposing changes in the log gender gap (the wage of men relative to women)

In spite of the fact that comparative advantage using computers does not vary by gender, changes in equipment productivity account for roughly 27 percent of the sum of the

impact of the demand-side forces. This is due to women having a comparative advantage in computer-intensive occupations which, together with an estimate of the elasticity of substitution across occupations larger than one, implies that computerization raises the wages of labor groups disproportionately employed in computer-intensive occupations.

Changes in occupation shifters account for roughly 40 percent of the sum of the forces decreasing the gender gap over the full sample. This is reflected in part by the fact that a number of male-intensive occupations—including, for example, mechanics/repairers as well as machine operators/assemblers/inspectors—contracted substantially between 1984 and 2003; see Figure 1 in Section 1 and Table 10 in the Appendix B for details. Again, however, with $\rho \neq 1$, all shocks contribute to these changes in occupation size.

Changes in labor productivity account for a sizable share, roughly 34 percent, of the impact of the demand-side forces on the gender gap, and play a central role in each sub-period except for 1997-2003. This suggests that factors such as changes in gender discrimination—if they affect labor productivity irrespective of the type of equipment used and the occupation of employment—may have played a substantial role in reducing the gender gap, especially early in our sample (in the 1980s and early 1990s); see e.g. Hsieh et al. (2016).

Thirty disaggregated labor groups. One of the advantages of our framework is that we can study the determinants of wage changes across a large number of labor groups. Table 7 presents evidence on the relative importance of the three demand-side shocks—occupation shifters, equipment productivity, and labor productivity—in explaining relative changes in wages across the thirty labor groups that we consider in our analysis. More precisely, we show in this table the results from decomposing the variance of the changes in relative wages predicted by our model, when all three demand-side shocks are active, into the covariances of these changes with those that are predicted by our model when only one of the three demand-side shocks is activated. Specifically, we present the three OLS estimates of the projection of each of the 30 changes in relative wages predicted by each of the three demand-side shocks included in our analysis on the 30 changes in relative wages predicted by our model when all the three demand-side shocks are taken into account.

The results show that, in the period 1984 to 2003, equipment productivity explains over 50 percent of the variance in the change in relative wages (implied by the combination of the three demand side shocks) across the thirty labor groups considered in the analysis. In this same time period, occupation shifters and labor productivity each explain slightly less than 25 percent. Whereas equipment productivity is also the most important demand-side contributor in the periods 1984-1989 and 1997-2003, labor productivity is the

	1984-2003	1984-1989	1989-1993	1993-1997	1997-2003
Equipment Productivity	53%	48%	33%	27%	50%
Occupation Shifter	24%	23%	15%	44%	20%
Labor Productivity	24%	29%	52%	28%	30%

Table 7: Variance Decomposition across 30 labor groups

Let $x_i(\lambda)$ denote the change in the average wage of group λ induced by demand-side shock i and let $y(\lambda) = \sum_{i=1}^3 x_i(\lambda)$. For each i and time period we report $cov(x_i(\lambda), y(\lambda)) / var(y(\lambda))$.

most important force in 1989-1993 and occupation shifters are the most important force in 1993-1997.

5 Open Economy Model

In this section we extend the model introduced in Section 3 to allow for international trade in equipment goods and occupation services.

5.1 Environment

We assume the final good is non-traded and produced according to the open economy counterpart of equation (1), with output in country n given by

$$Y_n = \left(\sum_{\omega} \mu_n(\omega)^{1/\rho} D_n(\omega)^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)},$$

where $D_n(\omega)$ denotes absorption of occupation ω in country n , and where we omit time subscripts to simplify notation. Absorption of each occupation is itself a CES aggregate of the services of these occupations sourced from all countries in the world,

$$D_n(\omega) = \left(\sum_i D_{in}(\omega)^{\frac{\eta(\omega)-1}{\eta(\omega)}} \right)^{\eta(\omega)/(\eta(\omega)-1)},$$

where $D_{in}(\omega)$ is absorption in country n of occupation ω sourced from country i , and $\eta(\omega) > 1$ is the elasticity of substitution across source countries for occupation ω .²⁶ Similarly, absorption of equipment of type κ in country n is a CES aggregate of equipment

²⁶The assumption that the final good is non-traded is without loss of generality for our results on relative wages. We assume an Armington trade model only for expositional simplicity. Our results would also hold in a Ricardian model as in Eaton and Kortum (2002).

sourced from all countries in the world,

$$D_n(\kappa) = \left(\sum_i D_{in}(\kappa)^{\frac{\eta(\kappa)-1}{\eta(\kappa)}} \right)^{\eta(\kappa)/(\eta(\kappa)-1)}. \quad (21)$$

Trade is subject to iceberg transportation costs, where $d_{ni}(\omega) \geq 1$ and $d_{ni}(\kappa) \geq 1$ denote the units of occupation ω output and equipment κ output, respectively, that must be shipped from origin country n in order for one unit to arrive in destination country i . Output of occupation ω and equipment κ in country n satisfy, respectively,

$$Y_n(\omega) = \sum_i d_{ni}(\omega) D_{ni}(\omega). \quad (22)$$

$$Y_n(\kappa) = \sum_i d_{ni}(\kappa) D_{ni}(\kappa). \quad (23)$$

Since the final good is non-traded, the resource constraint for the final good must satisfy equation (2). As country n 's trade costs limit to infinity, $d_{ni}(\omega) \rightarrow \infty$ and $d_{ni}(\kappa) \rightarrow \infty$ for all $n \neq i$, the economy limits to the autarkic version of the model presented in Section 3. For the exercises we consider below, we do not need to specify conditions on trade balance in each country.

5.2 Equilibrium

We now describe the differences between the equilibrium conditions of this open economy model and those of the closed economy model introduced in Section 3. In an open economy, we must distinguish between production prices and absorption prices. We denote by $p_{nn}(\omega)$ the output price of occupation ω in country n and by $p_n(\omega)$ its absorption price. Absorption in country n of occupation ω sourced from country i is given by

$$D_{in}(\omega) = \left(\frac{p_{in}(\omega)}{p_n(\omega)} \right)^{-\eta(\omega)} D_n(\omega) \quad (24)$$

and the occupation ω absorption price is given by

$$p_n(\omega) = \left[\sum_i p_{in}(\omega)^{1-\eta(\omega)} \right]^{\frac{1}{1-\eta(\omega)}}. \quad (25)$$

Here, $p_{in}(\omega)$ is the price of country i 's output of occupation ω in country n (inclusive of trade costs), which is related to the domestic output price by $p_{in}(\omega) = d_{in}(\omega) p_{ii}(\omega)$.

Production prices of equipment goods are given by $p_{nn}(\kappa) = \tilde{p}_n(\kappa) P_n$, as in the closed economy. Absorption in country n of equipment κ sourced from country i is given by

$$D_{in}(\kappa) = \left(\frac{p_{in}(\kappa)}{p_n(\kappa)} \right)^{-\eta(\kappa)} D_n(\kappa) \quad (26)$$

and the equipment κ absorption price is given by

$$p_n(\kappa) = \left[\sum_i p_{in}(\kappa)^{1-\eta(\kappa)} \right]^{\frac{1}{1-\eta(\kappa)}}. \quad (27)$$

As this expression shows, equipment absorption prices in country n depend on equipment output prices in all foreign countries and on trade costs between country n and all of its trading partners.

The relevant prices shaping allocations and average wages— $\pi_n(\lambda, \kappa, \omega)$ and $w_n(\lambda)$ in equations (3) and (4)—are absorption prices for equipment, $p_n(\kappa)$, since equipment is an input in production, and domestic production prices for occupations, $p_{nn}(\omega)$, since occupations are produced in each country. Occupation prices must therefore satisfy the open-economy versions of the occupation-market clearing condition in equation (5),

$$\sum_i \mu_i(\omega) p_{ni}(\omega)^{1-\eta(\omega)} p_i(\omega)^{\eta(\omega)-\rho} P_i^{\rho-1} E_i = \frac{1}{1-\alpha} \zeta_n(\omega), \quad (28)$$

where $\zeta_n(\omega)$ is labor income in occupation ω , E_n denotes total expenditure, and P_n is the final good price in country n (which was normalized to one in the closed economy). In general, solving for the level of occupation prices requires solving for prices in the full world equilibrium; these prices are functions of worldwide technologies, endowments, and trade costs.

6 International Trade Counterfactuals

The aim of this section is twofold. First, we show theoretically how the degree of openness may cause what we have treated in our closed economy model as primitive shocks to the cost of producing equipment and to the demand for occupations. Second, we present some simple calculations quantifying the impact of equipment and occupation trade on between-group inequality.

In general, calculating counterfactual changes in wages in a given country n in response to changes in trade costs or technologies either in n or in any other country in the

world requires detailed data on allocations and trade for all countries, which is unavailable in practice. Instead, with the aim of quantifying the impact of international trade in occupations and equipment on country n , we consider a specific counterfactual exercise that requires neither solving the full world general equilibrium nor estimating parameters for any country other than n . This counterfactual answers the following question: what are the differential effects on relative wages in country n of changes in primitives (i.e. technologies, labor compositions, and trade costs) in all countries in the world between any two periods t_0 and t_1 , relative to the effects of the same changes in primitives if country n were a closed economy?

Formally, define $w_n(\lambda; \Phi_t, \Phi_t^*, d_t)$ to be the average wage of labor group λ in country n if domestic technologies and endowments are given by Φ_t , foreign technologies and endowments are given by Φ_t^* , and the full matrix of world trade costs is d_t . Define $d_{n,t}^A$ to be an alternative matrix of world trade costs in which country n 's international trade costs (both for equipment and occupation trade) are infinite ($d_{in,t}(\kappa) = d_{in,t}(\omega) = \infty$ for all $i \neq n$). Our counterfactual calculates

$$\frac{w_n(\lambda; \Phi_{t_1}, \Phi_{t_1}^*, d_{t_1})}{w_n(\lambda; \Phi_{t_0}, \Phi_{t_0}^*, d_{t_0})} \bigg/ \frac{w_n(\lambda; \Phi_{t_1}, \Phi_{t_1}^*, d_{t_1}^A)}{w_n(\lambda; \Phi_{t_0}, \Phi_{t_0}^*, d_{t_0}^A)}. \quad (29)$$

Defining the impact on the wage of group λ of moving country n to autarky at time t as $\hat{w}_{n,t}^A(\lambda) \equiv w_n(\lambda; \Phi_t, \Phi_t^*, d_{n,t}^A) / w_n(\lambda; \Phi_t, \Phi_t^*, d_{n,t})$, notice that our counterfactual can be expressed succinctly as $\hat{w}_{n,t_0}^A(\lambda) / \hat{w}_{n,t_1}^A(\lambda)$. As computing the expression in equation (29) amounts to computing the impact of moving to autarky at two different points in time, we describe here how we calculate the counterfactual change in wages in country n from moving to autarky at a particular point in time t . Moving to autarky affects relative wages through two channels: changing (i) relative absorption prices across equipment types and (ii) relative demand across occupations. We now show that these two channels affect relative wages exactly like changes in (i) equipment productivity and (ii) occupation shifters in the closed economy.

Choosing the final good price in country n as the numeraire, the change in the absorption price of equipment, defined in equation (27), when moving to autarky at time t is

$$\hat{p}_n(\kappa) = s_{nn}(\kappa)^{\frac{1}{1-\eta(\kappa)}}, \quad (30)$$

where $s_{nn}(\kappa)$ denotes the fraction of expenditures on equipment κ in country n purchased

from itself (i.e. one minus the import share) at time t ,

$$s_{nn}(\kappa) = \frac{p_{nn}(\kappa) D_{nn}(\kappa)}{p_n(\kappa) D_n(\kappa)}.$$

In deriving equation (30) we have used equation (26), which implies

$$s_{nn}(\kappa) = (p_{nn}(\kappa) / p_n(\kappa))^{1-\eta(\kappa)}.$$

Changes in occupation output prices induced by moving to autarky at time t must satisfy the following occupation clearing condition,

$$f_{nn}(\omega) \left(\frac{\hat{p}_{nn}(\omega)}{\hat{p}_n(\omega)} \right)^{\rho-\eta(\omega)} \hat{p}_{nn}(\omega)^{1-\rho} \hat{E}_n = \frac{1}{1-\alpha} \hat{\zeta}_n(\omega), \quad (31)$$

where $f_{nn}(\omega)$ denotes the fraction of total sales of occupation ω in country n purchased from itself (i.e. one minus the export share) at time t ,

$$f_{nn}(\omega) = \frac{p_{nn}(\omega) D_{nn}(\omega)}{p_{nn}(\omega) Y_n(\omega)}.$$

Using equations (24) and (25) we can express equation (31) more simply as

$$f_{nn}(\omega) s_{nn}(\omega)^{\frac{\rho-\eta(\omega)}{\eta(\omega)-1}} \hat{p}_{nn}(\omega)^{1-\rho} \hat{E}_n = \frac{1}{1-\alpha} \hat{\zeta}_n(\omega), \quad (32)$$

where $s_{nn}(\omega)$ is defined analogously to $s_{nn}(\kappa)$ above.

Equations (30) and (32) allow us to calculate changes in relative wages in country n if n were to move to autarky at time t using the closed economy system of equations in which we imposed no restrictions on $T_t(\lambda, \kappa, \omega)$ —(6), (7) and (8)—and the following set of shocks

$$\begin{aligned} \hat{T}_n(\lambda, \kappa, \omega) &= \hat{L}_n(\lambda) = 1, \\ \hat{p}_n(\kappa) &= s_{nn}(\kappa)^{\frac{1}{1-\eta(\kappa)}}, \end{aligned} \quad (33)$$

and

$$\hat{p}_n(\omega) = f_{nn}(\omega) s_{nn}(\omega)^{\frac{\rho-\eta(\omega)}{\eta(\omega)-1}}, \quad (34)$$

where $s_{nn}(\kappa)$, $s_{nn}(\omega)$ and $f_{nn}(\omega)$ correspond to their time t levels.

The mapping between import and export shares at time t and the corresponding closed economy shocks is intuitive. First, if the import share of equipment type κ is high relative to κ' and trade elasticities are common across equipment goods, then moving

to autarky has an equivalent impact on relative wages as increasing the relative price of equipment κ relative to κ' in the closed economy. Second, if occupation ω has a low export share relative to occupation ω' , then moving to autarky is equivalent to increasing occupation demand for ω relative to ω' in the closed economy. Third, if occupation ω has a high import share relative to occupation ω' , trade elasticities are common across occupations, and $\eta(\omega) \geq \rho$ then moving to autarky has an equivalent impact on relative wages as increasing occupation demand for ω relative to ω' in the closed economy. If $\rho = 1$, then $\hat{\mu}_n(\omega)$ is equal to the ratio of the domestic absorption share to the domestic production share of occupation ω .

In conducting the counterfactual described above, we evaluate separately the impact of changes in trade costs in equipment (assuming that international trade costs in occupations are infinite during the whole sample period) and changes in trade costs in occupations (assuming that international trade costs in equipment are infinite during the whole sample period), imposing the values of θ and ρ that we estimated in Section 4.3.²⁷

6.1 Trade in Equipment

In order to compute the impact on relative wages in the US of counterfactual changes in international trade costs in equipment, we need measures of U.S. domestic expenditure shares by equipment type, $s_{nn}(\kappa)$, in 1984 and 2003, and estimates of the elasticity of substitution across countries of origin for each equipment type, $\eta(\kappa)$. For each equipment type κ , we measure $s_{nn}(\kappa)$ as

$$s_{nn,t}(\kappa) = 1 - \frac{Imports_{n,t}(\kappa)}{GrossOutput_{n,t}(\kappa) - Exports_{n,t}(\kappa) + Imports_{n,t}(\kappa)}.$$

We obtain Production, Export, and Import data from the OECD's Structural Analysis Database (STAN), which is arranged at the 2-digit level of the third revision of the International Standard Industrial Classification. We equate computers to industry 30 (Office, Accounting, and Computing Machinery) and non-computer equipment to industries 29, 31, 32, and 33 (Machinery and Equipment less Office, Accounting, and Computing Machinery). We observe that domestic expenditure shares for computers fell from 0.80 in 1984 to 0.26 in 2003, while domestic expenditure shares for non-computer equipment fell from 0.83 in 1984 to 0.59 in 2003; that is, the U.S. import share in computers rose significantly more than in non-computer equipment goods. Because of the lack of existing

²⁷In an open economy, without additional restrictions, it is possible that the exclusion restrictions underlying the identification of θ and ρ in the procedure described in Section 4.3 are violated.

estimates of computer and non-computer equipment trade elasticities in the literature, we estimate these elasticities using longitudinal data on trade volumes and import tariffs for computer and non-computer equipment manufacturing sectors.²⁸ Our baseline estimates of the trade elasticities are 3.2 for computers and 3.4 for non-computer equipment sectors.

	2003 to Autarky	1984 to Autarky	Impact of trade 84-03
Skill premium	-0.022	-0.001	0.021
Gender Gap	0.005	0.000	-0.005
HS dropout / Average	0.020	0.001	-0.019
HS grad / Average	0.007	0.000	-0.007
Some college / Average	-0.002	0.000	0.002
College / Average	-0.016	0.000	0.015
Grad training / Average	-0.017	-0.001	0.016

Table 8: Impact of moving from 2003 and 1984 to equipment autarky as well as the impact of trade in equipment between 1984 and 2003.

The high import share of computers relative to non-computer equipment in 2003 (together with the similar estimates of the trade elasticity for computers and non-computer equipment) implies that an increase in the international trade costs of equipment that makes trade in equipment prohibitively costly has an effect equivalent to that of an increase in the price of computers relative to non-computer equipment in the closed economy model; see equation (33). Column 1 in Table 8 shows that the skill premium would have fallen by roughly 2.2 percent moving from 2003 levels of equipment trade to autarky. On the other hand, given that U.S. import shares in 1984 in computer and non-computer equipment were much more similar, the impact on relative wages of moving to autarky in 1984 would have been much smaller. Column 3 combines the two previous trade-to-autarky counterfactuals to quantify how important was the rise of trade in equipment in generating relative wage changes between 1984 and 2003. Given changes in worldwide primitives between 1984 and 2003, our model predicts that the rise in the U.S. skill premium between these two years was 2.1 percent higher than it would have been had the

²⁸Specifically, for the period 1999-2003, we collect data on trade flows and effective tariffs from WITS by 2-digit rev. 3 ISIC sector and, in the case of non-computer equipment, aggregate the observations corresponding to industries 29, 31, 32, and 33. Given this data, we use ordinary least squares to estimate $\eta(\kappa)$ from the regression $\ln(X_{nit}(\kappa)) = f(n, t) + f(i, t) + f(n, i) - (\eta(\kappa) - 1) \ln(1 + d_{nit}(\kappa)) + \varepsilon_{nit}(\kappa)$, where $f(n, t)$, $f(i, t)$, and $f(n, i)$ denote, respectively, exporter-year, importer-year and exporter-importer fixed effects, $X_{nit}(\kappa)$ denotes the volume of trade from country n to country i at period t , and $d_{nit}(\kappa)$ denotes the effective tariff imposed by country n for imports of good κ from country i . Our point estimate of $\eta(\kappa) - 1$ is equal to 3.2 (0.56 standard error) for computers and 3.4 (0.22) for non-computer equipment.

US been in autarky over this time period. This 2.1 percentage point increase in the skill premium accounts for roughly 14 percent of the rise in the skill premium between 1984 and 2003 accounted for by changes in equipment productivity according to our closed economy calculations. The increase in the returns to college and graduate training (relative to the overall average change in wages) accounted for by trade in equipment was roughly 1.5 percent and the decrease in the gender gap was 0.5 percent. These numbers roughly double if we consider a lower trade elasticity of $\eta(\kappa) - 1 = 1.5$.

6.2 Trade in Occupations

In order to compute the impact on relative wages in the US of counterfactual changes in international trade costs in occupation services, we need both measures of U.S. domestic expenditure shares and domestic output shares by occupation ($s_{nn}(\omega)$ and $f_{nn}(\omega)$, respectively) in 1984 and 2003, and estimates of the elasticity of substitution across countries of origin for each occupation ($\eta(\omega)$). Obtaining these measures is a major challenge given the lack of readily available data on the occupation content of exports and imports; for a full discussion of the difficulties, see [Grossman and Rossi-Hansberg \(2007\)](#).

One possible route to measure the occupation content of exports and imports is to combine readily available information on trade by sector and on the occupation composition of labor payments of each sector, calculating the domestic share of absorption and the domestic share of output of occupation ω in country n as

$$s_{nn,t}(\omega) = 1 - \frac{\sum_{\sigma} v_{n,t}(\omega|\sigma) Imports_{n,t}(\sigma)}{\sum_{\sigma} v_{n,t}(\omega|\sigma) (GrossOutput_{n,t}(\sigma) - Exports_{n,t}(\sigma) + Imports_{n,t}(\sigma))} \quad (35)$$

$$f_{nn,t}(\omega) = 1 - \frac{\sum_{\sigma} v_{n,t}(\omega|\sigma) Exports_{n,t}(\sigma)}{\sum_{\sigma} v_{n,t}(\omega|\sigma) GrossOutput_{n,t}(\sigma)} \quad (36)$$

where $v_{n,t}(\omega|\sigma)$ denotes the share of labor payments to occupation ω in sector σ .²⁹ This approach is analogous to the basic calculation of the factor content of trade, replacing pay-

²⁹In Appendix H we present a version of our model that motivates these measures. In the model, the final good is produced according to a CES aggregator of tradable sectoral goods (allowing for time-varying shifters to the absorption of each sector), where each sectoral good is produced as the final good in our baseline model; i.e. according to a CES aggregator of non-traded occupation services. Occupations are produced exactly as in our baseline specification: a worker's productivity depends only on her occupation, and not on her sector of employment. The system of equations that solves for counterfactual wage changes caused by moving to autarky corresponds to the system of equations described for the closed economy version in which sectoral shifters are functions of sectoral export and import shares similarly to the expression for $\hat{\mu}_n(\omega)$ in equation (34). Moreover, if both elasticities of substitution across sectors and occupations are equal to one, then changes in relative wages in the autarky counterfactual in the model with trade in sectoral goods is identical to that in the model with trade in occupation services in which $s_{nn,t}(\omega)$ and $f_{nn,t}(\omega)$ are calculated using sectoral trade data as described above.

ments to a factor within each sector with payments to an occupation within each sector. We implement this approach using data for 30 2-digit rev. 3 ISIC manufacturing sectors and one aggregate services sector (obtaining U.S. export and import data from UN Comtrade, gross output data from UNIDO, and labor payments by sector and occupation from the CPS) and imposing a common trade elasticity across occupations of 5. Given that sectors in which the US has a comparative advantage tend to be intensive in occupations that disproportionately employ high skill workers, moving from trade to autarky tends to reduce the relative wage of more educated workers. However, the implied wage changes are very small. For example, starting in 2003, moving from trade to autarky reduces the skill premium by only 0.4 percent.

The computation of import and export shares by occupation in equations (35) and (36) assumes that, for each sector, exports and domestic production replaced by imports have the same occupation intensity in labor payments as does gross output. This assumption may be violated in practice. For example, if, within each sector, the US exports tasks that are more intensive in, e.g., abstract occupations than is gross output and imports tasks that are more intensive in, e.g., routine occupations than is gross output, then equation (35) will overstate (understate) the domestic share of absorption for routine (abstract) occupations and equation (36) will understate (overstate) the domestic share of output of routine (abstract) occupations. Hence, domestic absorption and output shares constructed as in equations (35) and (36) are likely to be biased.³⁰

We consider an alternative route to measure export and import shares by occupation in 2003 (that does not use information on trade by sector), making strong assumptions to classify occupations as exported, imported, and non-traded. Specifically, we split our set of 30 occupations, Ω , into three groups: the set of abstract occupations (professional, managerial and technical occupations) denoted by Ω^A , the set of routine occupations (clerical, administrative support, production, and operative occupations) denoted by Ω^R , and the set of non-traded occupations (protective services, food preparation, cleaning, personal services, and sales) denoted by Ω^N .³¹ We assume that abstract occupations are the only occupations that are exported and that all occupations within this set have the same export share; that routine occupations are the only occupations that are imported and that all occupations within this set have the same import share; and that non-traded occupations have a domestic share of absorption and output equal to one. For abstract

³⁰Closely related, [Burstein and Vogel \(2017\)](#) argue that measures of the factor content of trade constructed under the assumption that sectoral factor intensities are equal across destinations understate the impact of trade on the skill premium.

³¹This is the grouping used in [Acemoglu and Autor \(2011\)](#), except that we include sales in the group of non-traded occupations.

occupations, $\omega \in \Omega^A$, we set

$$f_{nn,t}(\omega) = 1 - \frac{Exports_{n,t}}{v_{n,t}(A) \times GrossOutput_{n,t}}$$

and $s_{nn,t}(\omega) = 1$, where $v_{n,t}(A)$ denotes labor payments across all abstract occupations relative to total labor payments. For routine occupations, $\omega \in \Omega^R$, we set $f_{nn,t}(\omega) = 1$ and

$$s_{nn,t}(\omega) = 1 - \frac{Imports_{n,t}}{v_{n,t}(R) \times GrossOutput_{n,t} + Imports_{n,t}}.$$

For non-traded occupations, $\omega \in \Omega^N$, we set $s_{nn,t}(\omega) = f_{nn,t}(\omega) = 1$. We calculate the share of labor payments in abstract and routine occupations from the CPS and total gross output, exports and imports from the BEA.³²

	2003 to Autarky	1984 to Autarky	Impact of trade 84-03
Skill premium	-0.063	-0.051	0.013
Gender Gap	0.010	0.003	-0.006
HS dropout / Average	0.025	0.023	-0.001
HS grad / Average	0.022	0.018	-0.004
Some college / Average	0.005	0.002	-0.004
College / Average	-0.037	-0.030	0.007
Grad training / Average	-0.064	-0.051	0.014

Table 9: Impact of moving from 2003 and 1984 to occupation autarky as well as the impact of trade in occupations between 1984 and 2003.

Here we combine occupations into three groups—abstract (which we assume are exported), routine (which we assume are imported), and non traded—and assume that all occupations within each group have the same import and export shares.

The results from this counterfactual are reported in Table 9. Recall from equation (34) that making trade in occupations prohibitively costly operates in our model exactly like a decline in occupation demand in the closed economy for those occupations with relatively high export shares and/or low import shares. Since abstract occupations (assumed to be only exported) tend to employ relatively more skilled workers than routine occupations (assumed to be only imported), the skill premium would have fallen by 6.3 percent if the US had moved to autarky in 2003, as reported in Column 1 of Table 9. The returns to graduate training (relative to the average change in wages) would have fallen by 6.4 percent, and the gender gap would have increased by 1 percent. These numbers are

³²Since the extent of offshoring of routine tasks was much less prevalent in the 1980s than in the 2000s (see e.g. Grossman and Rossi-Hansberg 2008), we believe that the strong assumptions that we use to construct occupation trade shares disproportionately increase the impact of moving to autarky in 1984.

an order of magnitude larger than those reported above using measures of occupation trade constructed from sectoral trade flows. If the US had moved to autarky in 1984, the skill premium would have fallen by 5.1 percent. Column 3 combines these two trade-to-autarky counterfactuals to quantify how important was the rise of trade in occupations in generating relative wage changes between 1984 and 2003. Given changes in world-wide primitives between 1984 and 2003, our model predicts that the rise in the U.S. skill premium between these two years was 1.3 percent higher than it would have been had the US been in autarky during this time period. This 1.3 percent increase in the skill premium accounts for roughly 26 percent of the rise in the skill premium between 1984 and 2003 accounted for by occupation shifters in our closed economy calculations. Given the crudeness of our measures of occupation trade, these calculations should be viewed as a first step to evaluate the role that occupation trade has had on the evolution of between-group inequality.

7 Conclusions

We provide an assignment model of the labor market extended to incorporate equipment types as another key dimension along which workers sort and international trade as a determinant of the equilibrium assignment of workers. We thus study within a unified framework the impact of changes in occupation demand shifters, computer productivity, labor productivity, labor composition, and international trade on changes in the relative average wages of multiple groups of workers.

We parameterize our model using detailed measures of computer usage within labor group-occupation pairs in the US between 1984 and 2003 as well as international trade data. Using the closed-economy version of our model, we conclude that computerization alone accounts for the majority of the observed rise in between-education-group inequality over this period and the differential growth of the skill premium between old and younger workers. The combination of computerization and occupation shifters explains roughly 80 percent of the rise in the skill premium, almost all of the rise in inequality across more disaggregated education groups, and the majority of the fall in the gender gap. We show how computerization and changes in occupation demand in our model can be caused by changes in the extent of international trade, and perform counterfactual exercises to quantify these effects.

The focus of this paper has been on the distribution of labor income between groups of workers with different observable characteristics. A fruitful avenue for future research is to extend our framework to address the changing distribution of income accruing to

labor and capital, as analyzed in e.g. Karabarbounis and Neiman (2014) and Oberfield and Raval (2014), as well as the changing distribution of income across workers within groups, as analyzed in e.g. Huggett et al. (2011), Hornstein et al. (2011), and Helpman et al. (2017).

References

- ACEMOGLU, D. (2002): "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, 40, 7–72.
- ACEMOGLU, D. AND D. AUTOR (2011): *Skills, Tasks and Technologies: Implications for Employment and Earnings*, Elsevier, vol. 4 of *Handbook of Labor Economics*, chap. 12, 1043–1171.
- ADAO, R. (2015): "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil," *mimeo*.
- AUTOR, D. (2014): "Skills, Education, and the Rise of Earnings Inequality Among the 'Other 99 Percent'," *Science*, 344, 843–851.
- AUTOR, D. H. AND D. DORN (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103, 1553–97.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content Of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118, 1279–1333.
- BEAUDRY, P. AND E. LEWIS (2014): "Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980-2000," *American Economic Journal: Applied Economics*, 6, 178–94.
- BUERA, F. J., J. P. KABOSKI, AND R. ROGERSON (2015): "Skill-Biased Structural Change and the Skill-Premium," *Working Paper*.
- BURSTEIN, A., J. CRAVINO, AND J. VOGEL (2013): "Importing Skill-Biased Technology," *American Economic Journal: Macroeconomics*, 5, 32–71.
- BURSTEIN, A. AND J. VOGEL (2017): "International Trade, Technology, and the Skill Premium," *Journal of Political Economy*, 125, 1356–1412.
- CARD, D. AND T. LEMIEUX (2001): "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis," *The Quarterly Journal of Economics*, 116, 705–746.
- COSTINOT, A. AND J. VOGEL (2010): "Matching and Inequality in the World Economy," *Journal of Political Economy*, 118, 747–786.
- (2015): "Beyond Ricardo: Assignment Models in International Trade," *Annual Review of Economics*, 7, 31–62.
- DEKLE, R., J. EATON, AND S. KORTUM (2008): "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, 55, 511–540.

- DINARDO, J. E. AND J.-S. PISCHKE (1997): "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?" *The Quarterly Journal of Economics*, 112, 291–303.
- DIX-CARNEIRO, R. AND B. K. KOVAK (2015): "Trade liberalization and the skill premium: A local labor markets approach," *The American Economic Review*, 105, 551–557.
- EATON, J. AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.
- ENTORF, H., M. GOLLAC, AND F. KRAMARZ (1999): "New Technologies, Wages, and Worker Selection," *Journal of Labor Economics*, 17, 464–491.
- GALLE, S., A. RODRÍGUEZ-CLARE, AND M. YI (2015): "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," *Working Paper*.
- GOOS, M., A. MANNING, AND A. SALOMONS (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104, 2509 – 2526.
- GORDON, R. J. (1990): *The Measurement of Durable Goods Prices*, National Bureau of Economic Research Research Monograph Series.
- GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (1997): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87, 342–62.
- GROSSMAN, G. AND E. ROSSI-HANSBERG (2007): *The Rise of Offshoring: It's Not Wine for Cloth Anymore*, Federal Reserve Bank of Kansas City, The New Economic Geography: Effects and Policy Implications.
- GROSSMAN, G. M. AND E. ROSSI-HANSBERG (2008): "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98, 1978–97.
- HELPMAN, E., O. ITSKHOKI, M.-A. MUENDLER, AND S. REDDING (2017): "Trade and Inequality: From Theory to Estimation," *Review of Economic Studies*, 84, 357–405.
- HORNSTEIN, A., P. KRUSELL, AND G. L. VIOLANTE (2011): "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *American Economic Review*, 101, 2873–98.
- HSIEH, C.-T., E. HURST, C. I. JONES, AND P. J. KLENOW (2016): "The Allocation of Talent and U.S. Economic Growth," Tech. rep., Mimeo.
- HUGGETT, M., G. VENTURA, AND A. YARON (2011): "Sources of Lifetime Inequality," *American Economic Review*, 101, 2923–54.
- KAMBOUROV, G. AND I. MANOVSKII (2009a): "Occupational Mobility and Wage Inequality," *Review of Economic Studies*, 76, 731–759.
- (2009b): "Occupational Specificity Of Human Capital," *International Economic Review*, 50, 63–115.
- KARABARBOUNIS, L. AND B. NEIMAN (2014): "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 129, 61–103.

- KATZ, L. F. AND K. M. MURPHY (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107, 35–78.
- KRUEGER, A. B. (1993): "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989," *The Quarterly Journal of Economics*, 108, 33–60.
- KRUSELL, P., L. E. OHANIAN, J.-V. RIOS-RULL, AND G. L. VIOLANTE (2000): "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68, 1029–1054.
- LAGAKOS, D. AND M. E. WAUGH (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, 103, 948–80.
- LEE, D. AND K. I. WOLPIN (2010): "Accounting for wage and employment changes in the US from 1968-2000: A dynamic model of labor market equilibrium," *Journal of Econometrics*, 156, 68–85.
- LEE, E. (2016): "Trade, Inequality, and the Endogenous Sorting of Heterogeneous Workers," Tech. rep.
- LEMIEUX, T. (2006): "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *American Economic Review*, 96, 461–498.
- OBERFIELD, E. AND D. RAVAL (2014): "Micro Data and Macro Technology," NBER Working Papers 20452, National Bureau of Economic Research, Inc.
- PARRO, F. (2013): "Capital-Skill Complementarity and the Skill Premium in a Quantitative Model of Trade," *American Economic Journal: Macroeconomics*, 5, 72–117.
- RAMONDO, N. AND A. RODRÍGUEZ-CLARE (2013): "Trade, Multinational Production, and the Gains from Openness," *Journal of Political Economy*, 121, 273 – 322.
- ROY, A. (1951): "Some thoughts on the distribution of earnings," *Oxford Economic Papers*, 3, 135–146.
- SAEZ, E. (2001): "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 68, 205–29.
- SAMPSON, T. (2014): "Selection into Trade and Wage Inequality," *American Economic Journal: Microeconomics*, 6, 157–202.
- TRAIBERMAN, S. (2016): "Occupations and Import Competition," Working Paper New York University.

A Derivations

Here we derive the equations in Section 3.2 for (i) the wage per efficiency unit of labor λ when teamed with equipment κ in occupation ω , $v_t(\lambda, \kappa, \omega)$; (ii) the probability that a randomly sampled worker, $z \in \mathcal{Z}_t(\lambda)$, uses equipment κ in occupation ω , $\pi_t(\lambda, \kappa, \omega)$; and (iii) the average wage of workers in group λ teamed with equipment κ in occupation ω , $w_t(\lambda, \kappa, \omega)$. We also show that $w_t(\lambda) = w_t(\lambda, \kappa, \omega)$ for all (κ, ω) .

Wage per efficiency unit of labor λ : $v_t(\lambda, \kappa, \omega)$. An occupation production unit hiring k units of equipment κ and l efficiency units of labor λ earns revenues $p_t(\omega) k^\alpha [T_t(\lambda, \kappa, \omega) l]^{1-\alpha}$ and incurs costs $p_t(\kappa) k + v_t(\lambda, \kappa, \omega) l$. The first-order condition for the optimal choice of k per unit of l for a given $(\lambda, \kappa, \omega)$ yields

$$k_t(l; \lambda, \kappa, \omega) = \left(\alpha \frac{p_t(\omega)}{p_t(\kappa)} \right)^{\frac{1}{1-\alpha}} T_t(\lambda, \kappa, \omega) l,$$

where the second-order condition is satisfied for any $\alpha < 1$. This implies that the production unit's revenue can be expressed as $p_t(\omega) k^\alpha [T_t(\lambda, \kappa, \omega) l]^{1-\alpha}$ and its cost as $[p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} (\alpha p_t(\omega))^{\frac{1}{1-\alpha}} T_t(\lambda, \kappa, \omega) + v_t(\lambda, \kappa, \omega)] l$. Hence, the zero profit condition requires that

$$v_t(\lambda, \kappa, \omega) = (1 - \alpha) \alpha^{\frac{\alpha}{1-\alpha}} p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} p_t(\omega)^{\frac{1}{1-\alpha}} T_t(\lambda, \kappa, \omega)$$

which is equivalent to the value in Section 3.2 given the definition of $\bar{\alpha} \equiv (1 - \alpha) \alpha^{\frac{\alpha}{1-\alpha}}$.

Labor allocation: $\pi_t(\lambda, \kappa, \omega)$. In what follows, denote by $\varphi \equiv (\kappa, \omega)$. A worker $z \in \mathcal{Z}_t(\lambda)$ chooses equipment-occupation pair φ if

$$v_t(\lambda, \varphi) \varepsilon(z, \varphi) > \max_{\varphi' \neq \varphi} \{v_t(\lambda, \varphi') \varepsilon(z, \varphi')\},$$

which is independent of $\varepsilon(z)$. The probability that a randomly sampled worker in group λ chooses φ is

$$\begin{aligned} \pi_t(\lambda, \varphi) &= \int_0^\infty \Pr \left[\varepsilon > \max_{\varphi' \neq \varphi} \left\{ \frac{\varepsilon(z, \varphi') v_t(\lambda, \varphi')}{v_t(\lambda, \varphi)} \right\} \right] dG(\varepsilon) \\ &= \int_0^\infty \exp \left[-\varepsilon^{-\theta(\lambda)} \left(\sum_{\varphi'} \left(\frac{v_t(\lambda, \varphi)}{v_t(\lambda, \varphi')} \right)^{-\theta(\lambda)} \right) \right] \theta(\lambda) \varepsilon^{-1-\theta(\lambda)} d\varepsilon \end{aligned}$$

Defining $n_t(\lambda, \varphi) \equiv \left(\sum_{\varphi'} \left(\frac{v_t(\lambda, \varphi)}{v_t(\lambda, \varphi')} \right)^{-\theta(\lambda)} \right)$, we have

$$\begin{aligned} \pi_t(\lambda, \varphi) &= - \int_0^\infty \exp\left(-\varepsilon^{-\theta(\lambda)} n_t(\lambda, \varphi)\right) (-\theta(\lambda)) \varepsilon^{-1-\theta(\lambda)} d\varepsilon \\ &= \frac{1}{n_t(\lambda, \varphi)} \end{aligned}$$

Substituting back in for $n_t(\lambda, \varphi)$, we have

$$\pi_t(\lambda, \varphi) = \frac{v_t(\lambda, \varphi)^{\theta(\lambda)}}{\sum_{\varphi'} v_t(\lambda, \varphi')^{\theta(\lambda)}} \quad (37)$$

Finally, substituting back for φ and for $v_t(\lambda, \kappa, \omega)$ and setting $\theta(\lambda) = \theta$ for all λ , we obtain equation (3) in Section 3.2.

Average wages: $w_t(\lambda, \kappa, \omega)$ and $w_t(\lambda)$. As in the previous derivation, denote by $\varphi \equiv (\kappa, \omega)$. The average efficiency units of each worker in $\mathcal{Z}_t(\lambda, \varphi)$, which denotes the set of workers $z \in \mathcal{Z}_t(\lambda)$ who choose φ , is

$$\begin{aligned} \mathbb{E}[\varepsilon(z) \varepsilon(z, \varphi) | z \in \mathcal{Z}_t(\lambda, \varphi)] &= \mathbb{E}[\varepsilon(z) | z \in \mathcal{Z}_t(\lambda, \varphi)] \times \mathbb{E}[\varepsilon(z, \varphi) | z \in \mathcal{Z}_t(\lambda, \varphi)] \\ &= \mathbb{E}[\varepsilon(z) | z \in \mathcal{Z}_t(\lambda)] \times \mathbb{E}[\varepsilon(z, \varphi) | z \in \mathcal{Z}_t(\lambda, \varphi)] \\ &= \mathbb{E}[\varepsilon(z, \varphi) | z \in \mathcal{Z}_t(\lambda, \varphi)]. \end{aligned}$$

where the first two equalities follow from both the assumption that $\varepsilon(z)$ is independent of $\varepsilon(z, \varphi)$ and the result above that the choice φ of each individual z does not depend on the value of $\varepsilon(z)$ and the third equality follows from normalizing $\mathbb{E}[\varepsilon(z) | z \in \mathcal{Z}_t(\lambda)]$ to be equal to 1.³³ In what follows let $\bar{\varepsilon}_t(\lambda, \varphi) \equiv \mathbb{E}[\varepsilon(z, \varphi) | z \in \mathcal{Z}_t(\lambda, \varphi)]$. We have

$$\bar{\varepsilon}_t(\lambda, \varphi) = \frac{1}{\pi_t(\lambda, \varphi)} \int_0^\infty \varepsilon \times \Pr \left[\varepsilon \geq \max_{\varphi' \neq \varphi} \left\{ \frac{\varepsilon(z, \varphi') v_t(\lambda, \varphi')}{v_t(\lambda, \varphi)} \right\} \right] \times dG(\varepsilon)$$

Hence, we have

$$\bar{\varepsilon}_t(\lambda, \varphi) = \frac{1}{\pi_t(\lambda, \varphi)} \int_0^\infty \exp \left[-\varepsilon^{-\theta(\lambda)} v_t(\lambda, \varphi)^{-\theta(\lambda)} \sum_{\varphi'} v_t(\lambda, \varphi')^{\theta(\lambda)} \right] \theta(\lambda) \varepsilon^{-\theta(\lambda)} d\varepsilon$$

³³This is a normalization because, for any value of $T'_t(\lambda)$ and $\mathbb{E}[\varepsilon'_t(z) | z \in \mathcal{Z}_t(\lambda)] \neq 1$, we can always define an alternative $T_t(\lambda)$ and $\varepsilon(z)$ such that $T_t(\lambda) = T'_t(\lambda) \times \mathbb{E}[\varepsilon'_t(z) | z \in \mathcal{Z}_t(\lambda)]$ and $\mathbb{E}[\varepsilon(z) | z \in \mathcal{Z}_t(\lambda)] = 1$.

Let $j = \varepsilon^{-\theta}$ and, as in the previous derivation, let $n_t(\lambda, \varphi) \equiv \left(\sum_{\varphi'} \left(\frac{v_t(\lambda, \varphi)}{v_t(\lambda, \varphi')} \right)^{-\theta(\lambda)} \right)$. Hence, we have

$$\bar{\varepsilon}_t(\lambda, \varphi) = \frac{1}{\pi_t(\lambda, \varphi)} \int_{-\infty}^0 \exp(-jz) j^{-1/\theta(\lambda)} (-dj) = \frac{1}{\pi_t(\lambda, \varphi)} \int_0^{\infty} j^{-1/\theta(\lambda)} \exp(-jz) dj$$

Let $y_t(\lambda, \varphi) = n_t(\lambda, \varphi) j$. Hence, we have

$$\begin{aligned} \bar{\varepsilon}_t(\lambda, \varphi) &= \frac{1}{\pi_t(\lambda, \varphi)} \int_0^{\infty} \left(\frac{y_t(\lambda, \varphi)}{n_t(\lambda, \varphi)} \right)^{-1/\theta(\lambda)} \exp(-y_t(\lambda, \varphi)) \frac{dy_t(\lambda, \varphi)}{n_t(\lambda, \varphi)} \\ &= \frac{1}{\pi_t(\lambda, \varphi)} n_t(\lambda, \varphi)^{\frac{1-\theta(\lambda)}{\theta(\lambda)}} \times \gamma(\lambda) \end{aligned}$$

where $\gamma(\lambda) \equiv \Gamma\left(1 - \frac{1}{\theta(\lambda)}\right)$ and $\Gamma(\cdot)$ is the Gamma function

$$\Gamma(x) \equiv \int_0^{\infty} t^{x-1} \exp(-t) dt.$$

Substituting in for $n_t(\lambda, \varphi)$, we obtain

$$\bar{\varepsilon}_t(\lambda, \varphi) = \gamma(\lambda) \pi_t(\lambda, \varphi)^{\frac{-1}{\theta(\lambda)}}.$$

Hence, the total income of workers in $\mathcal{Z}_t(\lambda)$ choosing φ , $L_t(\lambda) \pi_t(\lambda, \varphi) \bar{\varepsilon}_t(\lambda, \varphi) v_t(\lambda, \varphi)$, becomes $\gamma(\lambda) L_t(\lambda) \pi_t(\lambda, \varphi)^{\frac{\theta(\lambda)-1}{\theta(\lambda)}} v_t(\lambda, \varphi)$. Dividing by the mass of these workers, $L_t(\lambda) \pi_t(\lambda, \varphi)$, we obtain the wage rate

$$w_t(\lambda, \varphi) = \gamma(\lambda) v_t(\lambda, \varphi) \pi_t(\lambda, \varphi)^{-1/\theta(\lambda)}. \quad (38)$$

Substituting in for $v_t(\lambda, \varphi)$ and for φ and setting $\theta(\lambda) = \theta$ and $\gamma(\lambda) = \gamma$ for all λ , we obtain the un-numbered equation from Section 3.2:

$$w_t(\lambda, \kappa, \omega) = \bar{\alpha} \gamma T_t(\lambda, \kappa, \omega) p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} p_t(\omega)^{\frac{1}{1-\alpha}} \pi_t(\lambda, \kappa, \omega)^{-1/\theta}.$$

Finally, substituting in for $\pi_t(\lambda, \varphi)$ from equation (37) into equation (38), we obtain

$$w_t(\lambda) = w_t(\lambda, \varphi) = \gamma(\lambda) \left(\sum_{\varphi'} v_t(\lambda, \varphi')^{\theta(\lambda)} \right)^{1/\theta(\lambda)}.$$

Substituting in for $v_t(\lambda, \varphi)$ and for φ and setting $\theta(\lambda) = \theta$ and $\gamma(\lambda) = \gamma$ for all λ , we

obtain equation (4) from Section 3.2.

B Data details

Throughout, we restrict our sample by dropping workers who are younger than 17 or older than 64 years old, do not report positive paid hours worked, are self-employed, or are in the military.

MORG. We use the MORG CPS to form a sample of hours worked and income for each labor group. Specifically, we use the “hour wage sample” from [Acemoglu and Autor \(2011\)](#). Hourly wages are equal to the reported hourly earnings for those paid by the hour and the usual weekly earnings divided by hours worked last week for non-hourly workers. Top-coded earnings are multiplied by 1.5. Workers earning below \$1.675/hour in 1982 dollars are dropped, as are workers whose hourly wages exceed the number arising from multiplying the top-coded value of weekly earnings by 1/35 (i.e., workers paid by the hour whose wages are sufficiently high so that their weekly income would be top-coded if they worked at least 35 hours and were not paid by the hour). Observations with allocated earnings are excluded. Our measure of labor composition, $L_t(\lambda)$, is hours worked within each labor group λ (weighted by sample weights).

October Supplement. In 1984, 1989, 1993, 1997, and 2003, the October Supplement asked respondents whether they “have direct or hands on use of computers at work,” “directly use a computer at work,” or “use a computer at/for his/her/your main job.” Using a computer at work refers only to “direct” or “hands on” use of a computer with typewriter-like keyboards, whether a personal computer, laptop, mini computer, or mainframe.

Occupations. The occupations we include are listed in Table 10, where we also list the share of hours worked in each occupation by college educated workers and by women as well as the occupation share of labor payments in 1984 and in 2003. Our concordance of occupations across time is based on the concordance developed in [Autor and Dorn \(2013\)](#).

Composition-adjusted wages. We construct thirty labor groups defined by the intersection of five education, two gender, and three age categories. When we construct measures of changes in relative wages between broader groups that aggregate across our most disaggregated labor groups—e.g. the group of college educated workers combines ten of our thirty labor groups—we composition adjust wages by holding constant the relative employment shares of our thirty labor groups across all years of the sample. Specifically, after calculating mean log wages within each labor group (either from the model or the data), we construct mean wages for broader groups as fixed-weighted averages of the

Occupations	College intensity		Female intensity		Income share	
	1984	2003	1984	2003	1984	2003
Executive, administrative, managerial	0.48	0.58	0.32	0.40	0.12	0.16
Management related	0.53	0.61	0.43	0.54	0.05	0.05
Architect	0.86	0.88	0.15	0.24	0.00	0.00
Engineer	0.71	0.79	0.06	0.11	0.03	0.03
Life, physical, and social science	0.65	0.55	0.30	0.30	0.01	0.02
Computer and mathematical	0.86	0.91	0.31	0.41	0.01	0.01
Community and social services	0.76	0.73	0.46	0.58	0.01	0.02
Lawyers	0.98	0.98	0.24	0.35	0.01	0.01
Education, training, etc...*	0.90	0.87	0.63	0.68	0.05	0.06
Arts, design, entertainment, sports, media	0.49	0.57	0.39	0.45	0.01	0.01
Health diagnosing	0.96	0.98	0.20	0.33	0.01	0.01
Health assessment and treating	0.51	0.64	0.85	0.84	0.02	0.04
Technicians and related support	0.30	0.43	0.46	0.43	0.04	0.05
Financial sales and related	0.31	0.33	0.31	0.40	0.04	0.05
Retail sales	0.17	0.24	0.54	0.50	0.05	0.05
Administrative support	0.12	0.16	0.78	0.74	0.14	0.12
Housekeeping, cleaning, laundry	0.01	0.03	0.83	0.83	0.01	0.00
Protective service	0.16	0.21	0.11	0.19	0.02	0.02
Food preparation and service	0.05	0.06	0.61	0.51	0.02	0.02
Health service	0.04	0.08	0.90	0.90	0.01	0.01
Building, grounds cleaning, maintenance	0.04	0.05	0.20	0.21	0.02	0.01
Miscellaneous**	0.12	0.17	0.67	0.63	0.01	0.01
Child care	0.11	0.12	0.91	0.94	0.00	0.00
Agriculture and mining	0.05	0.06	0.10	0.16	0.01	0.01
Mechanics and repairers	0.04	0.07	0.03	0.04	0.05	0.04
Construction	0.04	0.05	0.01	0.02	0.05	0.04
Precision production	0.07	0.08	0.15	0.25	0.04	0.03
Machine operators, assemblers, inspectors	0.03	0.06	0.40	0.34	0.08	0.04
Transportation and material moving	0.03	0.05	0.06	0.09	0.05	0.04
Handlers, equip. cleaners, helpers, laborers	0.03	0.04	0.17	0.18	0.03	0.02

Table 10: Thirty occupations, their college and female intensities, and the occupational share of labor payments

*Education, training, etc... also includes library, legal support/assistants/paralegals

**Miscellaneous includes personal appearance, misc. personal care and service, recreation and hospitality

College intensity (Female intensity) indicates hours worked in the occupation by those with college degrees (females) relative to total hours worked in the occupation. Income share denotes labor payments in the occupation relative to total labor payments. Each is calculated using the MORG CPS.

relevant labor group means, using an average share of total hours worked by each labor group over 1984 to 2003 as weights. This adjustment ensures that changes in average wages across broader groups are not driven by shifts in the education \times age \times gender composition within these broader groups.

O*NET. We follow [Acemoglu and Autor \(2011\)](#) in our use of O*NET and construct composite measures of O*NET Work Activities and Work Context Importance scales: (i) non-routine cognitive analytical, (ii) non-routine cognitive interpersonal, (iii) routine manual, and (iv) non-routine manual physical. We aggregate their measures up to our thirty occupations and standardize each to have mean zero and standard deviation one.

C Measurement of shocks

Equations (11), (12), and (13) can be written so that changes in relative wages, $\hat{w}(\lambda) / \hat{w}(\lambda_1)$, relative transformed occupation price changes, $\hat{q}(\omega) / \hat{q}(\omega_1)$, and allocations, $\hat{\pi}(\lambda, \kappa, \omega)$, depend on relative shocks to labor composition, $\hat{L}(\lambda) / \hat{L}(\lambda_1)$, occupation shifters, $\hat{a}(\omega) / \hat{a}(\omega_1)$, equipment productivity, $\hat{q}(\kappa) / \hat{q}(\kappa_1)$, and labor productivity, $\hat{T}(\lambda) / \hat{T}(\lambda_1)$:

$$\frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} = \frac{\hat{T}(\lambda)}{\hat{T}(\lambda_1)} \frac{\left[\sum_{\kappa, \omega} \left(\frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} \frac{\hat{q}(\kappa)}{\hat{q}(\kappa_1)} \right)^\theta \pi_{t_0}(\lambda, \kappa, \omega) \right]^{1/\theta}}{\left[\sum_{\kappa', \omega'} \left(\frac{\hat{q}(\omega')}{\hat{q}(\omega_1)} \frac{\hat{q}(\kappa')}{\hat{q}(\kappa_1)} \right)^\theta \pi_{t_0}(\lambda, \kappa', \omega') \right]^{1/\theta}}$$

$$\hat{\pi}(\lambda, \kappa, \omega) = \frac{\left(\frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} \frac{\hat{q}(\kappa)}{\hat{q}(\kappa_1)} \right)^\theta}{\sum_{\kappa', \omega'} \left(\frac{\hat{q}(\omega')}{\hat{q}(\omega_1)} \frac{\hat{q}(\kappa')}{\hat{q}(\kappa_1)} \right)^\theta \pi_{t_0}(\lambda, \kappa', \omega')}$$

and

$$\frac{\hat{a}(\omega)}{\hat{a}(\omega_1)} \left(\frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} \right)^{(1-\alpha)(1-\rho)} = \frac{\zeta_{t_0}(\omega_1)}{\zeta_{t_0}(\omega)} \frac{\sum_{\lambda, \kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} \frac{\hat{L}(\lambda)}{\hat{L}(\lambda_1)} \hat{\pi}(\lambda, \kappa, \omega)}{\sum_{\lambda', \kappa'} w_{t_0}(\lambda') L_{t_0}(\lambda') \pi_{t_0}(\lambda', \kappa', \omega) \frac{\hat{w}(\lambda')}{\hat{w}(\lambda_1)} \frac{\hat{L}(\lambda')}{\hat{L}(\lambda_1)} \hat{\pi}(\lambda', \kappa', \omega)}.$$

Note that \hat{E} cancels out of this system of equations.

C.1 Baseline

Here we describe in detail the steps that we follow to obtain our measures of changes in labor composition, occupation shifters, equipment productivity and labor productivity.

First, the relative shocks to labor composition $\hat{L}(\lambda) / \hat{L}(\lambda_1)$ are directly observed in the data.

Second, we measure relative changes in equipment productivity (to the power θ) using equation (14) as

$$\frac{\hat{q}(\kappa_2)^\theta}{\hat{q}(\kappa_1)^\theta} = \exp \left(\frac{1}{N(\kappa_1, \kappa_2)} \sum_{\lambda, \omega} \log \frac{\hat{\pi}(\lambda, \kappa_2, \omega)}{\hat{\pi}(\lambda, \kappa_1, \omega)} \right),$$

dropping all (λ, ω) pairs for which $\pi_t(\lambda, \kappa_1, \omega) = 0$ or $\pi_t(\lambda, \kappa_2, \omega) = 0$ in either period t_0 or t_1 . $N(\kappa_1, \kappa_2)$ is the number of (λ, ω) pairs over which we average; in the absence of any zeros in allocations we have $N(\kappa_1, \kappa_2) = 900$, which is the number of labor groups multiplied by the number of occupations.

Third, we measure changes in transformed occupation prices relative to occupation ω_0 (to the power θ) using equation (16) as

$$\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_0)^\theta} = \exp \left(\frac{1}{N(\omega, \omega_0)} \sum_{\lambda, \kappa} \log \frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa, \omega_0)} \right).$$

dropping all (λ, κ) pairs for which $\pi_t(\lambda, \kappa, \omega_0) = 0$ or $\pi_t(\lambda, \kappa, \omega) = 0$ in either period t_0 or t_1 . $N(\omega, \omega_0)$ is the number of (λ, κ) pairs over which we average; in the absence of any zeros in allocations we have $N(\omega, \omega_0) = 60$, which is the number of labor groups multiplied by the number of equipment types. In our model, the estimates of the relative occupation shifters for any two occupations ω_A and ω_B should not depend on the choice of the reference category ω_0 . However, even if this prediction of the model is right, the fact that some of the values of $\pi_t(\lambda, \kappa, \omega_0)$ are equal to 0 in the data implies that the estimates of changes in relative transformed occupation prices (to the power θ) vary with the choice of ω_0 . In order to avoid this sensitivity to the choice of ω_0 , we compute changes in relative transformed occupation prices using the following geometric average

$$\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta} = \exp \left(\frac{1}{30} \sum_{\omega_0} \left(\log \frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_0)^\theta} - \log \frac{\hat{q}(\omega_1)^\theta}{\hat{q}(\omega_0)^\theta} \right) \right)$$

where, for each ω and ω_0 , $\hat{q}(\omega)^\theta / \hat{q}(\omega_0)^\theta$ is calculated as described above. This expression yields estimates that do not depend on the choice of ω_1 . Furthermore, as Section C.2 shows, this approach yields measures of relative changes in occupation shifters that are very similar to those that arise from projecting changes in allocations on a set of fixed effects.

Third, given our measures of changes in equipment and transformed occupation prices

(both to the power θ), we construct $\hat{s}(\lambda)$ using equation (18). Given $\hat{s}(\lambda)$, we estimate θ using equation (39) as described in Section 4.3.

Fourth, given the measures of changes in equipment productivity and transformed occupation prices (both to the power θ) in equations (14) and (16), the estimate of θ , and observed values both of the initial allocation $\pi_{t_0}(\lambda, \kappa, \omega)$ and changes in relative wages $\hat{w}(\lambda) / \hat{w}(\lambda_1)$, we measure changes in labor productivity $\hat{T}(\lambda) / \hat{T}(\lambda_1)$ using equation (17).

Fifth, using data on changes in payments to occupations, $\hat{\zeta}(\omega)$, the measures of changes in equipment productivity and transformed occupation prices (both to the power θ) in equations (14) and (16), and an estimate of θ , we estimate ρ as described in Section 4.3.

Finally, we measure changes in occupation shifters, $\hat{a}(\omega) / \hat{a}(\omega_1)$, using equation (15). A variable in this equation is the relative changes in total payments to occupations ω relative to those in a benchmark occupation ω_1 , $\hat{\zeta}(\omega) / \hat{\zeta}(\omega_1)$. We construct this variable as follows. The initial levels, $\zeta_{t_0}(\omega) / \zeta_{t_0}(\omega_1)$, are calculated directly using the observed values of $\pi_{t_0}(\lambda, \kappa, \omega)$, $w_{t_0}(\lambda)$, and $L_{t_0}(\lambda)$. The terminal levels, $\zeta_{t_1}(\omega) / \zeta_{t_1}(\omega_1)$, are constructed as

$$\frac{\zeta_{t_1}(\omega)}{\zeta_{t_1}(\omega_1)} = \frac{\sum_{\lambda, \kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega)}{\sum_{\lambda', \kappa'} w_{t_0}(\lambda') L_{t_0}(\lambda') \pi_{t_0}(\lambda', \kappa', \omega_1) \frac{\hat{w}(\lambda')}{\hat{w}(\lambda_1)} \hat{L}(\lambda') \hat{\pi}(\lambda', \kappa', \omega_1)},$$

where $\hat{\pi}(\lambda, \kappa, \omega)$ are those constructed by the model given the measures of $\hat{q}(\omega)^\theta / \hat{q}(\omega)^\theta$ and $\hat{q}(\kappa)^\theta / \hat{q}(\kappa_1)^\theta$. The correlation between $\log(\hat{\zeta}(\omega) / \hat{\zeta}(\omega_1))$ implied by the model and in the data is 0.77 between 1984 and 2003; the correlation between $\log(\hat{\zeta}(\omega) / \hat{\zeta}(\omega_1))$ implied by the model using the alternative approach in Appendix C.3 and in the data is 1 and the quantitative results we obtain from these two approaches are very similar.

C.2 Alternative approach 1: Regression based

Instead of using the expressions in equations (14) and (16), we can measure $\hat{q}(\omega)^\theta / \hat{q}(\omega_1)^\theta$ and $\hat{q}(\kappa)^\theta / \hat{q}(\kappa_1)^\theta$ using the coefficients of a regression of the observed changes in the log of factor allocations on labor group, occupation, and equipment type fixed effects. Specifically, we can express equation (12) as

$$\hat{\pi}(\lambda, \kappa, \omega) = \hat{q}(\lambda) \hat{q}(\omega)^\theta \hat{q}(\kappa)^\theta$$

where we define

$$\hat{q}(\lambda) \equiv \left[\sum_{\kappa', \omega'} \hat{q}(\omega')^\theta \hat{q}(\kappa')^\theta \pi_{t_0}(\lambda, \kappa', \omega') \right]^{-1}$$

Hence, in the presence of multiplicative measurement error $\iota_t(\lambda, \kappa, \omega)$ in the observed changes in allocations, we have

$$\log \hat{\pi}(\lambda, \kappa, \omega) = \log \hat{q}(\lambda) + \log \hat{q}(\omega)^\theta + \log \hat{q}(\kappa)^\theta + \iota_t(\lambda, \kappa, \omega).$$

Using this equation, we regress observed values of $\log \hat{\pi}(\lambda, \kappa, \omega)$ on labor group, equipment, and occupation effects. Exponentiating the resulting occupation and equipment fixed effects, we obtain estimates of $\hat{q}(\omega)^\theta / \hat{q}(\omega_1)^\theta$ and $\hat{q}(\kappa)^\theta / \hat{q}(\kappa_1)^\theta$. Using these estimates instead of those derived from equations (14) and (16), we can recover measures of occupation shifters, $\hat{a}(\omega) / \hat{a}(\omega_1)$, and labor productivity, $\hat{T}(\lambda) / \hat{T}(\lambda_1)$, as well as estimate ρ and θ , following the same steps outlined in Appendix C.1.

Our alternative and baseline approaches are identical in the absence of zeros in the allocation data. In practice, the correlation between the measures obtained using these two approaches is above 0.99 for both equipment productivity and transformed occupation prices (both to the power θ). We use the procedure described in Appendix C.1 as our baseline approach simply because, in our opinion, it more clearly highlights the variation in the data that is being used to identify the changes in occupation prices and equipment productivity (to the power θ).

C.3 Alternative approach 2: Matching income shares

Our baseline approach yields estimate of $\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta}$ and $\frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta}$ that do not exactly match observed changes in total labor income by occupation, $\zeta_t(\omega) \equiv \sum_{\lambda, \kappa} w_t(\lambda) L_t(\lambda) \pi_t(\lambda, \kappa, \omega)$, and by equipment type, $\zeta_t(\kappa) \equiv \sum_{\lambda, \omega} w_t(\lambda) L_t(\lambda) \pi_t(\lambda, \kappa, \omega)$. In this alternative approach we calibrate $\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta}$ and $\frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta}$ to match $\frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega_1)}$ and $\frac{\hat{\zeta}(\kappa)}{\hat{\zeta}(\kappa_1)}$ exactly.

For each time period we solve simultaneously for $\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta}$ and $\frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta}$ to match observed values of $\frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega_1)}$ and $\frac{\hat{\zeta}(\kappa)}{\hat{\zeta}(\kappa_1)}$. Specifically, for every t_0 , $\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta}$ and $\frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta}$ is the solution to the following non-linear system of equations:

$$\frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega_1)} = \frac{\zeta_{t_0}(\omega_1) \sum_{\lambda, \kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega)}{\zeta_{t_0}(\omega) \sum_{\lambda, \kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega_1) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega_1)}$$

$$\frac{\hat{\zeta}(\kappa)}{\hat{\zeta}(\kappa_1)} = \frac{\zeta_{t_0}(\kappa_1) \sum_{\lambda, \omega} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega)}{\zeta_{t_0}(\kappa) \sum_{\lambda, \omega} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa_1, \omega) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa_1, \omega)}$$

where $\hat{w}(\lambda) \hat{L}(\lambda)$, $\frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega_1)}$ and $\frac{\hat{\zeta}(\kappa)}{\hat{\zeta}(\kappa_1)}$ are observed in the data, and $\hat{\pi}(\lambda, \kappa, \omega) = \frac{(\hat{q}(\omega)\hat{q}(\kappa))^\theta}{\sum_{\kappa', \omega'} (\hat{q}(\omega')\hat{q}(\kappa'))^\theta \pi_{t_0}(\lambda, \kappa', \omega')}$ is constructed given $\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta}$ and $\frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta}$. After solving for $\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta}$ and $\frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta}$, the remaining shocks and parameters are determined exactly as in our baseline procedure. We also consider a variation in which we first measure $\frac{\hat{q}(\kappa)^\theta}{\hat{q}(\kappa_1)^\theta}$ using our baseline procedure, and then $\frac{\hat{q}(\omega)^\theta}{\hat{q}(\omega_1)^\theta}$ in order to match $\frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega_1)}$ in the data. Results using these alternative approaches are very similar to our baseline results.

D Estimation of elasticities

D.1 Baseline estimation of θ and ρ

This section provides a detailed description of the estimation approach described in Section 4.3. We estimate θ and ρ jointly using a Method of Moments (MM) estimator with two moment conditions.

In order to derive the first moment condition, we express equation (17) as

$$\log \hat{w}(\lambda, t) = \zeta_\theta(t) + \beta_\theta \log \hat{s}(\lambda, t) + \iota_\theta(\lambda, t), \quad (39)$$

We observe $\log \hat{w}(\lambda, t)$ in our data and construct $\log \hat{s}(\lambda, t)$ as indicated in equation (18). The parameter $\zeta_\theta(t) \equiv \log \hat{q}(\omega_1, t) \hat{q}(\kappa_1, t)$ is a time effect that is common across λ , $\beta_\theta \equiv 1/\theta$, and $\iota_\theta(\lambda, t) \equiv \log \hat{T}(\lambda, t)$ captures unobserved changes in labor group λ productivity. As shown in equation (17), measuring changes in labor productivity requires a value of θ and, therefore, we treat $\iota_\theta(\lambda, t)$ as unobserved when estimating θ . Our model predicts the observed covariate in equation (39), $\log \hat{s}(\lambda, t)$, to be correlated with its error term, $\iota_\theta(\lambda, t)$. From equation (18), $\log \hat{s}(\lambda, t)$ is a function of changes in transformed occupation prices, $\hat{q}(\omega, t)$, and, according to our model, these depend on changes in unobserved labor productivity, $\iota_\theta(\lambda, t)$. Specifically, our model implies that the error term $\iota_\theta(\lambda, t)$ and the covariate $\log \hat{s}(\lambda, t)$ are negatively correlated: the higher the growth in the productivity of a particular labor group, the lower the growth in the price of those occupations that use that type of labor more intensively. Therefore, we expect the Nonlinear Least Squares (NLS) estimate of β_θ to be biased downwards and, consequently, the estimate of θ to be biased upwards. To address the endogeneity of the covariate $\log \hat{s}(\lambda, t)$, we construct the following instrument for $\log \hat{s}(\lambda, t)$,

$$\chi_\theta(\lambda, t) \equiv \log \sum_{\kappa} \frac{\hat{q}(\kappa, t)^\theta}{\hat{q}(\kappa_1, t)^\theta} \sum_{\omega} \pi_{1984}(\lambda, \kappa, \omega),$$

which is a labor-group-specific average of the observed changes in equipment productivity to the power θ , $\hat{q}(\kappa, t)^\theta / \hat{q}(\kappa_1, t)^\theta$.³⁴ We use this instrument and equation (39) to build the following moment condition

$$\mathbb{E}_{\lambda, t} \left[\left(y_\theta(\lambda, t) - \frac{1}{\theta} x_\theta(\lambda, t) \right) \times z_\theta(\lambda, t) \right] = 0, \quad (40)$$

where: (1) $y_\theta(\lambda, t)$ is the (λ, t) OLS residual of a regression that projects the set of dependent variables $\log \hat{w}(\lambda, t)$, for all λ and t , on a set of year fixed effects; (2) $x_\theta(\lambda, t)$ is the (λ, t) OLS residual of a regression that projects the set of independent variables $\log \hat{s}(\lambda, t)$, for all λ and t , on a set of year fixed effects; (3) $z_\theta(\lambda, t)$ is the (λ, t) OLS residual of a regression that projects the set of independent variables $\chi_\theta(\lambda, t)$, for all λ and t , on a set of year fixed effects.³⁵ In order for the moment condition in equation 40 to correctly identify the parameter θ , after controlling for year fixed effects the variable $\chi_\theta(\lambda, t)$ must be correlated with $\log \hat{s}(\lambda, t)$ and uncorrelated with $\iota_\theta(\lambda, t)$. Our model predicts that the conditioning variable $\chi_\theta(\lambda, t)$ will be correlated with the endogenous covariate $\log \hat{s}(\lambda, t)$, as an increase in the relative productivity of equipment κ between t_0 and t_1 raises the wage of group λ relatively more if a larger share of λ workers use equipment κ in period t_0 . Equation (40) implicitly imposes that the shock $\chi_\theta(\lambda, t)$ is mean independent of the labor productivity shock $\log \hat{T}(\lambda, t)$ across labor groups and time periods. A sufficient condition for this mean independence condition to hold is that, between any two periods t_0 and t_1 , the change in unobserved labor productivity, $\log \hat{T}(\lambda, t)$, and the weighted changes in equipment productivity are uncorrelated across labor groups.

In order to derive the second moment condition, we express equation (13) as

$$\log \hat{\zeta}(\omega, t) = \zeta_\rho(t) + \beta_\rho \log \frac{\hat{q}(\omega, t)^\theta}{\hat{q}(\omega_1, t)^\theta} + \iota_\rho(\omega, t). \quad (41)$$

We observe $\log \hat{\zeta}(\omega, t)$ in the MORG CPS and measure $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$ follow-

³⁴In constructing the instrument for $\log \hat{s}(\lambda, t)$ between any two periods t_0 and t_1 , we could also have used the observed labor allocations at period t_0 . However, in order to minimize the correlation between a possibly serially correlated $\iota_\theta(\lambda, t)$ and the instrument, we construct our instrument for $\log \hat{s}(\lambda, t)$ between any two sample periods t_0 and t_1 using allocations in the initial sample year, 1984.

³⁵By projecting first on a set of year effects and using the residuals from this projection in the moment condition in equation (40), we simplify significantly the computational burden involved in estimating both the parameter of interest θ and the set of incidental parameters $\{\zeta_\theta(t)\}_t$. The Frisch-Waugh-Lovell Theorem guarantees that the resulting estimate of θ is consistent and has identical asymptotic variance to the alternative GMM estimator that estimates the year fixed effects $\{\zeta_\theta(t)\}_t$ and the parameter θ jointly.

ing the procedure indicated in Section 4.2. The parameter $\varsigma_\rho(t)$ is a time effect that is common across ω and given by $\varsigma_\rho(t) \equiv \log \hat{E} + (1 - \alpha)(1 - \rho) \log \hat{q}(\omega_1, t)$, $\beta_\rho \equiv (1 - \alpha)(1 - \rho) / \theta$, and $\iota_\rho(\omega, t) \equiv \log \hat{a}(\omega, t)$ captures unobserved changes in occupation shifters. As shown in equation (15), measuring changes in occupation shifters requires a value of ρ and, therefore, we treat $\iota_\rho(\omega, t)$ as unobserved when estimating ρ . Our model predicts that the observed covariate in equation (41), $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$, and the error term $\iota_\rho(\omega, t)$ are correlated: according to our model, changes in equilibrium transformed occupation prices, $\hat{q}(\omega, t)$, depend on changes in unobserved occupation shifters, $\iota_\rho(\omega, t)$. Specifically, we expect the error term $\iota_\rho(\omega, t)$ and the covariate $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$ to be positively correlated: the higher the growth in the shifter of a particular occupation, the higher the growth in the occupation price. Therefore, given any value of α and θ , we expect the NLS estimate of β_ρ to be biased upwards and the resulting estimate of ρ to be biased downwards. To address the endogeneity of the covariate $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$, we construct the following Bartik-style instrument for $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$,

$$\chi_\rho(\omega, t) \equiv \log \sum_\kappa \frac{\hat{q}(\kappa, t)^\theta}{\hat{q}(\kappa_1, t)^\theta} \sum_\lambda \frac{L_{1984}(\lambda) \pi_{1984}(\lambda, \kappa, \omega)}{\sum_{\lambda', \kappa'} L_{1984}(\lambda') \pi_{1984}(\lambda', \kappa', \omega)},$$

which is an occupation-specific average of observed changes in equipment productivity to the power θ , $\hat{q}(\kappa, t)^\theta / \hat{q}(\kappa_1, t)^\theta$.³⁶ We use this instrument and equation (41) to build the following moment condition

$$\mathbb{E}_{\omega, t} \left[\left(y_\rho(\omega, t) - (1 - \alpha)(1 - \rho) \frac{1}{\theta} x_\rho(\omega, t) \right) \times z_\rho(\omega, t) \right] = 0, \quad (42)$$

where: (1) $y_\rho(\omega, t)$ is the (ω, t) OLS residual of a regression that projects the set of dependent variables $\log \hat{\zeta}(\omega, t)$, for all ω and t , on a set of year fixed effects; (5) $x_\rho(\omega, t)$ is the (ω, t) OLS residual of a regression that projects the set of independent variables $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$, for all ω and t , on a set of year fixed effects; and (6) $z_\rho(\omega, t)$ is the (ω, t) OLS residual of a regression that projects the set of instruments $\chi_\rho(\omega, t)$, for all ω and t , on a set of year fixed effects. In order for the moment condition in equation 42 to correctly identify the parameter ρ , after controlling for year fixed effects the variable $\chi_\rho(\lambda, t)$ must be correlated with $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$ and uncorrelated with $\iota_\rho(\lambda, t)$. Our model predicts that the conditioning variable $\chi_\rho(\omega, t)$ will be correlated with the endogenous covariate $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$ as an increase in the relative productivity of κ

³⁶In order to minimize the correlation between a possibly serially correlated $\iota_\rho(\omega, t)$ and the instrument, we construct $\chi_\rho(\omega, t)$ using allocations in 1984: $L_{1984}(\lambda) \pi_{1984}(\lambda, \kappa, \omega)$ is the number of λ workers using equipment κ employed in occupation ω in 1984 and the denominator in the expression for $\chi_\rho(\omega, t)$ is total employment in occupation ω in 1984.

raises occupation ω 's output—and, therefore, reduces its price—relatively more if a larger share of workers employed in occupation ω use equipment κ in period t_0 . Equation (42) implicitly imposes that the shock $\chi_\rho(\lambda, t)$ is mean independent across occupations and time periods of the occupation shifter $\log \hat{a}(\omega, t)$. A sufficient condition for this mean independence condition to hold is that, for any given pair of years t_0 and t_1 , the change in the (unobserved) occupation shifter and the weighted changes in equipment productivity are uncorrelated across occupations.

We estimate θ and ρ using the sample analogue of the moment conditions in equations (40) and (42). These two moment conditions exactly identify the parameter vector (θ, ρ) . In order to build these sample analogues, we use data on four time periods: 1984-1989, 1989-1993, 1993-1997, and 1997-2003. As discussed in Section 4.3, we obtain a point estimate of θ equal to 1.81 (standard error equal to 0.28) and a point estimate of ρ equal to 1.81 (standard error equal to 0.36).

Interestingly, the difference between our baseline estimates and analogous non-linear least squares NLS estimates—i.e., non-linear estimation without instruments—is consistent with the bias that, according to our model, should affect these NLS estimates. The fact that the NLS estimate of θ is higher than its baseline counterpart is consistent with the prediction of our model that the error term $\iota_\theta(\lambda, t)$ is negatively correlated with the covariate $\log \hat{s}(\lambda, t)$. The fact that the NLS estimate of ρ is lower than its baseline counterpart is consistent with the prediction of our model that the error term $\iota_\rho(\omega, t)$ is negatively correlated with the covariate $\log \hat{q}(\omega, t)^\theta / \hat{q}(\omega_1, t)^\theta$.

D.2 Estimation of θ and ρ allowing for time trends

Here we discuss estimates of θ and ρ that result when we add as controls a labor-group-specific time trend in equation (39) and an occupation-specific time trend in equation (41). Allowing for these time trends relaxes the orthogonality restrictions imposed to derive the moment conditions used in our baseline analysis (equations (40) and (42)).

In reviewing Krusell et al. (2000), Acemoglu (2002) raises the concern that the presence of common trends in unobserved labor-group-specific productivity, explanatory variables, and instruments may bias the estimates of wage elasticities. In order to address this concern, we follow Katz and Murphy (1992), Acemoglu (2002), and the estimation of the canonical model more generally. Specifically, we express $\hat{T}(\lambda, t)$ as following a λ -specific time trend with deviations around this trend, $\log \hat{T}(\lambda, t) = \beta_\theta(\lambda) \times (t_1 - t_0) + \iota_{\theta 1}(\lambda, t)$, and build a moment condition that is identical to that in equation (40) except for the fact that each of the variables $y_\theta(\lambda, t)$, $x_\theta(\lambda, t)$, and $z_\theta(\lambda, t)$ is now defined as the (λ, t)

OLS residual of a regression that projects $\log \hat{w}(\lambda, t)$, $\log \hat{s}(\lambda, t)$, and $\chi_\theta(\lambda, t)$, respectively, on a set of year fixed effects and on a set of labor-group-specific time trends. The resulting moment condition therefore assumes that, after controlling for year fixed effects, deviations from a labor-group-specific linear time trend in the labor-group-specific productivities $\log \hat{T}(\lambda, t)$ are mean independent of the deviations from a labor-group-specific linear time trend in the labor-group specific average of equipment shocks $\chi_\theta(\lambda, t)$. This orthogonality restriction is weaker than that imposed in our baseline estimation to derive the moment condition in equation (40). Specifically, explicitly controlling for labor-group specific time trends in the wage equation guarantees that the resulting estimates of θ will be consistent even if it were to be true that those labor groups whose productivity, $\log \hat{T}(\lambda, t)$, has grown more during the 20 years between 1984 and 2003 also happen to be the labor groups that in 1984 were more intensively using those types of equipment whose productivities, $\hat{q}(\kappa, t)$, have also grown systematically more during the 1984-2003 time period. As an example, if it were to be true that (i) highly educated workers used computers more in 1984, (ii) they experienced a large average growth in their productivities between 1984 and 2003, and (iii) computers also had a relatively large growth in their productivity in the same sample period, then our baseline estimates of θ would be biased but the estimates that control for labor-group-specific time trends would not, unless it were true that those specific years within the period 1984-2003 with higher growth of the productivity of educated workers were precisely also the years in which the productivity of computers also happened to grow above its 1984-2003 trend.

In the same way in which we allow for a labor-group specific time trend in equation (39), we also additionally control for an occupation-specific time trend in equation (41). Specifically, we express the unobserved changes in occupation shifters, $\log \hat{a}(\omega, t)$, as the sum of a ω -specific time trend and deviations around this trend, $\log \hat{a}(\omega, t) = \beta_\rho(\omega) \times (t_1 - t_0) + \iota_{\rho 1}(\omega, t)$. We then build a moment condition that is analogous to that in equation (42) except for the fact that each of the variables $y_\rho(\omega, t)$, $x_\rho(\omega, t)$, and $z_\rho(\omega, t)$ are now defined as the residuals of a linear projection of each them on year fixed effects and an occupation-specific time trend. The orthogonality restriction implied by the resulting moment condition is weaker than that in our baseline estimation; specifically, it would not be violated in the hypothetical case in which those occupations whose idiosyncratic productivity grew systematically more during the period 1984-2003 happen to also be the occupations that in 1984 used more intensively the types of equipment whose idiosyncratic productivity also grew more on average during this same period.

The estimates of θ and ρ that result from adding as controls a labor-group-specific time trend in equation (39) and an occupation-specific time trend in equation (41) are,

respectively 1.26, with a standard error of 0.31, and 2.10, with a standard error of 0.75, as discussed in Section 4.3. The fact that the estimate of θ is smaller than that obtained without controlling for labor-group specific time trends and the estimate of ρ is larger than that obtained without controlling for occupation-group specific time trends is consistent with the hypothesis that our baseline estimates are affected by a weaker version of the same kind of bias affecting the NLS estimates. However, note that allowing for time trends does not have a large quantitative impact in our estimates of θ and ρ : the two estimates of θ are within two standard deviations of each other and the two estimates of ρ are even within one standard deviation of each other. Furthermore, when computing the effects of shocks on relative wages, the estimates that result from controlling for worker-group-specific time trends in equation (39) actually imply a smaller role for changes in labor productivity (the residual), as we show in Appendix E.1.

D.3 Estimation of θ using within-worker group distribution of wages

Here we consider an approach—based on Lagakos and Waugh (2013) and Hsieh et al. (2016)—to estimate θ under a very different set of restrictions than those imposed above. This approach identifies θ from moments of the unconditional distribution of observed wages within each labor group λ and, therefore, trivially allows one to estimate values of θ that vary by labor group.

Recall that in our baseline approach, a worker $z \in \mathcal{Z}_t(\lambda)$ supplies $\epsilon(z) \times \epsilon(z, \kappa, \omega)$ efficiency units of labor if teamed with equipment κ in occupation ω . Hence, in spite of the fact that each worker $z \in \mathcal{Z}_t(\lambda)$ draws $\epsilon(z, \kappa, \omega)$ across (κ, ω) pairs from a Fréchet distribution with CDF $G(\epsilon) = \exp(\epsilon^{-\theta})$, the introduction of $\epsilon(z)$ allows for correlation in efficiency units across (κ, ω) pairs within a given worker in an unrestricted way.

A more typical and restrictive approach to allow for correlation—see e.g. Ramondo and Rodríguez-Clare (2013) and Hsieh et al. (2016)—assumes away $\epsilon(z)$ (i.e. assumes its distribution across z is degenerate) and instead uses a more parametric assumption: each worker $z \in \mathcal{Z}_t(\lambda)$, draws the vector $\{\epsilon(z, \kappa, \omega)\}_{\kappa, \omega}$ from a multivariate Fréchet distribution,

$$G(\epsilon(z); \lambda) = \exp \left(- \left(\sum_{\kappa, \omega} \epsilon(z, \kappa, \omega)^{-\tilde{\theta}(\lambda)/1-\nu(\lambda)} \right)^{1-\nu(\lambda)} \right).$$

The parameter $\tilde{\theta}(\lambda) > 1$ governs the λ -specific dispersion of efficiency units across (κ, ω) pairs; a higher value of $\tilde{\theta}(\lambda)$ decreases this dispersion. The parameter $0 \leq \nu(\lambda) \leq$

1 governs the λ -specific correlation of each worker's efficiency units across (κ, ω) pairs; a higher value of $\nu(\lambda)$ increases this correlation. We define $\theta(\lambda) \equiv \tilde{\theta}(\lambda) / (1 - \nu(\lambda))$. In what follows, we use this generalized distribution.

Allowing θ to vary by λ , our assumption on the distribution of idiosyncratic productivity implies that the distribution of wages within labor group λ is Fréchet with shape parameter $\tilde{\theta}(\lambda)$, where $\theta(\lambda) \equiv \tilde{\theta}(\lambda) / (1 - \nu(\lambda))$. We can therefore use the empirical distribution of wages within each λ to estimate $\tilde{\theta}(\lambda)$, separately for each labor group λ . Specifically, we jointly estimate the shape and scale parameter for each λ in each year t using maximum likelihood. Figure 3 plots the empirical and predicted wage distributions

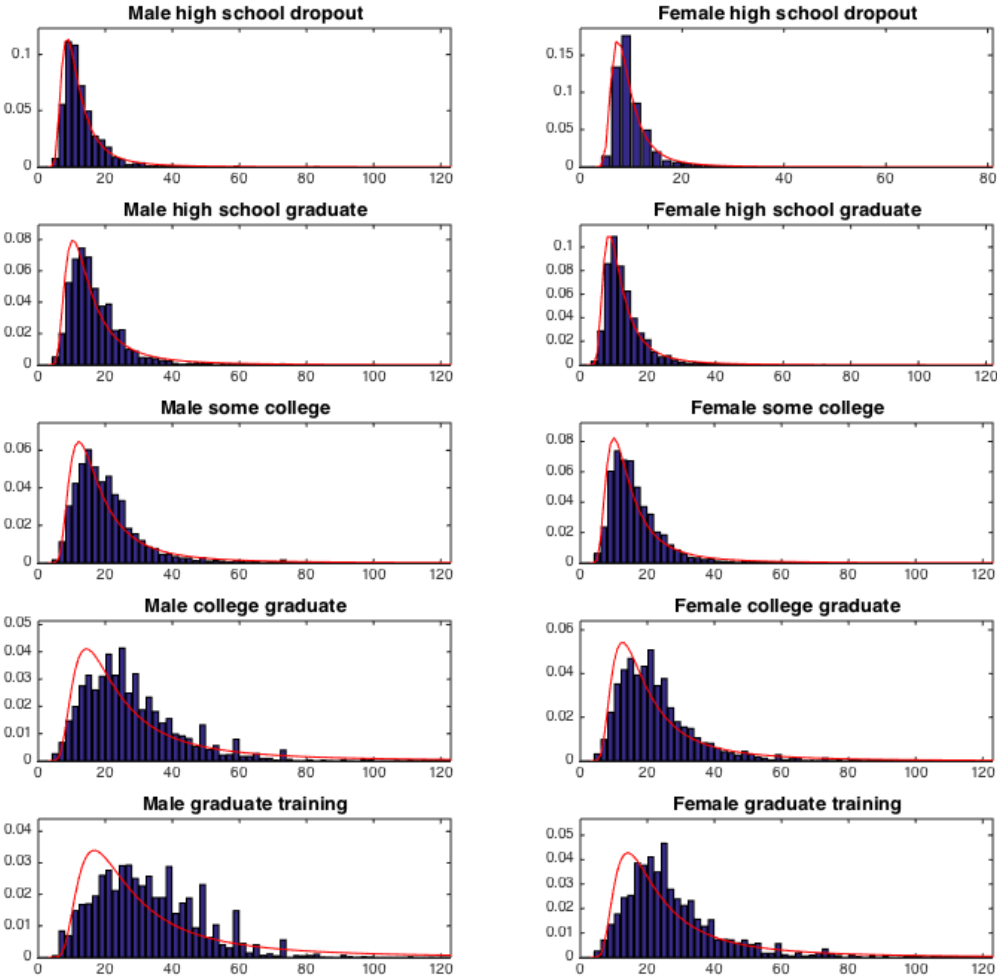


Figure 3: Empirical and predicted (Fréchet distribution estimated using maximum likelihood) wage distributions for all middle-aged labor groups in 2003

for all middle-aged workers in 2003. We then average across years our estimates of the shape parameter to obtain $\tilde{\theta}(\lambda)$. Finally, we obtain an estimate of $\theta(\lambda)$ from $\tilde{\theta}(\lambda)$ using Hsieh et al.’s (2016) implied estimate of $\nu \equiv \nu(\lambda) \approx 0.1$.

Consistent with the observation that higher earning labor groups have more within-group wage dispersion, see e.g. Lemieux (2006), we find that $\theta(\lambda)$ is lower (i.e. more within group wage inequality) for more educated groups than less educated groups—averaging within each of the five education groups across age and gender, we obtain estimates that fall monotonically from 3.46 amongst high school dropouts to 2.03 amongst those with graduate training—for men than women—averaging within each gender across age and education, we obtain an estimate of 2.41 for men and 2.82 for women. The average across λ varies very little over time.

E Robustness and sensitivity analyses

In this section we consider four types of sensitivity exercises. First, we perform sensitivity to different values of ρ and θ . Second, we illustrate the importance of accounting for all three forms of comparative advantage by performing similar exercises to those described in Section 4.4 in versions of our model that omit some of them. Third, we allow for changes in comparative advantage over time (i.e. relaxing assumption (10)). Finally, we perform sensitivity to allowing for θ to vary across λ , as described in Appendix D.3.

E.1 Alternative parameter values

We first consider the sensitivity of our results for the skill premium and gender gap over the period 1984-2003 to alternative estimated values of the parameters θ and ρ . To demonstrate the role of each parameter, we then show the impact on our results of varying one parameter at a time.

Alternative estimated values of θ and ρ . In Table 11, we decompose changes in the skill premium and gender gap between 1984 and 2003 using our alternative estimates of θ and ρ . The first row uses our benchmark GMM estimates. The second row uses our GMM estimates of θ and ρ when adding as controls a labor-group-specific time trend in equation (39) and an occupation-specific time trend in equation (41). The third row uses the value of θ estimated from moments of the unconditional distribution of observed wages within each labor group λ and our baseline value of ρ .

Our main results are robust for all these alternative estimates of the parameter vector (θ, ρ) . First, computerization is the most important force accounting for the rise in

(θ, ρ)	Skill premium				Gender gap			
	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Baseline	-0.113	0.049	0.159	0.055	0.041	-0.069	-0.046	-0.058
Time trends	-0.117	-0.012	0.251	0.028	0.043	-0.058	-0.084	-0.035
Wage distribution	-0.093	0.057	0.109	0.077	0.034	-0.060	-0.030	-0.077

Table 11: Decomposing changes in the log skill premium and gender gap between 1984 and 2003 for alternative estimates of (θ, ρ) . Estimates of (θ, ρ) are: baseline (1.81, 1.81), time trends (1.26, 2.10), and wage distribution (2.62, 1.81).

between-education-group inequality between 1984 and 2003. Alone, it accounts for between roughly 45% and 94% of the demand-side forces raising the skill premium. Second, residual labor productivity accounts for no more than one-third of the demand-side forces raising the skill premium. Third, changes in each of the demand-side forces play an important role in accounting for the reduction in the gender gap between 1984 and 2003. Specifically, changes in labor productivity play a larger role in accounting for the reduction in the gender gap than in accounting for the rise in the skill premium. Alone, they account for between roughly 20% and 46% of the demand-side forces reducing the gender gap.

The role of θ and ρ in shaping our decomposition. To provide intuition for the roles of θ and ρ , we recompute our counterfactuals varying either θ or ρ to take the values that correspond to the endpoints of the 95% confidence interval that is implied by our baseline estimation.³⁷ Whereas the first row of Table 12 replicates our baseline decomposition, the second and third rows fix ρ at our baseline level and vary θ , whereas the fourth and fifth rows fix θ at our baseline level and vary ρ .

(θ, ρ)	Skill premium				Gender gap			
	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Baseline	-0.113	0.049	0.159	0.055	0.041	-0.069	-0.046	-0.058
(2.36, 1.81)	-0.099	0.056	0.121	0.071	0.036	-0.063	-0.034	-0.072
(1.26, 1.81)	-0.132	0.024	0.229	0.028	0.048	-0.075	-0.072	-0.035
(1.81, 2.52)	-0.090	-0.005	0.186	0.058	0.033	-0.044	-0.063	-0.060
(1.81, 1.10)	-0.153	0.140	0.111	0.050	0.056	-0.112	-0.019	-0.056

Table 12: Decomposing changes in the log skill premium and gender gap between 1984 and 2003 for extreme values of θ or ρ . Baseline estimates of (θ, ρ) are (1.81, 1.81).

³⁷This is intended to provide intuition for how our results vary with alternative values of θ and ρ . The numbers in Table 12 should not be interpreted as confidence intervals for our decomposition.

As is evident from rows two and three of Table 12, a higher value of θ implies a larger role for changes in labor productivity and a smaller role for the other shocks in accounting for changes in the skill premium and the gender gap. The intuition is straightforward. According to equation (17), the elasticity of changes in average wages of workers in labor group λ , $\hat{w}(\lambda)$, to changes in the measured labor-group-specific average of equipment productivities and transformed occupation prices (both to the power θ), $\hat{s}(\lambda)$, is $1/\theta$. Because our measure of $\hat{s}(\lambda)$ is independent of θ , a higher value of θ reduces the impact on wages of changes in the labor-group-specific average of equipment productivities and transformed occupation prices and, therefore, increases the impact of changes in labor productivity, identified as a residual to match observed changes in average wages.

The value of ρ may potentially affect the contribution of each shock to relative wages through two channels: by affecting the measured shock itself and by affecting the elasticity of occupation prices to these measured shocks. As shown in Section 4.2, ρ does not affect our measurement of either the labor composition or equipment productivity shock; hence, ρ affects the importance of these shocks for relative wages only through the elasticity of occupation prices. Because labor composition only affects relative wages through occupation prices, a higher value of ρ reduces the impact of labor composition on relative wages. As described in Section 3.4, computerization has two effects. First, it raises the relative wages of labor groups that disproportionately use computers. Second, by lowering the prices of occupations in which computers are disproportionately used, it lowers the wages of labor groups that are disproportionately employed in these occupations. A higher value of ρ mitigates the second effect and, therefore, strengthens the impact of computerization on the skill premium and gender gap.

On the other hand, the value of ρ impacts occupation shifters both through the magnitude of the measured shocks (see equation (15)) and through the elasticity of occupation prices to these measured shocks. In practice, a higher value of ρ yields measured occupation shifters that are less biased towards educated workers (in fact, occupation shifters reduce the skill premium for sufficiently high values of ρ) and tends to reduce the effect of occupation shifters on the gender gap by reducing the elasticity of occupation prices to shocks.

E.2 Sources of comparative advantage

To demonstrate the importance of including each of the three forms of comparative advantage, we perform two exercises. We first assume there is no comparative advantage related to occupations and then we redo the decomposition under the assumption that

there is no comparative advantage related to equipment. In all cases, we hold the values of α , ρ , and θ fixed to the same values employed in Section 4.4.

Table 13 reports our baseline decomposition between 1984 and 2003 both for the skill premium (in the left panel) and the gender gap (in the right panel) as well as decompositions under the restriction that there is comparative advantage only between labor and equipment or only between labor and occupations.

	Skill premium				Gender gap			
	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Baseline	-0.113	0.049	0.159	0.055	0.041	-0.069	-0.046	-0.058
Only labor-equip. CA	0	0	0.237	-0.085	0	0	-0.103	-0.030
Only labor-occ. CA	-0.113	0.120	0	0.143	0.041	-0.059	0	-0.116

Table 13: Decomposing changes in the log skill premium and log gender gap between 1984 and 2003 under different assumptions on the evolution of comparative advantage

Abstracting from any comparative advantage at the level of occupations (i.e. assuming away worker-occupation and equipment-occupation comparative advantage) has two effects. First—because changes in labor composition and occupation shifters affect relative wages only through occupation prices—it implies that the labor composition and occupation shifters components of our decomposition go to zero. This affects the labor productivity component, since changes in labor productivity are identified as a residual to match observed changes in wages. Second, it implies that worker-equipment comparative advantage is the only force giving rise to the observed allocation of labor groups to equipment types. This affects the inferred strength of worker-equipment comparative advantage and, therefore, affects both the equipment and labor productivity components of the decomposition.

Row 2 of Table 13 shows that if we were to abstract from any comparative advantage at the level of occupations, we would incorrectly conclude that all of the rise in the skill premium has been driven by changes in relative equipment productivities. Similarly, because we would infer that women have a strong comparative advantage with computers, we would incorrectly conclude that changes in equipment productivity account for almost all of the fall in the gender gap.

Similarly, assuming there is no comparative advantage at the level of equipment implies that the equipment productivity component of our decomposition is zero and that the only force giving rise to the allocation of labor groups to occupations is worker-occupation comparative advantage. Row 3 of Table 13 shows that abstracting from any comparative advantage at the level of equipment magnifies the importance of labor pro-

ductivity in explaining the rise of the skill premium and the fall in the gender gap. The impact of occupation shifters on the gender gap does not change significantly.

In summary, abstracting from comparative advantage at the level of either occupations or equipment has a large impact on the decomposition of changes in between-group inequality. It does so by forcing changes in labor productivity to absorb the impact of the missing component(s) and by changing the importance of the remaining source of comparative advantage.

E.3 Evolving comparative advantage

In our baseline model we imposed that the only time-varying components of productivity are multiplicatively separable between labor, equipment, and occupation components. In practice, over time some labor groups may have become relatively more productive in some occupations or using some types of equipment, perhaps caused by differential changes in discrimination of labor groups across occupations, by changes in occupation characteristics that affect labor groups differentially, or by changes in the characteristics of equipment.

In the most general case, we could allow $T_t(\lambda, \kappa, \omega)$ to vary freely over time. In this case, we would match $\hat{\pi}(\lambda, \kappa, \omega)$ exactly in each time period. The impact of labor composition would be exactly the same as in our baseline. However, we would only be able to report the joint effects of the combination of all λ -, κ -, and ω -specific shocks on relative wages. Instead, here we generalize our baseline model to incorporate changes over time in comparative advantage in a restricted manner. Specifically, we consider separately three extensions of our baseline model:

$$T_t(\lambda, \kappa, \omega) = \begin{cases} T_t(\kappa) T_t(\lambda, \omega) T(\lambda, \kappa, \omega) & \text{case 1} \\ T_t(\omega) T_t(\lambda, \kappa) T(\lambda, \kappa, \omega) & \text{case 2} \\ T_t(\lambda) T_t(\kappa, \omega) T(\lambda, \kappa, \omega) & \text{case 3} \end{cases}$$

We allow for changes over time in comparative advantage between workers and occupations in case 1, workers and equipment in case 2, and equipment and occupations in case 3. Table 14 reports our results from decomposing changes in the skill premium between 1984 and 2003 in our baseline exercise as well as in cases 1, 2, and 3. In all cases, we hold the values of α , ρ , and θ fixed to the same values employed in Section 4.4.

Our results are largely unchanged and the intuition for why is straightforward in cases 1 and 2. In all three cases, our measures of initial factor allocations and changes in labor

Changes in CA	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Labor-occ.	Labor-equip.	Equip.-occ.
None (baseline)	-0.113	0.049	0.159	0.055	-	-	-
Worker-occ. (case 1)	-0.113	-	0.159	-	0.094	-	-
Worker-equip. (case 2)	-0.113	0.046	-	-	-	0.221	-
Equip.-occ. (case 3)	-0.113	-	-	0.044	-	-	0.220

Table 14: Decomposing changes in the log skill premium between 1984 and 2003 allowing comparative advantage to evolve over time

composition as well as the system of equations that determines the impact of changes in labor composition on relative wages are exactly the same as in our baseline model. Hence, the labor composition component of our baseline decomposition is unchanged if we incorporate time-varying comparative advantage. Similarly, our measure of changes in equipment productivity as well as the system of equations that determines their impact are exactly the same in case 1 as in our baseline model. Hence, the equipment productivity component in case 1 is unchanged from the baseline. In case 2, whereas our measure of changes in transformed occupation prices is exactly the same as in our baseline model, our measure of changes in occupation labor payment shares—and, therefore, our measure of occupation shifters—differs slightly from our baseline, since predicted allocations in period t_1 differ slightly. However, since these differences aren't large and since the system of equations determining the impact of occupation shifters is the same, our results on occupation shifters in case 2 are very similar to those in the baseline. Finally, since (when fed in one at a time) the sum of all four components of our decomposition in the baseline model match the change in relative wages in the data well and the sum of all three components of our decomposition in the extensions considered here match the data reasonably well (in each case they match wage changes perfectly when fed in together), the change in wages resulting from the sum of the labor productivity and occupation productivity components in our baseline (when fed in one at a time) must closely match the change in wages from the labor-occupation component in case 1; similarly, the sum of the labor productivity and equipment productivity components in our baseline must closely match the labor-equipment component in case 2.

In what follows we show how to measure the relevant shocks and how to decompose changes in between-group inequality into labor composition, occupation shifter, and labor-equipment components in case 2. Details for cases 1 and 3 are similar.

Details for case 2. In case 2 we impose the following restriction

$$T_t(\lambda, \kappa, \omega) \equiv T_t(\omega) T_t(\lambda, \kappa) T(\lambda, \kappa, \omega). \quad (43)$$

The equilibrium conditions are unchanged: equations (3), (4), and (5) hold as in our baseline model. However, we can re-express the system in changes as follows. Defining

$$q_t(\lambda, \kappa) \equiv T_t(\lambda, \kappa) p_t(\kappa)^{\frac{-\alpha}{1-\alpha}},$$

equations (11) and (12) become

$$\hat{w}(\lambda) = \left(\sum_{\kappa, \omega} \pi_{t_0}(\lambda, \kappa, \omega) (\hat{q}(\lambda, \kappa) \hat{q}(\omega))^{\theta(\lambda)} \right)^{1/\theta(\lambda)}, \quad (44)$$

$$\hat{\pi}(\lambda, \kappa, \omega) = \frac{(\hat{q}(\omega) \hat{q}(\lambda, \kappa))^{\theta(\lambda)}}{\sum_{\kappa', \omega'} (\hat{q}(\omega') \hat{q}(\lambda, \kappa'))^{\theta(\lambda)} \pi_{t_0}(\lambda, \kappa', \omega')}, \quad (45)$$

whereas equation (13) remains unchanged. Expressing equation (44) in relative terms yields

$$\frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} = \frac{\hat{q}(\lambda, \kappa_1)}{\hat{q}(\lambda_1, \kappa_1)} \frac{\left(\sum_{\kappa, \omega} \pi_{t_0}(\lambda, \kappa, \omega) \left(\frac{\hat{q}(\lambda, \kappa)}{\hat{q}(\lambda, \kappa_1)} \frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} \right)^{\theta(\lambda)} \right)^{1/\theta(\lambda)}}{\left(\sum_{\kappa', \omega'} \pi_{t_0}(\lambda_1, \kappa', \omega') \left(\frac{\hat{q}(\lambda_1, \kappa')}{\hat{q}(\lambda_1, \kappa_1)} \frac{\hat{q}(\omega')}{\hat{q}(\omega_1)} \right)^{\theta(\lambda_1)} \right)^{1/\theta(\lambda_1)}}. \quad (46)$$

Hence, the decomposition requires that we measure $\hat{q}(\lambda, \kappa) / \hat{q}(\lambda, \kappa_1)$ for each (λ, κ) as well as $\hat{q}(\lambda, \kappa_1) / \hat{q}(\lambda_1, \kappa_1)$ for each λ .

Here we provide an overview—similar in structure to that provided in Section 4.2—of how we measure shocks taking as given the parameters α , ρ , and θ . Equations (3) and (43) give us

$$\frac{\hat{q}(\lambda, \kappa_1)^\theta}{\hat{q}(\lambda, \kappa_2)^\theta} = \frac{\hat{\pi}(\lambda, \kappa_1, \omega)}{\hat{\pi}(\lambda, \kappa_2, \omega)}$$

for each λ and ω . Hence, we can measure $\hat{q}(\lambda, \kappa_1)^\theta / \hat{q}(\lambda, \kappa_2)^\theta$ for each λ as the exponential of the average across ω of the log of the right-hand side of the previous expression. We can recover changes in transformed occupation prices to the power θ and use these to measure changes in occupation shifters exactly as in our baseline. Finally, given these measures, we can recover $\hat{q}(\lambda, \kappa_1)^\theta / \hat{q}(\lambda_1, \kappa_1)^\theta$ to match changes in relative wages using equation (46).

E.4 Letting θ vary by λ

Here we perform our decomposition allowing allowing $\theta(\lambda)$ to vary across λ using the multivariate Fréchet described above in Appendix D.3.

Our baseline equilibrium equations in levels—(3), (4), and (5)—and in changes—(11), (12), and (13)—are unchanged except for $\theta(\lambda)$ replacing θ . The key distinction between our baseline and extended model is the parameterization. Appendix D.3 describes how to use the empirical distribution of wages to estimate $\theta(\lambda)$ separately for each labor group λ . Given values of $\theta(\lambda)$, we measure changes in equipment productivity and transformed occupation prices, not to the power $\theta(\lambda)$, using the following variants of equations (14) and (16)

$$\frac{\hat{q}(\kappa)}{\hat{q}(\kappa_1)} = \left(\frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa_1, \omega)} \right)^{1/\theta(\lambda)}$$

and

$$\frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} = \left(\frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa, \omega_1)} \right)^{1/\theta(\lambda)}.$$

Given changes in transformed occupation prices, we measure changes in occupation shifters using equation (15).

Table 15 reports the results of our decomposition of the skill premium and the gender gap over the period 1984-2003 under three alternative specifications. The first row reports our baseline results in which θ is constant across all groups and estimated as described in Section 4.3. The second row reports results in which $\theta(\lambda)$ is estimated separately for each λ , but we use the average value θ for each λ , reported in Appendix D.3. Finally the final row reports results using distinct values of $\theta(\lambda)$ across each λ . In each case, we use our baseline value of $\rho = 1.81$. The key message of Table 15 is that our results are robust. This is particularly true comparing between the second and the third rows of Table 15, in which the average value of $\theta(\lambda)$ is constant by construction.

	Skill premium				Gender gap			
	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Baseline	-0.113	0.049	0.159	0.055	0.041	-0.069	-0.046	-0.058
$\theta = 2.62$	-0.093	0.057	0.109	0.077	0.034	-0.060	-0.030	-0.077
$\theta(\lambda)$	-0.098	0.047	0.111	0.090	0.035	-0.063	-0.036	-0.068

Table 15: Decomposing changes in the log skill premium and gender gap between 1984 and 2003 allowing $\theta(\lambda)$ to vary with θ

F Additional details

F.1 Computer use decomposition

We decompose the difference in aggregate computer use between women and men as well as between college-educated and non-college-educated workers in Panel B of Table 1. Let $\pi_t(\lambda, c)$ denote the aggregate share of hours group λ spends working with computers in year t , $\pi_t(\lambda, c) \equiv \sum_{\omega} \pi_t(\lambda, c, \omega)$; $\pi_t(\lambda, \omega)$ denote the share of hours group λ spends working in occupation ω in year t , $\pi_t(\lambda, \omega) \equiv \sum_{\kappa} \pi_t(\lambda, \kappa, \omega)$, and $s_t(\lambda, c, \omega)$ denote the share of hours within occupation ω that group λ spends working with computers in year t , $s_t(\lambda, c, \omega) \equiv \pi_t(\lambda, c, \omega) / \pi_t(\lambda, \omega)$. We then have

$$\pi_t(\lambda, c) = \sum_{\omega} s_t(\lambda, c, \omega) \pi_t(\lambda, \omega).$$

Hence, we can decompose the difference between the aggregate computer use of two groups λ' and λ as

$$\Delta\pi_t(\cdot, c) = \sum_{\omega} \bar{s}_t(\cdot, c, \omega) \Delta\pi_t(\cdot, \omega) + \sum_{\omega} \Delta s_t(\cdot, c, \omega) \bar{\pi}_t(\cdot, \omega),$$

where for any variable x , $\bar{x}(\cdot) = (x(\lambda') + x(\lambda)) / 2$ is the average of λ' and λ and $\Delta x(\cdot) = x(\lambda') - x(\lambda)$ is the difference between λ' and λ . We refer to the first term of the above decomposition as the between component, since it reflects the share of the difference in aggregate computer usage between λ' and λ that reflects differences in allocations across occupations at common (average) computer usage within each occupation. We refer to the second term of the above decomposition as the within component, since it reflects the share of the difference in aggregate computer usage between λ' and λ that reflects differences in computer usage within each occupation at common (average) allocations across occupations.

F.2 Occupation characteristics

As documented in Figure 1 and in Table 10 in Appendix B, certain occupations—including, for example, executive, administrative, managerial as well as health assessment and treating—have grown disproportionately over the last three decades. As discussed in, e.g., Autor et al. (2003), these changes have been systematically related to the task content of each occupation; for example, there has been an expansion of occupations intensive in non-routine cognitive analytical and non-routine cognitive interpersonal tasks and a

corresponding contraction in occupations intensive in routine manual and non-routine manual physical tasks, as defined following [Acemoglu and Autor \(2011\)](#).

Here we use the characteristics of our thirty occupations, derived from O*NET as described in [Appendix B](#), to understand how each shock shapes the observed evolution of labor income shares across occupations. The first row of [Table 16](#) shows that, if we regress the change in the share of labor income earned in each occupation between 1984 and 2003, measured using the MORG CPS, separately on each of the four occupation characteristics, we find a systematic contraction in occupations that are intensive in routine manual as well as non-routine manual physical tasks and a systematic expansion of occupations that are intensive in non-routine cognitive analytical and non-routine cognitive interpersonal tasks. This growth pattern of different occupations depending on their task content has been previously documented in a large literature. Rows two through five replicate this exercise, but instead of using the change in labor income shares across occupations from the data, we use the change predicted by our model in response to each shock separately. If $\rho = 1$, then changes in labor income shares across occupations are generated by occupation demand shifters alone. Because $\rho \neq 1$, changes in labor income shares across occupations are also caused by other shocks.

In practice, we find that changes in equipment productivity, labor composition, and labor productivity play a significant role in the systematic contraction of occupations that are intensive in routine manual as well as non-routine manual physical tasks and the systematic expansion of occupations that are intensive in non-routine cognitive analytical and non-routine cognitive interpersonal tasks. On the other hand, occupation demand shifters do not have a statistically significant effect.

F.3 Worker aggregation

In theory we could incorporate as many labor groups, equipment types, and occupations as the data permits without complicating our measurement of shocks or our estimation of parameters. In practice, as we increase the number of labor groups, equipment types, or occupations, we also increase both the share of $(\lambda, \kappa, \omega)$ triplets for which $\pi_t(\lambda, \kappa, \omega) = 0$ and measurement error in factor allocations in general.

Our objective here is to understand the extent to which our particular disaggregation may be driving our results. To do so, we decrease the number of labor groups from 30 to 10 by dropping age as a characteristic. In this case, the share of $(\lambda, \kappa, \omega)$ observations for which $\pi_t(\lambda, \kappa, \omega) = 0$ falls from (roughly) 26 percent to 12 percent. Because, in our baseline, we compositionally adjust the skill premium and the gender gap using gender,

	Non-routine cogn. anlyt.	Non-routine cogn. inter.	Routine man.	Non-routine man. phys.
Data	0.217***	0.233***	-0.285***	-0.210***
Labor composition	0.050***	0.050***	-0.044***	-0.040***
Occupation shifters	-0.065	0.050	-0.088	-0.035
Equipment prod.	0.088***	0.057***	-0.065***	-0.070***
Labor productivity	0.003**	0.004***	-0.004***	-0.005***

Table 16: The evolution of labor income shares across occupations in the data and predicted separately by each shock. Each cell represents the coefficient estimated from a separate OLS regression across thirty occupations of the change in the income share between 1984 and 2003—either in the data (using only the MORG CPS) or predicted in the model by each shock—on a constant and a single occupation characteristic derived from O*NET.

Non-routine cogn. anlyt. refers to Non-routine cognitive analytical; Non-routine cogn. inter. refers to Non-routine cognitive interpersonal; and Non-routine man. phys. refers to Non-routine manual physical

*** Significant at the 1 percent level, ** Significant at the 5 percent level, * Significant at the 10 percent level

education, and age, whereas here we only use gender and education, we find slightly different changes in the skill premium, 16.1 instead of 15.1 percent, and the gender gap, -13.2 instead of -13.3 percent, between 1984 and 2003.

	Data	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
	Skill premium				
Baseline 30 groups	0.151	-0.113	0.049	0.159	0.055
10 groups baseline θ, ρ	0.161	-0.096	0.062	0.158	0.036
10 groups re-estimate θ, ρ	0.161	-0.112	0.065	0.195	0.012
	Gender gap				
Baseline 30 groups	0.133	0.041	-0.069	-0.046	-0.058
10 groups baseline θ, ρ	0.132	-0.040	0.064	0.047	0.059
10 groups re-estimate θ, ρ	0.132	-0.047	0.075	0.059	0.044

Table 17: Decomposing changes in the log skill premium and gender gap between 1984 and 2003 with 10 rather than 30 labor groups

We conduct our decomposition with 10 labor groups using two different approaches. In both approaches we re-measure all shocks. However, in one approach we use our baseline values of θ and ρ estimated with 30 labor groups, $\theta = 1.81$ and $\rho = 1.81$, whereas in the other approach we re-estimate these parameters with 10 labor groups using our baseline estimation approach, yielding $\theta = 1.41$ and $\rho = 1.70$. The standard error on the estimate of ρ is almost twice as large as in our baseline with 30 labor groups. We report results for both approaches and our baseline in Table 17. Our baseline results are largely

robust to decreasing the number of labor groups.

G Between-within wage decomposition

Our baseline model implies that the average wage of workers in group λ is the same across all equipment-occupation pairs, an implication that is clearly rejected in the data. Here we argue that, in the data, differences in wage levels across occupation (together with worker re-allocation across occupations) are not the main driver of changes in labor group average wages over time. To do so, we conduct a between-within decomposition of changes in the average wage of group λ , $w_t(\lambda)/w_t$, where w_t is the composition-adjusted average wage across all labor groups. We consider variation in average wages within a labor group across occupations but not across equipment types, $w_t(\lambda, \omega)$, because the October CPS contains wage data for only a subset of observations (those respondents in the Outgoing Rotation Group). Measures of average wages across workers employed in particular (κ, ω) pairs would therefore likely be noisy.

The following accounting identity must hold at each t ,

$$\frac{w_t(\lambda)}{w_t} = \sum_{\omega} \frac{w_t(\lambda, \omega)}{w_t} \pi_t(\lambda, \omega),$$

and, therefore, we can write

$$\Delta \frac{w_t(\lambda)}{w_t} = \sum_{\omega} \Delta \frac{w_t(\lambda, \omega)}{w_t} \bar{\pi}_t(\lambda, \omega) + \sum_{\omega} \frac{\overline{w_t(\lambda, \omega)}}{w_t} \Delta \pi_t(\lambda, \omega), \quad (47)$$

where $\Delta x_t = x_{t_1} - x_{t_0}$ and $\bar{x}_t = (x_{t_1} + x_{t_0})/2$. The first term on the right-hand side of the equation (47) is the within component whereas the second term is the between component. According to the model, the contribution to changes in wages of the within component should be 100 percent for each labor group. We conduct this decomposition using the MORG CPS data between 1984 and 2003 for each of 30 labor groups and find that the median contribution across labor groups of the within component is above 86 percent. Hence, while in practice there are large differences in average wages for a labor group across occupations, these differences do not appear to be first order in explaining changes in labor group average wages over time.

H International trade in sectoral goods

We consider an extension of the model that incorporates sectors whose output is physically traded. We show that we can disaggregate $\mu_t(\omega)$ into sector shifters and within-sector occupation shifters. Moreover, moving to autarky corresponds to setting occupation shifters as functions of sectoral export and import shares in the initial open-economy equilibrium, similarly to the expression for $\hat{\mu}_n(\omega)$ in equation (34) above. Finally, we show that if both elasticities of substitution across sectors and occupations are equal to one, then changes in relative wages in the autarky counterfactual in the model with trade in sectoral goods is identical to that in the model with trade in occupation services (but no trade in sectoral output) in which export and import shares by occupation are calculated using sectoral trade data according to equations (35) and (36).

Sectors are denoted by σ . As in Section 5, we omit time subscripts. The final good in country n combines absorption from each sector, $D_n(\sigma)$, according to

$$Y_n = \left(\sum_{\sigma} \mu_n(\sigma)^{1/\rho_{\sigma}} D_n(\sigma)^{(\rho_{\sigma}-1)/\rho_{\sigma}} \right)^{\rho_{\sigma}/(\rho_{\sigma}-1)} \quad (48)$$

where ρ_{σ} is the elasticity of substitution across sectors. Absorption of sectoral σ in country n is a CES aggregate of sector σ goods sourced from all countries in the world,

$$D_n(\sigma) = \left(\sum_i D_{in}(\sigma)^{\frac{\eta(\sigma)-1}{\eta(\sigma)}} \right)^{\eta(\sigma)/(\eta(\sigma)-1)}, \quad (49)$$

where $D_{in}(\sigma)$ is absorption in country n of sector σ sourced from country i , and $\eta(\sigma) > 1$ is the elasticity of substitution across source countries for equipment σ . $d_{ni}(\sigma) \geq 1$ denotes the units of σ output that must be shipped from origin country n in order for one unit to arrive in destination country i ; with $d_{nn}(\sigma) = 1$ for all n .

Sector σ output is produced, as in our baseline model, combining the service of occupations according to

$$Y_n(\sigma) = \left(\sum_{\omega} \mu_n(\omega, \sigma)^{1/\rho} Y_n(\omega, \sigma)^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)}, \quad (50)$$

where $Y_n(\omega, \sigma) \geq 0$ denotes the services of occupation ω used in the production of sector σ , and $\mu_n(\omega, \sigma) \geq 0$ is an exogenous demand shifter for occupation ω in sector σ . Output

of sector σ in country n must equal its use across countries,

$$Y_n(\sigma) = \sum_i d_{ni}(\sigma) D_{ni}(\sigma). \quad (51)$$

Given that we abstract from international trade in occupations, total output of occupation ω must be equal to its demand across sectors, $Y_n(\omega) = \sum_\sigma Y_n(\omega, \sigma)$. Occupations are produced exactly as in our baseline specification: a worker's productivity depends only on her occupation ω , and not on her sector of employment σ .³⁸ Total output of occupation ω , $Y_n(\omega)$, is equal to the sum of output across all workers employed in ω .

The equilibrium quantity of occupation ω used in the production of sector σ in country n is

$$Y_n(\omega, \sigma) = \mu_n(\omega, \sigma) \left(\frac{p_n(\omega)}{p_{nn}(\sigma)} \right)^{-\rho} Y_n(\sigma), \quad (52)$$

where $p_{nn}(\sigma)$ denotes the output price of sector σ in country n , given by

$$p_{nn}(\sigma) = \left(\sum_\omega \mu_n(\omega, \sigma) p_n(\omega)^{1-\rho} \right)^{\frac{1}{1-\rho}}. \quad (53)$$

Absorption of sector σ in country n is

$$D_n(\sigma) = \mu_n(\sigma) \left(\frac{p_n(\sigma)}{p_n} \right)^{-\rho\sigma} Y_n, \quad (54)$$

where the absorption price of sector σ in country n , $p_n(\sigma)$, is given by

$$p_n(\sigma) = \left[\sum_i (d_{in}(\sigma) p_{ii}(\sigma))^{1-\eta(\sigma)} \right]^{\frac{1}{1-\eta(\sigma)}} \quad (55)$$

and absorption of sector σ in country n sourced from country i is given by

$$D_{in}(\sigma) = \left(\frac{d_{in}(\sigma) p_{ii}(\sigma)}{p_n(\sigma)} \right)^{-\eta(\sigma)} D_n(\sigma) \quad (56)$$

³⁸Accordingly, for example, an individual worker provides the same efficiency units of labor as an executive in an airplane-producing sector or as an executive in a textile-producing sector; although the airplane-producing sector may demand relatively more output from the executive occupation. It is straightforward to assume, alternatively, that worker productivity depends both on occupation and sector of employment, $T_t(\lambda, \kappa, \omega, \sigma) \varepsilon(z, \kappa, \omega, \sigma)$. Our estimation approach extends directly to this alternative assumption; however, in practice, the data become sparser, in the sense that there are many $(\lambda, \kappa, \omega, \sigma, t)$ for which $\pi_t(\lambda, \kappa, \omega, \sigma) = 0$.

The equations determining the allocation of workers over equipment types and occupations, $\pi_n(\lambda, \kappa, \omega)$, and the average wage $w_n(\lambda)$ are the same as in the baseline model and are given by (3) and (4). The occupation market clearing condition is given by

$$p_n(\omega) \sum_{\sigma} Y_n(\omega, \sigma) = \frac{1}{1-\alpha} \zeta_n(\omega), \quad (57)$$

Using equations (52)-(56), we can re-write equation (57) as the closed-economy counterpart (i.e. equation 5),

$$\mu_n(\omega) p_n^{1-\rho}(\omega) E_n = \frac{1}{1-\alpha} \zeta_n(\omega)$$

where the occupation shifter (which we treat as a parameter in the closed economy model without sectors) is given by³⁹

$$\mu_n(\omega) = \sum_{\sigma} \mu_n(\omega, \sigma) p_{nn}(\sigma)^{\rho-\eta(\sigma)} \sum_i d_{ni}(\sigma)^{1-\eta(\sigma)} p_i(\sigma)^{\eta(\sigma)-\rho\sigma} \mu_i(\sigma) p_i^{\rho\sigma-1} \frac{E_i}{E_n}. \quad (58)$$

Suppose that country n is in autarky, that is $d_{ni}(\sigma) \rightarrow \infty$ and $d_{ni}(\sigma) \rightarrow \infty$ for all $n \neq i$. Normalizing $p_n = 1$, the occupation shifter $\mu_n(\omega)$ is

$$\mu_n(\omega) = \sum_{\sigma} \mu_i(\sigma) \mu_n(\omega, \sigma) p_{nn}(\sigma)^{\rho-\rho\sigma}$$

and if $\rho = \rho_{\sigma}$,

$$\mu_n(\omega) = \sum_{\sigma} \mu_i(\sigma) \mu_n(\omega, \sigma)$$

In this case, we can disaggregate $\mu_n(\omega)$ into sector shifters and within-sector occupation shifters.⁴⁰

Suppose now that country n is initially open to trade at time t and then moves to autarky. Setting $\hat{d}_{in}(\sigma) = \hat{d}_{ni}(\sigma) = \infty$ for all $i \neq n$, the change in the occupation shifters defined in equation (58) is

$$\hat{\mu}_n(\omega) = \sum_{\sigma} v_n(\sigma|\omega) \hat{p}_{nn}(\sigma)^{\rho-\rho\sigma} f_{nn}(\sigma) s_{nn}(\sigma)^{\frac{\eta(\sigma)-\rho\sigma}{1-\eta(\sigma)}} \quad (59)$$

³⁹If all sectors share the same occupation intensities, that is $v_n(\omega|\sigma) \equiv \frac{p_n(\omega)Y_n(\omega,\sigma)}{\sum_{\omega'} p_n(\omega')Y_n(\omega',\sigma)}$ is independent of σ , and if $\hat{\mu}_n(\omega, \sigma) = \hat{\mu}_n(\omega)$, then irrespective of the value of ρ and ρ_{σ} , the model with sectors is equivalent to the baseline model where the occupation shifter is given by $\hat{\mu}_n(\omega)$.

⁴⁰In autarky, if all sectors share the same occupation intensities, that is $v_n(\omega|\sigma) \equiv \frac{p_n(\omega)Y_n(\omega,\sigma)}{\sum_{\omega'} p_n(\omega')Y_n(\omega',\sigma)}$ is independent of σ , and if $\hat{\mu}_n(\omega, \sigma) = \hat{\mu}_n(\omega)$ then, irrespective of the values of ρ and ρ_{σ} , this model is equivalent to the baseline model where occupation shifters are given by $\hat{\mu}_n(\omega)$.

where the change in the output price of sector σ is

$$\hat{p}_{nn}(\sigma) = \left[\sum_i v_n(\omega|\sigma) \hat{p}_n(\omega)^{1-\rho} \right]^{\frac{1}{1-\rho}}.$$

In equation (59), $v_n(\sigma|\omega) \equiv \frac{Y_n(\omega,\sigma)}{Y_n(\omega)}$ is the share of occupation ω employed in sector σ , $v_n(\omega|\sigma) \equiv \frac{p_{nn}(\omega)Y_n(\omega,\sigma)}{p_{nn}(\sigma)Y_n(\sigma)}$ is the occupation ω intensity in the production of sector σ in country n , $f_{nn}(\sigma)$ is defined analogously to $f_{nn}(\omega)$ as the fraction of total sales of sector σ in country n purchased from itself,

$$f_{nn}(\sigma) = \frac{p_{nn}(\sigma) D_{nn}(\sigma)}{p_{nn}(\sigma) Y_n(\sigma)}$$

and $s_{nn}(\sigma)$ is defined analogously to $s_{nn}(\omega)$ as the the fraction of expenditures on sector σ in country n purchased from itself,

$$s_{nn}(\sigma) = \frac{p_{nn}(\sigma) D_{nn}(\sigma)}{p_n(\sigma) D_n(\sigma)}.$$

The intuition for the mapping between import and export shares in the initial equilibrium and the corresponding closed economy occupation shifters is very similar to that in the model with occupation trade, discussed in Section 6.

In the special case in which $\rho = \rho_\sigma = 1$, the expression for the changes in occupation shifters simplifies to

$$\hat{p}_n(\omega) = \sum_\sigma v_n(\sigma|\omega) \frac{f_{nn}(\sigma)}{s_{nn}(\sigma)},$$

where $f_{nn}(\sigma) / s_{nn}(\sigma)$ is the ratio of absorption to production in sector σ . Substituting the definitions above,

$$\hat{p}_n(\omega) = \frac{\sum_\sigma v_n(\omega|\sigma) (p_n(\sigma) D_n(\sigma))}{p_n(\omega) Y_n(\omega)}. \quad (60)$$

Now consider the specification of our model with occupation trade (and no sectoral trade). If $\rho = 1$, then by equation (34), occupation shifters in the autarky counterfactual are given by

$$\hat{p}_n(\omega) = \frac{f_{nn}(\omega)}{s_{nn}(\omega)}.$$

If import and export shares by occupation are calculated according to equations (35) and

(36), then

$$\begin{aligned}\hat{\mu}_n(\omega) &= \frac{f_{nn}(\omega)}{s_{nn}(\omega)} = \frac{\sum_{\sigma} \frac{p_n(\omega)Y_n(\omega,\sigma)}{p_{nn}(\sigma)Y_n(\sigma)} (p_n(\sigma) D_n(\sigma))}{\sum_{\sigma} \frac{p_n(\omega)Y_n(\omega,\sigma)}{p_{nn}(\sigma)Y_n(\sigma)} (p_{nn}(\sigma) Y_n(\sigma))} \\ &= \frac{\sum_{\sigma} v_n(\omega|\sigma) (p_n(\sigma) D_n(\sigma))}{p_n(\omega) Y_n(\omega)}\end{aligned}$$

which coincides with expression (60).