NBER WORKING PAPER SERIES

ACCOUNTING FOR CHANGES IN BETWEEN-GROUP INEQUALITY

Ariel Burstein
Eduardo Morales
Jonathan Vogel

Working Paper 20855
http://www.nber.org/papers/w20855

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2015

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, and Mike Waugh for helpful comments and David Autor and David Dorn for making their code available. Vogel thanks the Princeton International Economics Section for their support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Accounting for Changes in Between-Group Inequality
Ariel Burstein, Eduardo Morales, and Jonathan Vogel
NBER Working Paper No. 20855
January 2015, Revised January 2016
JEL No. E24,F16,J2

ABSTRACT

We perform a quantitative analysis of observed changes in U.S. between-group inequality between 1984 and 2003. We use an assignment framework with many labor groups, equipment types, and occupations in which changes in inequality are caused by changes in workforce composition, occupation demand, computerization, and labor productivity. We parameterize our model using direct measures of computer usage within labor group-occupation pairs and quantify the impact of each shock for various measures of between-group inequality. We find, for instance, that the combination of computerization and shifts in occupation demand account for roughly 80% of the rise in the skill premium, with computerization alone accounting for roughly 60%. We show theoretically how computerization and changes in occupation demand may be caused by international trade and quantify the impact of trade in computers on U.S. inequality.

Ariel Burstein
Department of Economics
Bunche Hall 8365
Box 951477
UCLA
Los Angeles, CA 90095-1477
and NBER
arielb@econ.ucla.edu

Eduardo Morales
Department of Economics
Princeton University
310 Fisher Hall
Princeton, NJ 08544
and NBER
ecmorales@princeton.edu

Jonathan Vogel
Department of Economics
Columbia University
420 West 118th Street
New York, NY 10027
and NBER
jvogel@columbia.edu

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*); see, e.g. Acemoglu and Autor (2011). For example, 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.[1]

Following Katz and Murphy (1992), 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). Related to the first hypothesis, Table 1 shows that computer use rose dramatically between 1984 and 2003 and that computers are used more intensively by educated workers and women.[2] Related to the second hypothesis, Figure 1 shows that education- and female-intensive occupations grew relatively quickly over the same time period; see Table 12 in Appendix B for details.

		1984	1989	1993	1997	2003
All		27.4	40.1	49.8	53.3	57.8
Gender	Female	32.8	47.6	57.3	61.3	65.1
	Male	23.6	34.5	43.9	47.0	52.1
Education	College degree	45.5	62.5	73.4	79.8	85.7
	No college degree	22.1	32.7	41.0	43.7	45.3

Table 1: Share of hours worked with computers

The goal of our paper is to use a general equilibrium model together with detailed data on factor allocation to quantify the impact of computerization and changes in occupation demand on between-group inequality in the United States; in addition, we aim to quantify the role of international trade in shaping these forces. We base our analysis

[1]The relative importance of between- and within-group inequality is an area of active research. Autor (2014) concludes: "In the U.S., for example, 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." On the other hand, Helpman et al. (2012) conclude: "Residual wage inequality is at least as important as worker observables in explaining the overall level and growth of wage inequality in Brazil from 1986-1995."

[2]We describe our data sources in depth in Section 4.1 and Appendix B.

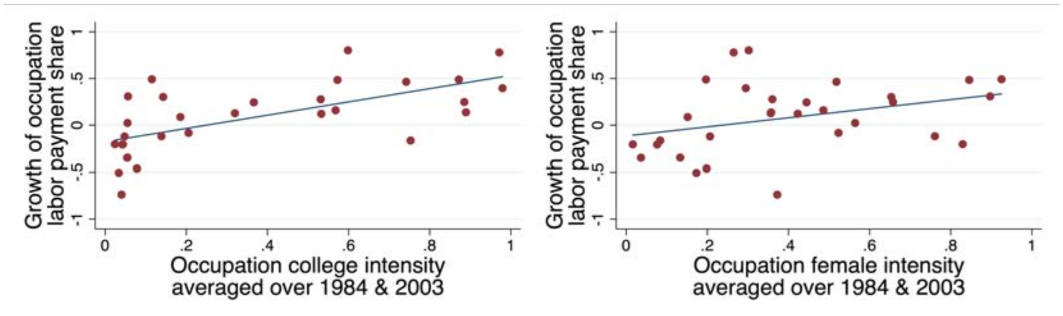

Figure 1: Growth (1984-2003) of the occupation share of labor payments and the average (1984 & 2003) of the share of workers in the occupation who have a college degree (left) and are female (right)

on an assignment model with many groups of workers and many occupations—building on Eaton and Kortum (2002), Lagakos and Waugh (2013), and Hsieh et al. (2013)—which we extend to incorporate many types of equipment. We use this framework because it allows for potentially rich patterns of complementarity between computers, labor groups and occupations and, in spite of its high dimensionality, remains tractable enough to estimate its key parameters, perform aggregate counterfactuals, and incorporate international trade in a parsimonious manner. Moreover, the model's aggregate implications for relative wages nest those of workhorse models of between-group inequality, e.g. Katz and Murphy (1992) and Krusell et al. (2000).

The impact of changes in the economic environment on between-group inequality is shaped by comparative advantage between labor groups, equipment types, and occupations. Consider, for example, the potential impact of an increase in the productivity of computers on a given labor group—such as educated workers or women—that uses computers intensively. A labor group may use computers intensively for two reasons. First, it may have a comparative advantage with computers, in which case this group would use computers relatively more within occupations, as we show is the case in the data for more educated workers. In this case, our model predicts that computerization increases the relative wage of this labor group. Second, a labor group may have a comparative advantage in occupations in which computers have a comparative advantage, in which case this group would be allocated disproportionately to occupations in which all workers are relatively more likely to use computers, as we show is the case in the data for women. In this case, our model predicts that computerization may increase or decrease the relative wage of this group depending on the degree of substitutability between occupations.[3]

[3]This implies that our model is flexible enough to allow computerization to increase the relative wage of workers who are relatively productive using computers and reduce the relative wage of workers employed

To quantify the impact that computerization and changes in occupation demand, labor supply, and labor productivity have had on relative wages in the U.S. over recent decades, we therefore must: measure these changes; identify comparative advantage between labor groups, equipment types, and occupations; and estimate the elasticity of substitution between occupations (which shapes the responsiveness of occupation prices to shocks) and an elasticity shaping the within-worker dispersion of productivity across occupation-equipment type pairs (which shapes the responsiveness of group average wages to our measures of changes in occupation prices and equipment productivity).

Comparative advantage can be inferred directly from data on the allocation of workers to equipment type-occupation pairs. Over any time period, we measure changes in the the determinants of relative wages as follows. Changes in "equipment productivity" that result in computerization can be inferred from changes in the allocation of workers to computers within labor group-occupation pairs; focusing on within labor group-occupation pairs is important because aggregate computer usage rise in the absence of changes in equipment productivity if either labor groups that have a comparative advantage using computers or occupations that have a comparative advantage with computers grow. Changes in occupation demand ("occupation shifters") can be inferred from changes in the allocation of workers to occupations within labor group-equipment type pairs and in the labor income shares across occupations. Changes in labor supply ("labor composition") are directly observed in the data. Finally, we measure labor productivity as a residual to exactly match changes in the average wage of each labor group.[4]

In order to estimate the two key elasticities, we derive moment conditions that are consistent with equilibrium relationships generated by our model. We identify the structural elasticity of substitution between occupations using the reduced-form elasticity of occupation labor income to measured occupation prices. We identify the structural elasticity shaping the within-worker dispersion of productivity across occupation-equipment type pairs using the reduced-form elasticity of labor-group-specific average wages to labor-group-specific weighted averages of changes in occupations' prices and equipment types' productivities. Because of the endogeneity of occupation prices, when estimating both elasticities we use a function of our measures of changes in equipment productivity as an instrument.

in occupations in which computers are particularly productive, as described by, e.g., Autor et al. (1998) and Autor et al. (2003).

[4]Residual labor productivity affects the relative productivity of labor groups, independent of the equipment they use and occupations in which they are employed. Factors shaping residual labor productivity include, for example, discrimination and the quality of education and health systems; see e.g. Card and Krueger (1992) and Goldin and Katz (2002).

As is evident from the previous discussion, our procedure crucially requires information for multiple years on the shares of workers of each labor group that belong to each equipment type-occupation pair. We obtain this information for the U.S. from the October Current Population Survey (CPS) Computer Use Supplement, which provides data for five years (1984, 1989, 1993, 1999, and 2003) on whether a worker has direct or hands on use of a computer at work—be it a personal computer, laptop, mini computer, or mainframe—and on worker characteristics, hours worked, and occupation; see, e.g., Krueger (1993) and Autor et al. (1998) for previous studies using the October Supplement. For our purposes, this data is not without limitations: it imposes a narrow view of computerization that does not capture, e.g., automation of assembly lines; it only provides information on the allocation of workers to one type of equipment, computers; it does not detail the share of each worker's time at work spent using computers; and it does not extend beyond 2003.[5]

We find that computerization alone accounts for roughly 60% 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 inequality (e.g., the relative wage of workers with graduate training relative to high school dropouts). This result from our model is driven by the following three observations in the data. First, we observe 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). Second, more educated workers use computers within occupations relatively more than less educated workers, which—together with computerization—yields a rise in their relative wages according to our model. Third, more educated workers are also disproportionately employed in occupations in which all workers use computers relatively intensively, which—together with computerization and an estimated elasticity of substitution between occupations greater than one—also yields a rise in their relative wages according to our model. The combination of computerization and occupation shifters accounts for roughly 80% of the rise in the skill premium, leaving only 20% to be explained by labor productivity, the shock measured as a residual to match 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

[5]Our sample period, 1984-2003, accounts for a substantial share of the increase in the skill premium and reduction in the gender gap in the U.S. since the late 1960s. In Appendix D we show that the German *Qualification and Working Conditions* survey, which alleviates some of the limitations of the CPS data, reveals similar patterns of comparative advantage in Germany as in the U.S.

male to female workers). Computerization reduces the gender gap in spite of the fact that, unlike educated workers, women do not have a comparative advantage using computers. 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.

Whereas in our baseline model we treat computerization as an exogenous change, in Section 7 we study the extent to which this observed change in computer productivity is a consequence of international trade in equipment. We focus on equipment trade both because computerization is a dominant force in accounting for the rise in the skill premium (as we show in the present paper), and because many countries import a large share of their equipment (as shown in Eaton and Kortum (2001)). Theoretically, we show that the procedure to quantify the impact on relative wages of moving to autarky in equipment trade is equivalent to the procedure we follow in our baseline model to calculate changes in relative wages in a closed economy with the only difference that the computerization shock is now measured as a simple function of import shares of different equipment types in the open economy. 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. Using this latter result, we quantify the impact of trade in equipment goods on between-group inequality in the U.S. between 1984 and 2003 for different values of the trade elasticity. In Appendix I, we extend the model to additionally allow for trade in occupations and in intermediate industries and show how the occupation shifters introduced in our baseline model become simple functions of occupations' and industries' import and export shares. In practice, however, occupation-specific import and export shares are hard to obtain, and this prevents us from measuring the impact that trade in occupations has had on wage inequality in the U.S.

Our paper is organized as follows. In Section 2, we discuss the related literature. We describe our framework, characterize its equilibrium, and discuss its mechanisms in Section 3. We parameterize the model in Section 4, describe our baseline closed-economy results in Section 5, and consider various robustness exercises and sensitivity analyses in Section 6. Finally, we extend our model to incorporate and quantify the impact of international trade in equipment in Section 7 and conclude in Section 8. Additional details and robustness exercises are relegated to appendices.

2 Literature

We follow Hsieh et al. (2013) and Lagakos and Waugh (2013) in using an assignment model of the labor market parameterized with a Fréchet distribution. We extend these models by introducing equipment types as another dimension along which workers sort, and international trade as another set of forces determining the equilibrium assignment of workers to occupations and equipment types. This framework allows us to incorporate multiple labor groups (enabling us to study within a unified framework the sources of changes in inequality in the United States between multiple groups of workers, such as the decline in the gender gap and the rise in the return to education across disaggregated education groups) and multiple occupations and equipment types (allowing us to study within a unified framework the impact of changes in occupation demand shifters, computer productivity, labor productivity and labor composition). Furthermore, the framework guides us in how to use detailed data on the allocation of different labor groups to occupations and types of equipment to measure comparative advantage between labor groups, occupations and types of equipment, the changes in occupation shifters and equipment productivity, and key elasticities.

In trying to explain the evolution of between-group inequality as a function of changes in observables, rather than explaining patterns in the skill premium and gender gap through latent skill- or gender-biased technological change, 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 which 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. We corroborate the findings in Krusell et al. (2000) and extend them by additionally considering the impact of equipment productivity growth on the gender gap and other measures of between-group inequality. However, we rely on a distinct methodology. Whereas Krusell et al. (2000) identify the degree of capital-skill complementarity exclusively using aggregate time series data, our approach additionally leverages information on the allocation of workers to computers and occupations and, consequently, yields parameter estimates shaping the degree of equipment-labor group complementarity that are robust to allowing for exogenous 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.

Lee and Wolpin (2010) use a dynamic model of endogenous human capital accumulation to study the evolution of relative wages and labor supply and find that skill-biased technical change (the residual) plays the central role in explaining changes in the skill premium. By allowing for a greater degree of disaggregation (e.g. 30 occupations) and exploiting detailed data on factor allocation, our results substantially reduce the role of changes in the residual (labor productivity) in shaping changes in the skill premium. On the other hand, in contrast to Lee and Wolpin (2010), we treat labor composition as exogenous.[6]

Two important related papers use differential 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 the effect of computerization (as well as other shocks) on changes over time in between-group inequality. Instead of relying on regional variation in the exposure to computerization, we make use of detailed data on computer usage within labor group-occupation pairs.

Our focus on occupation shifters is related to a broad literature using shift-share analyses. For instance, Autor et al. (2003) use a shift-share analysis and find that occupation shifters account for more than 50% of the relative demand shift favoring college labor between 1970 and 1988; we arrive at a similar result using a shift-share analysis between 1984 and 2003 in our data even though our decomposition exercise implies that occupation shifters account for a relatively small share of the rise in the relative wage of college labor. Under some assumptions (Cobb-Douglas utility and production functions), shift-share analyses structurally decompose changes in wage bill shares, i.e. changes in labor income for one labor group relative to the sum of labor payments across all labor groups, into within and between occupation shifters; see e.g. Katz and Autor (1999). However, changes in wage bill shares can be very different from changes in relative wages (on which we focus), especially when changes in labor composition are large, as they are in the data. Therefore, the results of shift-share analyses are not inconsistent with the finding in our decomposition exercise that occupation shifters account for only 19% of the rise in the

[6]Extending 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. In this case, our measures of shocks—to occupation shifters, equipment productivity, and labor productivity—and our estimates of model parameters are robust to extending our model to endogenize the supply of each labor group. In our counterfactual exercises, we fix labor composition to isolate the direct effect of individual shocks to occupation shifters, equipment productivity, and labor productivity on labor demand and wages. In order to take into consideration the accumulation of occupation-specific human capital as studied in, e.g., Kambourov and Manovskii (2009a) and Kambourov and Manovskii (2009b), 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. We leave this for future work.

relative wage of college labor.[7]

The 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, for each labor group, we estimate measures of comparative advantage that vary flexibly across occupations, independent of the similarity in the task composition of these occupations. Similarly, our procedure to measure occupation-specific demand shocks does not use any information on the task composition of occupations. Consequently, the counterfactual exercises that we perform have implications for the growth of different occupations that do not rely on information on their task composition. As we show in Appendix F.1, one can use the predictions of our model for the growth rates of each occupation we consider to understand which of our measured shocks are causing the correlation between occupation growth and occupation characteristics documented in Autor et al. (2003). For instance, we find that computerization generates an expansion in those occupations that are intensive in non-routine cognitive analytical and interpersonal tasks and a contraction in occupations that are intensive in non-routine manual physical tasks.

In relying on detailed 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 has identified 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. Instead, a component of our structural estimation builds on the empirical approach suggested by Acemoglu and Autor (2011) as a stylized example of how their assignment model might be brought to the data. Specifically, we regress changes in labor-group-specific wages on interactions between beginning-of-sample measures of the specialization of the different labor groups across occupations and equipment types and measures of changes in those occupations' prices and equipment types' productivities. The possible endogeneity concerns affecting this regression are very different from those described in DiNardo and Pischke (1997) and are discussed in detail in Section 4.3.

In modeling international trade, we empirically study the theoretical insights of Costinot and Vogel (2010) and Costinot and Vogel (Forthcoming) regarding the impact of international trade on inequality in a high-dimensional environment. We show, as in concurrent work by Galle et al. (2015) and Lee (2015), that one can use a similar approach to that

[7] Firpo et al. (2011) uses a statistical model of wage setting to investigate the contribution of changes in the returns to occupational tasks compared to other explanations such as de-unionization and changes in the labor market wide returns to general skills. Our paper complements theirs by incorporating general equilibrium effects and explicitly modeling the endogenous allocation of factors.

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. Finally, in modeling international trade in equipment, we extend the quantitative analyses of Burstein et al. (2013) and Parro (2013), who study the impact of trade in capital equipment on the skill premium using the model of Krusell et al. (2000).

3 Model

In this section we introduce the baseline version of our model, characterize its equilibrium, and show how to use it to decompose observed changes in relative average wages into four types of changes in the economic environment: labor composition, equipment productivity, occupation shifters, and labor productivity. Finally, we provide intuition for how each of these changes affects relative wages.

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.

Occupations are used to produce a single final good according to a constant elasticity of substitution (CES) production function

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

where $\rho > 0$ is the elasticity of substitution across occupations, $Y_t(\omega) \geq 0$ is the endogenous output of occupation ω, and $\mu_t(\omega) \geq 0$ is an exogenous demand shifter for occupation ω.[8] The final good is used to produce consumption, C_t, and equipment, $Y_t(\kappa)$,

[8] We show in Appendix I that we can disaggregate $\mu_t(\omega)$ further into sector shifters and within-sector occupation shifters. We also show how changes in the extent of international trade/offshoring in sectoral output and occupation output may generate changes in these sector shifters and within-sector occupation shifters. For now, however, we combine sector and within-sector occupation shifters and treat them as exogenous.

according to the resource constraint

$$Y_t = C_t + \sum_{\kappa} p_t(\kappa) Y_t(\kappa), \tag{2}$$

where $p_t(\kappa)$ denotes the cost of a unit of equipment κ in terms of units of the final good.[9]

Occupation output is produced by perfectly competitive production units. A unit hiring k units of equipment type κ and l efficiency units of labor group λ produces $k^{\alpha}\left[T_t(\lambda,\kappa,\omega)\, l\right]^{1-\alpha}$ units of output, where α denotes the output elasticity of equipment in each occupation 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.[10] 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 $T_t(\lambda',\kappa',\omega)/T_t(\lambda',\kappa,\omega) \geq T_t(\lambda,\kappa',\omega)/T_t(\lambda,\kappa,\omega)$. Labor-occupation and equipment-occupation comparative advantage are defined symmetrically.

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 $\epsilon(z) \in (\underline{\epsilon}_{\lambda}, \bar{\epsilon}_{\lambda})$—with $0 < \underline{\epsilon}_{\lambda} \leq \bar{\epsilon}_{\lambda} < \infty$—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 (and prove this is without loss of generality in 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 $\epsilon(z)$ is independent of the distribution of $\varepsilon(z,\kappa,\omega)$ across (κ,ω). Second, each $\varepsilon(z,\kappa,\omega)$ is drawn independently from a Fréchet distribution with cumulative distribution function $G(\varepsilon) = \exp\left(\varepsilon^{-\theta}\right)$, where a higher value of $\theta > 1$ decreases the within-worker dispersion of efficiency units across (κ,ω) pairs.[11]

[9] Here we assume that $Y_t(\kappa)$ denotes the supply of equipment κ in period t. This is equivalent to assuming that equipment fully depreciates every period. Alternatively, we could assume that $Y_t(\kappa)$ denotes investment in capital equipment κ, which depreciates at a given rate. All our counterfactual exercises are consistent with this alternative model with capital accumulation, interpreting them as comparisons across balanced growth paths in which the real interest rate and the growth rate of relative productivity across equipment types are both constant over time. In our baseline model, we treat the cost of producing equipment as exogenous. We show in Section 7 how changes in the extent of international trade in equipment generates endogenous changes in $p_t(\kappa)$.

[10] We restrict α to be common across ω because we do not have the data to estimate a different value of $\alpha(\omega)$ to each ω. Moreover, without affecting any of our results, we can extend the model to incorporate other inputs such as structure or intermediate inputs s, which are produced linearly using the final good and which enter the production function multiplicatively as $s^{1-\eta}\left(k^{\alpha}\left[T_t(\lambda,\kappa,\omega)\, l\right]^{1-\alpha}\right)^{\eta}$. In either case, α is the share of equipment relative to the combination of equipment and labor.

[11] See Adao (2015) for an approach relaxing these two restrictions in an environment in which each worker faces exactly two choices. In Appendix E we show that our results hold exactly if we also allow

The assumption that $\varepsilon(z, \kappa, \omega)$ is distributed Fréchet is made for analytical tractability; it implies that the average wage of a labor group is a CES function of occupation prices and equipment productivity.[12]

Total output of occupation ω, $Y_t(\omega)$, is the sum of output across all units producing occupation ω using any labor group λ and equipment type κ. 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 of two frameworks that have been used commonly to study between-group inequality: the *canonical model*, as named in Acemoglu and Autor (2011), and an extension of the canonical model that incorporates capital-skill complementarity; see e.g. Katz and Murphy (1992) and Krusell et al. (2000), respectively.[13]

3.2 Equilibrium

We characterize the competitive equilibrium, first in partial equilibrium—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 equipment κ relative to the price of the final good (which we normalize to one) is simply $p_t(\kappa)$. 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 where $p_t(\omega)$ 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

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

---

for correlation for each z of $\varepsilon(z, \kappa, \omega)$ across (κ, ω) pairs. Moreover, we also conduct our analysis allowing for variation across labor groups in the dispersion parameter θ and show that our quantitative results are robust.

[12] The wage distribution implied by this assumption is a good approximation to the observed distribution of individual wages; see e.g. Saez (2001) and Figure 5 in the Appendix.

[13] The aggregate implications of our model for relative wages are equivalent to those of the canonical model if we assume no equipment (i.e. $\alpha = 0$) and two labor groups, each of which has a positive productivity in only one occupation. The model of capital-skill complementarity is an extension of the canonical model in which there is one type of equipment and the equipment share is positive in one occupation and zero in the other (i.e. $\alpha = 0$ for the latter occupation).

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

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

The higher is θ—i.e. the less dispersed are efficiency units across (κ, ω) pairs—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

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

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. Similar conditions hold for the allocation of workers to occupations (within an equipment type) and for the allocation of equipment to occupations (within a labor group).

The average wage of workers in group λ teamed with equipment κ in occupation ω, denoted by $w_t(\lambda, \kappa, \omega)$, 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), where[14]

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

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 occupation ω,

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

where $E_t \equiv (1-\alpha)^{-1} \sum_\lambda w_t(\lambda) L_t(\lambda)$ is total income and $\zeta_t(\omega) \equiv \sum_{\lambda,\kappa} w_t(\lambda) L_t(\lambda) \pi_t(\lambda,\kappa,\omega)$ is total labor income in occupation ω. The left-hand side of equation (5) is expenditure on occupation ω and the right-hand side is total income earned by factors employed in occupation ω. In equilibrium, the aggregate quantity of the final good is $Y_t = E_t$, 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 determined by equation (2).

3.3 Decomposing changes in relative wages

Our objective in this paper is to use the framework described above to quantify the relative importance of different changes in the economic environment in determining changes in relative wages between labor groups. In what follows, we impose the assumption that $T_t(\lambda,\kappa,\omega)$ can be expressed as

$$T_t(\lambda,\kappa,\omega) \equiv T_t(\lambda) T_t(\kappa) T_t(\omega) T(\lambda,\kappa,\omega). \tag{6}$$

Accordingly, whereas we allow labor group, $T_t(\lambda) \geq 0$, equipment type, $T_t(\kappa) \geq 0$, and occupation, $T_t(\omega) \geq 0$, productivity to vary over time, we impose that the interaction between labor group, equipment type, and occupation productivity, $T(\lambda,\kappa,\omega) \geq 0$, is

[14] The implication that average wage *levels* are common across occupations within λ (which is inconsistent with the data) obviously implies that the changes in wages will also be common across occupations within λ. In Appendix G we decompose the observed *changes* in average wages for each λ into a between occupation component (which is zero in our model) and a within occupation component and show that the between component is small. Furthermore, we show that by incorporating preference shifters for working in different occupations that generate compensating differentials, similar to Heckman and Sedlacek (1985), our model can generate differences in average wage levels across occupations within a labor group.

constant over time. That is, we assume that comparative advantage is fixed over time. This restriction allows us to separate the impact on relative wages of λ-, κ-, and ω-specific productivity shocks. In Section 6.3 we relax this restriction.

Given the restriction in equation (6), one could use the model described above to decompose changes in relative average wages for any two labor groups observed in a particular economy between any two points in time into changes in labor composition, $L_t(\lambda)$; changes in labor productivity, $T_t(\lambda)$; changes in occupation demand, $\mu_t(\omega)$; changes in occupation productivity, $T_t(\omega)$; changes in equipment productivity, $T_t(\kappa)$; and changes in equipment production costs, $p_t(\kappa)$. This requires being able to measure the changes in these six exogenous components that have taken place in the economy of interest between the same two points in time. However, given available data for the U.S., we cannot separately identify changes in occupation demand and occupation productivity; instead, we combine them in a composite occupation shifter, $a_t(\omega) \equiv \mu_t(\omega) T_t(\omega)^{(1-\alpha)(\rho-1)}$. Similarly, we do not separately identify changes in equipment productivity and equipment production costs; instead, we combine them in a composite equipment productivity, $q_t(\kappa) \equiv p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} T_t(\kappa)$. Hence, in our empirical analysis, we decompose the observed changes in relative wages in the U.S. into four shocks: (i) labor composition, (ii) equipment productivity, (iii) occupation shifters, and (iv) labor productivity.[15]

Given measures of changes in these shocks and model parameters, we will quantify the impact of observed changes in each shock on relative wages across labor groups. To solve for changes in relative wages as a function of changes in any of these shocks, 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—using equations (4) and (6)—are given by

$$\hat{w}(\lambda) = \hat{T}(\lambda) \left[\sum_{\kappa,\omega} \left(\hat{q}(\omega) \hat{q}(\kappa) \right)^{\theta} \pi_{t_0}(\lambda, \kappa, \omega) \right]^{1/\theta}, \tag{7}$$

where $q_t(\omega)$ denotes transformed occupation prices, i.e. $q_t(\omega) \equiv p_t(\omega)^{1/(1-\alpha)} T_t(\omega)$.

[15]The two components of the computerization shock, $T_t(\kappa)$ and $p_t(\kappa)$, could be measured separately using information on the price of computers and other types of equipment, which is subject to known quality-adjustment issues raised by, e.g., Gordon (1990). Similarly, the two components of the occupation shock, $T_t(\omega)$ and $\mu_t(\omega)$, could be measured separately using information on changes in occupation prices, which are hard to measure in practice. Notice that, 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 counterfactual or hypothetical changes in the economic environment on relative average wages, then it would not be necessary to combine changes in equipment productivity and equipment production costs into a single shock, or similarly to combine changes in occupation demand and occupation productivity into a single composite shock.

Changes in transformed occupation prices—using equations (3), (5), and (6)—are determined by the following system of equations

$$\hat{\pi}(\lambda,\kappa,\omega) = \frac{(\hat{q}(\omega)\,\hat{q}(\kappa))^{\theta}}{\sum_{\kappa',\omega'}(\hat{q}(\omega')\,\hat{q}(\kappa'))^{\theta}\,\pi_{t_0}(\lambda,\kappa',\omega')}, \tag{8}$$

$$\hat{a}(\omega)\,\hat{q}(\omega)^{(1-\alpha)(1-\rho)}\,\hat{E} = \frac{1}{\zeta_{t_0}(\omega)}\sum_{\lambda,\kappa} w_{t_0}(\lambda)\,L_{t_0}(\lambda)\,\pi_{t_0}(\lambda,\kappa,\omega)\,\hat{w}(\lambda)\,\hat{L}(\lambda)\,\hat{\pi}(\lambda,\kappa,\omega). \tag{9}$$

In this system, the forcing variables are the four shocks mentioned above: $\hat{L}(\lambda)$, $\hat{T}(\lambda)$, $\hat{a}(\omega)$, and $\hat{q}(\kappa)$. Given these variables, equations (7)-(9) yield the model's implied values of changes in average wages, $\hat{w}(\lambda)$, allocations, $\hat{\pi}(\lambda,\kappa,\omega)$, and transformed occupation prices, $\hat{q}(\omega)$.[16]

3.4 Intuition

The impact of shocks on wages. In partial equilibrium—i.e. for given changes in transformed occupation prices—changes in wages are proportional to changes in labor productivity, $\hat{T}(\lambda)$, and to a CES combination of changes in transformed occupation prices and equipment productivities, where the weight given to changes in each of these components depends on factor allocations in the initial period t_0, $\pi_{t_0}(\lambda,\kappa,\omega)$; see equation (7). An increase in occupation ω's transformed price or equipment κ's productivity between 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. In what follows we describe the general equilibrium impact on relative wages of changes in occupation shifters, labor composition, labor productivity, and equipment productivity.

Changes in occupation shifters and labor composition affect wages only indirectly (i.e. in general equilibrium) through transformed occupation prices.[17] Consider an increase in the occupation shifter $a_t(\omega)$, i.e. $\hat{a}(\omega) > 1$, arising either from an increase in the demand shifter for occupation ω or an increase (decrease) in the productivity of all workers employed in occupation ω if $\rho > 1$ ($\rho < 1$). This shock raises the transformed price

[16] Given available data, we can only measure relative shocks to occupation shifters, equipment productivity, and labor productivity. Specifically, rather than measure $\hat{T}(\lambda)$, $\hat{a}(\omega)$, and $\hat{q}(\kappa)$, we can only measure $\hat{T}(\lambda)/\hat{T}(\lambda_1)$, $\hat{a}(\omega)/\hat{a}(\omega_1)$, and $\hat{q}(\kappa)/\hat{q}(\kappa_1)$ for any arbitrary choice of a benchmark labor group, λ_1, occupation, ω_1, and type of equipment, κ_1. However, we can re-express equations (7), (8), and (9)—see Appendix C—so that changes in relative wages, $\hat{w}(\lambda)/\hat{w}(\lambda_1)$, transformed occupation prices, $\hat{q}(\omega)/\hat{q}(\omega_1)$, 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, $\hat{T}(\lambda)/\hat{T}(\lambda_1)$.

[17] In Section 4.3, we provide empirical evidence that supports our model's implication that changes in labor composition affect wages only indirectly through occupation prices.

of occupation ω (an increase in occupation productivity reduces the *primitive* occupation price, $p_t(\omega)$) and, therefore, the relative wages of labor groups that are disproportionately employed in occupation ω. Similarly, an increase in labor supply $L_t(\lambda)$ reduces the transformed prices of occupations in which group λ is disproportionately employed. This lowers the relative wage not only of group λ, but also of other labor groups employed in similar occupations as λ. An increase in labor productivity $T_t(\lambda)$ directly raises the relative wage of group λ and reduces the transformed prices of occupations in which this group is disproportionately employed, thus reducing the relative wages of labor groups employed in similar occupations as λ.[18]

Finally, consider the impact of a change in the productivity of equipment κ, i.e. $\hat{q}(\kappa) > 1$. The direct (i.e. partial equilibrium) impact of this shock is to raise the relative wages of labor groups that use κ intensively. In general equilibrium this shock also reduces the transformed prices of occupations in which κ is used intensively, lowering the relative wages of labor groups that are disproportionately employed in these occupations. Overall, the impact on relative wages of changes in equipment productivity depends on the value of ρ and on whether aggregate patterns of labor allocation across equipment types are mostly a consequence of variation in labor-equipment comparative advantage or mainly determined by variation in labor-occupation and equipment-occupation comparative advantage. While in practice all sources of comparative advantage are active, it is useful to consider two extreme cases.

If the only form of comparative advantage is between workers and equipment, then an increase in $q_t(\kappa)$ does not affect relative occupation prices (since all occupations are equally intensive in κ) and, therefore, the partial equilibrium effect—i.e. for given changes in transformed occupation prices—is the same as the general equilibrium effect. In this case an increase in $q_t(\kappa)$ will increase the relative wages of worker groups that use equipment κ more intensively in the initial period.

On the other hand, if there is no comparative advantage between workers and equipment but there is comparative advantage between workers and occupations and between equipment and occupations, then an increase in $q_t(\kappa)$ directly increases the relative wages of worker groups that use equipment κ intensively (those disproportionately employed in κ-intensive occupations) and indirectly, through transformed occupation prices, reduces the relative wages of worker groups employed in κ-intensive occupations. The relative

[18] Costinot 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.

strength of the direct and indirect channels depends on ρ. The relative wages of worker groups employed in κ-intensive occupations rise—i.e. the direct effect dominates the indirect occupation price effect—if and only if $\rho > 1$. Intuitively, an increase in $q_t(\kappa)$ acts like a positive productivity shock to the occupations in which κ has a comparative advantage. If $\rho > 1$, this increases employment and the relative wages of labor groups disproportionately employed in the occupations in which κ has a comparative advantage.

The role of θ and ρ. The parameter θ, which governs the degree of within-worker dispersion of productivity across occupation-equipment type pairs, determines the extent of worker reallocation in response to changes in transformed occupation prices and equipment productivities: a higher dispersion of idiosyncratic draws, as given by a lower value of θ, results in less worker reallocation. Concerning the role of θ for changes in labor group average wages, taking a first-order approximation of equation (7) at period t_0 allocations, we obtain:

$$\log \hat{w}(\lambda) = \log \hat{T}(\lambda) + \sum_{\kappa,\omega} \pi_{t_0}(\lambda,\kappa,\omega) \left(\log \hat{q}\left(\omega\right) + \log \hat{q}\left(\kappa\right)\right). \tag{10}$$

Therefore, for given changes in transformed occupation prices and equipment productivities, the change in average wages does not depend on θ, up to a first-order approximation.[19] 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 by affecting occupation price changes, the value of θ affects the response of average wages to shocks even up a first-order approximation.

The parameter ρ determines the elasticity of substitution between occupations, with a higher value of ρ reducing the responsiveness of occupation prices to shocks and hence, reducing the impact through occupation price changes of shocks on relative wages. For example, because labor composition only affects relative wages through occupation prices, a higher value of ρ reduces the impact of labor composition on relative wages. Similarly, as we described above, a higher value of ρ reduces the negative effects of computerization on wages of labor groups with a comparative advantage on occupations that use computers intensively.

[19]More generally, the shape of the distribution of $\varepsilon\left(z,\kappa,\omega\right)$ —which we have assumed to be Fréchet— does not matter for the first-order effect of any shock on average wage changes (given changes in occupation prices and equipment productivities) because, for any worker group λ, the marginal worker's wage is equal across occupations and equipment types (i.e. applying an envelope condition on the workers' assignment problem).

4 Parameterization

According to equations (7)-(9), the impact of changes in the economic environment between any two periods t_0 and t_1 on relative wages depends on: (i) period t_0 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, $\zeta_{t_0}(\omega)$; (ii) 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, $\hat{T}(\lambda)/\hat{T}(\lambda_1)$; and (iii) estimates of the parameters α, ρ, and θ.

4.1 Data

We use data from the Combined CPS May, Outgoing Rotation Group (MORG CPS) and the October CPS Supplement (October Supplement) in 1984, 1989, 1993, 1997, and 2003. We restrict our sample by dropping workers who are younger than 17 years old, do not report positive paid hours worked, or are self employed. Here we briefly describe our use of these sources; we provide further details in Appendix B. After cleaning, the MORG CPS and October Supplement contain data for roughly 115,000 and 50,000 individuals, respectively, in each year.

We divide workers into 30 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 and older). We consider two types of equipment: computers and other equipment. We use thirty occupations, which we list, together with summary statistics, in Table 12 in Appendix B.

We use the MORG CPS to construct total hours worked and average hourly wages by labor group by year.[20] We use the October Supplement to construct 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." 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. We construct $\pi_t(\lambda, \kappa', \omega)$ as the hours worked in occupation ω by λ workers who report

[20]We measure wages using the MORG CPS rather than the March CPS. Both datasets imply similar changes in average wages within a labor group. 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, which we will use in one of our sensitivity analyses; see Lemieux (2006).

that they use a computer κ' at work relative to the total hours worked by labor group λ in year t. Similarly, we construct $\pi_t(\lambda, \kappa'', \omega)$ as the hours worked in occupation ω by λ workers who report that they do not use a computer at work (where κ'' = other equipment) relative to the total hours worked by labor group λ in year t.[21]

Constructing factor allocations, $\pi_t(\lambda, \kappa, \omega)$, as we do introduces four limitations. First, our view of computerization is narrow. Second, at the individual level our computer-use variable takes only two values: zero or one. Third, we are not using any information on the allocation of non-computer equipment. Finally, the computer use question was discontinued after 2003.[22]

Factor allocation. By aggregating $\pi_t(\lambda, \kappa, \omega)$ across ω and λ, we showed in Table 1 that in the aggregate women and more educated workers use computers more intensively than men and less educated workers, respectively. Here we identify a few key patterns in the disaggregated $\pi_t(\lambda, \kappa, \omega)$ data.

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') relatively more than non-computer equipment (κ) within occupations (ω), the left panel of Figure 2 plots the histogram of

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

across all five years, thirty occupations, and six gender-age groups described above. Clearly, college educated workers are relatively more likely to use computers within occupations compared to high school educated workers. A similar pattern holds comparing across other education groups.

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 on average there is no clear difference in computer usage across genders within occupations (i.e. the histogram is roughly centered around zero). Hence, in order to account for the

[21] On average across the five years considered in the analysis, we measure $\pi_t(\lambda, \kappa, \omega) = 0$ for roughly 27% of the $(\lambda, \kappa, \omega)$ triplets. As a robustness check, in Appendix F.2 we drop age as a characteristic defining a labor group and redo our analysis with the resulting 10 labor groups. With only 10 labor groups, the share of measured allocations $\pi_t(\lambda, \kappa, \omega)$ that are equal to 0 is substantially smaller.

[22] The German *Qualification and Working Conditions* survey, used in e.g. DiNardo and Pischke (1997), helps mitigate the second and third concerns by providing data on worker usage of multiple types of equipment and, in 2006, the share of hours spent using computers. In Appendix D, we show using this more detailed German data similar patterns of comparative advantage—between computers and education groups and between computers and gender—as in the U.S. data. We do not accounting exercise for Germany because wage data is noisy in publicly available datasets, see e.g. Dustmann et al. (2009), and because the German *Qualification and Working Conditions* survey contains many fewer observations than the October Supplement (depending on the year, between roughly 10,700 and 21,150 observations after cleaning).

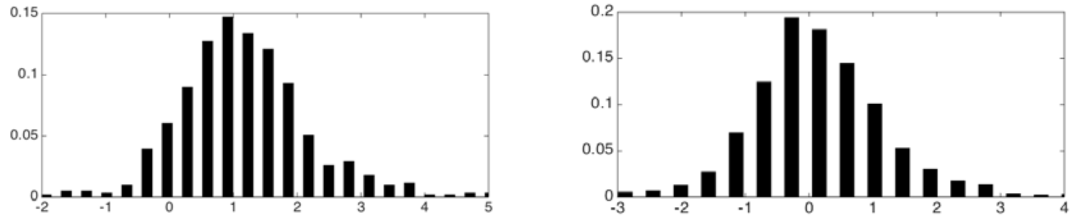

Figure 2: Computer relative to non-computer usage for college degree relative to high school degree workers (female relative to male workers) in the left (right) panel

fact that women use computers more than men at the aggregate level—see Table 1—women must have a comparative advantage in occupations in which computers have a comparative advantage.

We can similarly study the extent to which labor groups differ in their allocations across occupations conditional on computer usage and the extent to which computers differ in their allocations across occupations conditional on labor groups. For instance, using similar histograms we can show 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 computers are much more likely to be used in administrative support than in construction occupations, conditional on labor group. These comparisons provide an example of a more general relationship: women tend to be employed in occupations in which all labor groups are relatively more likely to use computers.

4.2 Measuring shocks

Here we describe our baseline procedure to measure shocks to labor composition, $\hat{L}(\lambda)/\hat{L}(\lambda_1)$, equipment productivity to the power 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 over time 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 our baseline procedure in Appendix C.1 and variations of this procedure in Appendices C.2 and C.3 (both of which yield very similar results).

Consider first our measure of changes in equipment productivity to the power θ. Equations (3) and (6) imply

$$\frac{\hat{q}(\kappa)^{\theta}}{\hat{q}(\kappa_1)^{\theta}} = \frac{\hat{\pi}(\lambda,\kappa,\omega)}{\hat{\pi}(\lambda,\kappa_1,\omega)} \tag{11}$$

for any (λ, ω) pair. Hence, if computer productivity rises relative to other equipment between t_0 and t_1, then the share of λ hours spent working with computers relative to other equipment in occupation ω will increase. 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 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. To construct changes in equipment productivity to the power θ, we combine the data on the right-hand side of equation (11) over (λ, ω) pairs following the procedure described in Appendix C.1.

Second, consider our measure of changes in occupation shifters. Equation (9) implies

$$\frac{\hat{a}(\omega)}{\hat{a}(\omega_1)} = \frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega_1)} \left(\frac{\hat{q}(\omega)}{\hat{q}(\omega_1)}\right)^{(1-\alpha)(\rho-1)}. \tag{12}$$

We construct the right-hand side of equation (12) as follows. Equations (3) and (6) provide a measure of changes in transformed occupation prices to the power θ between t_0 and t_1 for any (λ, κ) pair,

$$\frac{\hat{q}(\omega)^{\theta}}{\hat{q}(\omega_1)^{\theta}} = \frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa, \omega_1)}. \tag{13}$$

As described in Appendix C.1, we combine these (λ, κ) pair specific measures to obtain a unique measure of changes in transformed occupation prices to the power θ, $\hat{q}(\omega)^{\theta}/\hat{q}(\omega_1)^{\theta}$.[23] Intuitively, if the share of λ hours spent working with κ in occupation ω relative to in occupation ω_1 increased, then it must be that the transformed price of ω rose relative to ω_1. As above, it is important to condition on (λ, κ) when identifying changes in transformed occupation prices. Given values of α, ρ, and θ, we recover $(\hat{q}(\omega)/\hat{q}(\omega_1))^{(1-\alpha)(\rho-1)}$. Next, 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 (9). Note that if $\rho = 1$, changes in occupation shifters depend only on changes in the share of labor payments across occupations.

Finally, we measure changes in labor productivity as a residual to match changes in relative wages, expressing equation (7) as

$$\frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} = \frac{\hat{T}(\lambda)}{\hat{T}(\lambda_1)} \left(\frac{\hat{s}\,(\lambda)}{\hat{s}\,(\lambda_1)}\right)^{1/\theta}. \tag{14}$$

The variable $\hat{s}\,(\lambda)$ is a labor-group-specific weighted average of changes in equipment

[23] In calculating changes in relative wages in response to a subset of shocks, we solve for counterfactual changes in transformed occupation prices and allocations using equations (8) and (9).

productivity and transformed occupation prices, both raised to the power θ,

$$\hat{s}(\lambda) = \sum_{\kappa,\omega} \frac{\hat{q}(\omega)^{\theta}}{\hat{q}(\omega_1)^{\theta}} \frac{\hat{q}(\kappa)^{\theta}}{\hat{q}(\kappa_1)^{\theta}} \pi_{t_0}(\lambda,\kappa,\omega), \tag{15}$$

which we can construct using observed allocations and our measures of changes in equipment productivity to the power θ and transformed occupation prices to the power θ. From equation (14), we also require a value of θ to measure changes in labor productivity; we describe how we estimate θ below. Note that only our measures of changes in labor productivities are directly a function of the observed changes in relative wages that we will decompose below: given parameters α, ρ, and θ, our measures of wage changes have no effect on our measures of changes in labor composition or equipment productivity to the power 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

Baseline estimation. 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 U.S. 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.

The parameter θ determines the within-worker dispersion of productivity across occupation-equipment type pairs, and 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 (we conduct sensitivity analyses using a range of values of θ and ρ in Section 6.1). Equation (14) can be expressed as

$$\log \hat{w}(\lambda,t) = \varsigma_{\theta}(t) + \beta_{\theta} \log \hat{s}(\lambda,t) + \iota_{\theta}(\lambda,t), \tag{16}$$

where $\hat{x}(t)$ denotes the relative change in a variable x between any two periods t and $t' > t$.[24] We observe $\log \hat{w}(\lambda,t)$ in our data and construct $\log \hat{s}(\lambda,t)$ as indicated in equa-

[24]While we have argued—using equation (10)—that the change in average wages in response to given changes in transformed occupation prices and equipment productivities, $\hat{q}(\omega,t)$ and $\hat{q}(\kappa,t)$, does not depend—to a first-order approximation—on the value of θ, we will measure $\hat{q}(\omega,t)^{\theta}$ and $\hat{q}(\kappa,t)^{\theta}$. Hence,

tion (15). The parameter $\varsigma_{\theta}(t) \equiv \log \hat{q}(\omega_1, t)\, \hat{q}(\kappa_1, t)$ 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 (14), measuring changes in labor productivity requires a value of θ and, therefore, we treat $\iota_{\theta}(\lambda, t)$ as unobserved when estimating θ. Similarly, equation (9) can be expressed as

$$\log \hat{\zeta}(\omega, t) = \varsigma_{\rho}(t) + \beta_{\rho} \log \frac{\hat{q}(\omega, t)^{\theta}}{\hat{q}(\omega_1, t)^{\theta}} + \iota_{\rho}(\omega, t). \tag{17}$$

We directly observe $\log \hat{\zeta}(\omega, t)$ in the MORG CPS and measure $\log \hat{q}(\omega, t)^{\theta} / \hat{q}(\omega_1, t)^{\theta}$ following 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 (12), measuring changes in occupation shifters requires a value of ρ and, therefore, we treat $\iota_{\rho}(\omega, t)$ as unobserved when estimating ρ.

Equations (16) and (17) may be used to jointly identify θ and ρ. According to our model, however, using Nonlinear Least Squares (NLS) to estimate θ and ρ would yield biased estimates of these parameters. The reason is that the observed covariate in equation (16), $\log \hat{s}(\lambda, t)$, is predicted to be correlated with its error term, $\iota_{\theta}(\lambda, t)$, and the observed covariate in equation (17), $\log \hat{q}(\omega, t)^{\theta} / \hat{q}(\omega_1, t)^{\theta}$, is predicted to be correlated with its error term, $\iota_{\rho}(\omega, t)$.

Concerning the endogeneity of $\log \hat{s}(\lambda, t)$ in equation (16), note from equation (15) that this covariate 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 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 productiv-

the response of wages to *measured* shocks does depend, even to a first order, on the value of θ. This logic explains the intuition behind using equation (16) to identify θ.

ity to the power θ, $\hat{q}(\kappa,t)^{\theta}/\hat{q}(\kappa_1,t)^{\theta}$.[25] We use this instrument and equation (16) to build the following moment condition

$$\mathbb{E}_{\lambda,t}\left[\left(y_{\theta}(\lambda,t)-\frac{1}{\theta}x_{\theta}(\lambda,t)\right)\times z_{\theta}(\lambda,t)\right]=0, \tag{18}$$

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.[26] In order for the moment condition in equation 18 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 (18) 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.

Concerning the endogeneity of $\log\hat{q}(\omega,t)^{\theta}/\hat{q}(\omega_1,t)^{\theta}$ in equation (17), note that, 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 $\log\hat{q}(\omega,t)^{\theta}/\hat{q}(\omega_1,t)^{\theta}$,

[25] In 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.

[26] By projecting first on a set of year effects and using the residuals from this projection in the moment condition in equation (18), we simplify significantly the computational burden involved in estimating both the parameter of interest θ and the set of incidental parameters $\{\varsigma_{\theta}(t)\}_t$. 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 $\{\varsigma_{\theta}(t)\}_t$ and the parameter θ jointly.

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

which is an occupation-specific average of observed changes in equipment productivity to the power θ, $\hat{q}(\kappa,t)^\theta/\hat{q}(\kappa_1,t)^\theta$.[27] We use this instrument and equation (17) to build the following moment condition

$$\mathbb{E}_{\omega,t}\left[\left(y_\rho(\omega,t) - (1-\alpha)(1-\rho)\frac{1}{\theta}x_\rho(\omega,t)\right) \times z_\rho(\omega,t)\right] = 0, \tag{19}$$

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 19 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 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 (19) 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.[28]

We estimate θ and ρ using the sample analogue of the moment conditions in equations (18) and (19). 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. We report the point estimates and standard errors

[27] In 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.

[28] Theoretically, if the U.S. specializes in traded sectors that employ a large share of computer-intensive occupations, then a rise in trade between 1984 and 2003 might generate a demand shift towards computer-intensive occupations within traded sectors, inducing a potential correlation between occupation shifters and weighted changes in equipment productivity. In practice, however, we find that computer-intensive occupations do not grow faster relative to non-computer-intensive occupations in manufacturing than in non-manufacturing, suggesting that this theoretical concern is not a problem in our setting.

in the top row of Table 1: the resulting estimate of θ is 1.78 with a standard error of 0.29; the estimate of ρ is also 1.78 but with a slightly larger standard error of 0.35.

Parameter	Time Trend?	Estimate	(SE)
(θ, ρ)	NO	(1.78, 1.78)	(0.29, 0.35)
(θ, ρ)	YES	(1.13, 2.00)	(0.32, 0.71)

Table 2: Parameter estimates based on joint estimation

Alternative estimations of θ and ρ. Here we present estimates of θ and ρ that rely on alternative identification assumptions. In Section 6.1, we show how our results on the decomposition of observed wage changes are affected if we use these alternative estimates in our analyses.

First, we add as controls a labor-group-specific time trend in equation (16) and an occupation-specific time trend in equation (17). Allowing for these time trends relaxes the orthogonality restrictions imposed to derive the moment conditions used in our baseline analysis (equations 18 and 19).

In particular, 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 and 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 (18) 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 (18). 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 (16), we also additionally control for an occupation-specific time trend in equation (17). 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 (19) 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 that 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 GMM estimates of θ and ρ that result from adding as controls a labor-group-specific time trend in equation (16) and an occupation-specific time trend in equation (17) are, respectively 1.13, with a standard error of 0.32, and 2, with a standard error of 0.71, as displayed in row 2 of Table 2. 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 (16) actually imply a smaller role for changes in labor productivity (the residual), as we show in Section 6.1.

Second, whereas in our baseline analysis we derive moment conditions from a wage equation and an occupation-expenditure share equation expressed in time differences (see equations 16 and 17), in Appendix C.4 we estimate θ and ρ using versions of both equation (16) and equation (17) in levels. When doing so, we also expand the right-hand side variables in equation (16) with λ-specific fixed effects and the right-hand side variables in equation (17) with ω-specific fixed effects. The resulting GMM estimates are θ =1.57, with a standard error of 0.14, and $\rho = 3.27$, with a standard error of 1.34. The details of this alternative estimation procedure are contained in Appendix C.4.

Third, there is an alternative approach—based on Lagakos and Waugh (2013) and Hsieh et al. (2013)—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. When following this approach—see Appendix E—we obtain an estimate θ equal to 2.62.[29]

Specification tests. As discussed above, one implication of our model is that, if we had estimated the parameter vector (θ, ρ) using a NLS estimator, we would have obtained an upward biased estimate of θ and a downward biased estimate of ρ. In fact, when estimating the parameter vector (θ, ρ) using NLS, we obtain an estimate of $\theta = 2.61$, with a standard error of 0.57, and an estimate of $\rho = 0.21$, with a standard error of 0.45. The fact that the NLS estimate of θ is higher than its GMM 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 GMM counterpart is consistent with the prediction of our model that the error term $\iota_\rho (\omega, t)$ is negatively correlated with the covariate $\log \hat{q} (\omega, t)^\theta / \hat{q} (\omega_1, t)^\theta$.

Another implication of our model, described in Section 3, is that labor composition only affects wages indirectly through occupation prices. We test this prediction by including changes in labor supply as an additional explanatory variable in equation (16) and computing the two-stage least squares estimates of both β_θ and the coefficient on labor supply. This yields an estimate of $\theta = 1.84$, which is not statistically different from our baseline (which was $\theta = 1.78$). Moreover, we cannot reject at the 10% significance level the null hypothesis that the effect of changes in labor supply on changes in wages, conditional on the composite term $\log \hat{s}(\lambda, t)$, is equal to zero.

Relation to the literature. Equation (16) is related to two distinct approaches in the literature studying the impact of worker allocations (across either computers or occupations) on worker wages. The first approach regresses wages of different workers on a dummy for computer usage and worker characteristics; see e.g. Krueger (1993). In critiquing this

[29] In Appendix E we also also apply this alternative estimation approach to obtain estimates of θ that are λ-specific.

work, DiNardo and Pischke (1997) argue that the coefficient on computers in a regression of wages of different workers on a dummy for computer usage will combine two effects: *(a)* the causal effect of computers on wages and *(b)* a selection effect according to which workers of different groups have some unobserved characteristic that both impacts their wages and their likelihood of using computers; one source of this selection effect may be workers of different groups selecting into occupations in which computers are more likely to be used, as is the case in our model. If the unobserved worker characteristic that makes a worker group more likely to use computers has a direct positive impact on wages, the coefficient on computers would overestimate the causal effect of computers on wages. In our empirical approach, we never regress wage levels for different labor groups at a point in time on the intensity of computer usage by that labor group at the same point in time and, therefore, the critique in DiNardo and Pischke (1997) does not directly apply to our estimates. The key difference between equation (16) and the approach criticized by DiNardo and Pischke (1997) is that—instead of regressing changes in wages by labor group on changes in the intensity of computer usage by labor group—we project changes in wages by labor group on fixed-over-time measures of labor-group-specific computer usage interacted with our measures of changes in computer productivity (common to all labor groups). Therefore, endogenous changes in idiosyncratic labor productivity that correlate with changes in labor-group-specific computer usage will not bias our estimates.[30]

Equation (16) is most similar to the estimating equation that Acemoglu and Autor (2011) offer as a stylized example of how their assignment model might be brought to the data. Acemoglu and Autor (2011) suggest regressing changes in labor-group-specific wages on beginning-of-sample measures of the specialization of the different labor groups across abstract-intensive, routine-intensive, and manual-intensive occupations. A log-linearized version of equation (16)—our equation (10)—actually provides a micro-foundation for their regression model—their equation (40). Our approach, however, differs from the regression suggested by Acemoglu and Autor (2011) in that: *(a)* the included covariates measure specialization by worker types not only across three types of tasks but across thirty occupations and two types of equipment; *(b)* the included covariates are not just

[30] Two less central differences are as follows. A first difference between equation (16) and that criticized by DiNardo and Pischke (1997) is that equation (16) explores the impact of computerization on wage changes, instead of levels. As DiNardo and Pischke (1997) mention, this by itself is not likely to solve the bias due to self-selection: "if the returns to unobserved skills change, then differencing the data will not eliminate the effect of unobserved wage determinants that might be correlated with computer use." Another difference between equation (16) and that criticized by DiNardo and Pischke (1997) is that our regression is performed at the labor group level instead of at the worker level; again, this is unlikely to matter much if the main concern is self-selection of workers into using computers.

the beginning-of-sample measures of specialization but the interaction of these measures of specialization with measures of changes in occupation prices and equipment productivities. While explicitly controlling for these measured changes in aggregate occupation prices and equipment productivities is not important if the aim is to use equation (16) as evidence of the importance of comparative advantage for the evolution wages by skill group (as is the intention of Acemoglu and Autor (2011)), it is key for our purposes. The reason is that we use equation (16) to identify structural parameters of our model. If we had not explicitly controlled for observed measures of changes in equipment productivity and occupation prices and had just introduced measures of comparative advantage as covariates, the resulting estimates would be a combination of both the structural parameters of interest and the particular changes in equipment productivity and occupation prices that had taken place during our sample period. These reduced-form estimates would be a function of endogenous variables—the change in occupation prices depends on all four shocks—and, therefore, would not provide enough information to perform the wage decompositions and counterfactuals that we provide in Sections 5, 6, and 7. This is why it is crucial for our exercise that we first obtain measures of changes in equipment productivity and occupation prices on which we can condition when estimating equation (16).

5 Results

In this section we summarize our baseline closed economy results, quantifying the implications for the observed changes in relative wages in the U.S. between 1984 and 2003 of the changes in labor composition, occupation shifters, equipment productivity, and labor productivity that took place. We construct various measures of changes in between-group inequality, each of them aggregating wage changes across our thirty labor groups in different ways (e.g., the skill premium). As is standard, when doing so, both in the model and in the data, we construct *composition-adjusted* wage changes; that is, for each aggregated measure we average 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.[31] We also report the log change over each sub-period in our data.

Skill premium. We begin by decomposing changes in the skill premium over the full

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

sample period and over each sub-period. The first column in Table 3 reports the change in the data, which is also the change predicted by the 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 points between 1984 and 2003, with the largest increases occurring between 1984 and 1993. The subsequent four columns summarize the counterfactual change in the skill premium predicted by the model if only one of the exogenous shocks is considered (i.e. holding the other exogenous parameters at their t_0 level).

	Data	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
1984 - 1989	0.057	-0.031	0.026	0.052	0.009
1989 - 1993	0.064	-0.017	-0.009	0.045	0.046
1993 - 1997	0.037	-0.023	0.044	0.021	-0.005
1997 - 2003	-0.007	-0.043	-0.011	0.042	0.006
1984 - 2003	0.151	-0.114	0.049	0.159	0.056

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

Changes in labor composition decrease the skill premium over each sub-period in response to the large increase in the share of hours of more educated workers. The increase in hours worked by those with college degrees relative to those without of 47.4 log points between 1984 and 2003 decreases the skill premium by 11.4 log points. 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.

Changes in equipment productivity, i.e. computerization, account for roughly 60% of the sum of the demand-side forces pushing the skill premium upwards: $0.60 \simeq 0.159 / (0.049 + 0.159 + 0.056)$. Over sub-periods, changes in equipment productivity are particularly important in generating increases in the skill premium over the years in which the skill premium rose most dramatically: 1984-1989 and 1989-1993. These are precisely the years in which the overall share of workers using computers rose most rapidly; see Table 1. The intuition behind our result that computerization had a large impact on the skill premium is the following. We measure large growth in computer productivity using the procedure described in Section 4.2.[32] Computerization raises the skill premium

[32]This 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 procedure. The decline over time in the U.S. in the price of equipment relative to structures—see e.g. Greenwood et al.

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 Figure 2. Second, educated workers have a comparative advantage in occupations in which computers have a comparative advantage, 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 19% of the sum of the forces pushing the skill premium upwards over the full sample. This result is intuitively related to the expansion of certain occupations. As documented in Figure 1 and in Table 12 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, non-routine cognitive interpersonal, and socially perceptive tasks and a corresponding contraction in occupations intensive in routine manual and non-routine manual physical tasks, as defined using O*NET constructed task measures following Acemoglu and Autor (2011) (see Appendix F.1 for details). In our model, occupation shifters generate the largest movements in income shares across occupations, and these movements are biased towards occupations intensive in certain characteristics, e.g. social perceptiveness, as we show in Appendix F.1. However, with $\rho \neq 1$ all shocks other than changes in occupation shifters also affect income shares across occupations. Specifically, we find that computerization and labor composition play significant roles in explaining the observed relationship between occupations' and growth. For example, both computerization and changes in labor composition have contributed to the increase in the size of occupations intensive in non-routine cognitive analytical and non-routine cognitive interpersonal tasks as well as the decrease in the size of occupations intensive in routine manual and non-routine manual physical tasks, as we show in Appendix F.1.

Finally, labor productivity, the residual to match observed changes in relative wages, accounts for roughly 21% of the sum of the effects of the three demand-side mechanisms.

Gender gap. The average wage of men relative to women, the gender gap, declined

(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).

by 13.3 log points between 1984 and 2003. Table 4 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 log points between 1984 and 2003 increased the gender gap by 4.2 log points. 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 of the gender gap in the data.

	Data	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
1984 - 1989	-0.056	0.012	-0.009	-0.016	-0.044
1989 - 1993	-0.052	0.013	-0.035	-0.014	-0.016
1993 - 1997	-0.003	0.006	0.015	-0.005	-0.020
1997 - 2003	-0.021	0.012	-0.038	-0.012	0.019
1984 - 2003	-0.133	0.042	-0.067	-0.047	-0.061

Table 4: Decomposing changes in the log gender gap (the wage of men relative to women)

Changes in equipment productivity, i.e. computerization, account for roughly 27% of the sum of the forces decreasing the gender gap over the full sample in spite of the fact that women do not have a comparative advantage using computers. This results from the finding that women have a comparative advantage in the occupations in which computers have a comparative advantage, 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 38% of the sum of the forces decreasing the gender gap over the full sample. This is driven 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 Table 12 in the Appendix B for details.

Changes in labor productivity account for a sizable share, roughly 35%, of the impact of the demand-side forces affecting 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. (2013).

Five education groups. Table 5 decomposes changes in between-education-group wage inequality at a higher level of worker disaggregation than what the skill premium captures, studying changes in average wages across the five education groups considered in our analysis.

	Data	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
HS grad / HS dropout	0.037	-0.049	0.022	0.128	-0.060
Some college / HS dropout	0.074	-0.095	0.050	0.231	-0.110
College / HS dropout	0.174	-0.161	0.062	0.296	-0.022
Grad training / HS dropout	0.232	-0.189	0.104	0.310	0.009

Table 5: Decomposing changes in log relative wages across education groups between 1984 and 2003

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

Thirty disaggregated labor groups. One of the advantages of our framework is that we can solve for wage changes across a large number of labor groups. Whereas above we quantify the impact of shocks on measures of between-group inequality that aggregate across a number of labor groups, Table 6 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, in this table we show 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 each time. In order to perform this decomposition, we present three coefficients that arise from projecting each of the three sets of 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. This decomposition approach is analogous to that previously performed in Klenow and Rodriǵuez-Clare (1997).

The results show that, in the period 1984 to 2003, equipment productivity explains over 50% of the variance in the change in relative wages (implied by the combination of the 3 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%. Whereas equipment productivity is also the most important demand-side contributor in the periods 1984-1989 and 1997-2003, labor productivity is the most important force in 1989-1993 and occupation shifters are the most important force in 1993-1997.

	1984-2003	1984-1989	1989-1993	1993-1997	1997-2003
Equipment Productivity	52.93%	47.91%	33.15%	27.70%	51.38%
Occupation Shifter	23.80%	24.57%	10.54%	48.50%	14.92%
Labor Productivity	23.27%	27.52%	56.30%	23.80%	29.19%

Table 6: 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 $cov(x_i(\lambda), y(\lambda)) / var(y(\lambda))$.

6 Robustness and sensitivity analyses

In this section we consider three 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 5 in versions of our model that omit some of them. Finally, we allow for changes in comparative advantage over time.

6.1 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 7, 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 (16) and an occupation-specific time trend in equation (17). The third row uses our GMM estimates when using versions of both equation (16) and equation (17) in levels rather than in time differences. The final row uses the value of θ estimated from moments of the unconditional distribution of observed wages within each labor group λ and the value of ρ that arises from the moment condition in equation (19) taking as given the value of θ.

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 between-education-group inequality between 1984 and 2003. Alone, it accounts for between roughly 50% and 99% 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

(θ, ρ)	Skill premium Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Gender gap Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Baseline: (1.78, 1.78)	-0.114	0.049	0.159	0.056	0.042	-0.067	-0.047	-0.061
Time trends: (1.13, 2.00)	-0.159	0.066	0.221	0.021	0.058	-0.097	-0.063	-0.031
Levels: (1.57, 3.27)	-0.126	-0.018	0.272	0.022	0.046	-0.057	-0.092	-0.031
Wage distribution: (2.62, 2.15)	-0.084	0.037	0.117	0.080	0.031	-0.048	-0.036	-0.080

Table 7: Decomposing changes in the log skill premium and gender gap between 1984 and 2003 for alternative estimates of (θ, ρ).

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 16% and 49% 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.[33] Whereas the first row of Table 8 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 Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Gender gap Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Baseline: (1.78, 1.78)	-0.114	0.049	0.159	0.056	0.042	-0.067	-0.047	-0.061
(1.21, 1.78)	-0.135	0.021	0.236	0.027	0.050	-0.072	-0.075	-0.036
(2.35, 1.78)	-0.099	0.057	0.121	0.072	0.037	-0.061	-0.034	-0.074
(1.78, 1.09)	-0.155	0.142	0.111	0.051	0.057	-0.110	-0.019	-0.057
(1.78, 2.47)	-0.091	-0.005	0.188	0.058	0.034	-0.041	-0.064	-0.062

Table 8: Decomposing changes in the log skill premium and gender gap between 1984 and 2003 for extreme values of θ or ρ.

As is evident from rows two and three of Table 8, 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.

[33] This is intended to provide intuition for how our results vary with alternative values of θ and ρ. The numbers in Table 8 should not be interpreted as confidence intervals for our decomposition.

According to equation (14), 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 must reduce the impact of labor composition on relative wages, as is confirmed in rows four and five of Table 8. 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, as reported in rows four and five of Table 8.

On the other hand, the value of ρ impacts occupation shifters both through the magnitude of the measured shocks (see equation (12)) 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.

6.2 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 5.

Table 9 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	-0.114	0.049	0.159	0.056	0.042	-0.067	-0.047	-0.061
Only labor-equip. CA	0	0	0.240	-0.088	0	0	-0.105	-0.029
Only labor-occ. CA	-0.114	0.116	0	0.149	0.042	-0.056	0	-0.120

Table 9: 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 9 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 9 shows that abstracting from any comparative advantage at the level of equipment magnifies the importance of labor productivity 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.

6.3 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 (e.g. job flexibility, which women may value relatively more, see e.g. Goldin 2014), 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. In Appendix H 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. Table 10 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 5.

Changes in CA	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Labor-occ.	Labor-equip.	Equip.-occ.
None (baseline)	-0.114	0.049	0.159	0.056	-	-	-
Worker-occ. (case 1)	-0.114	-	0.159	-	0.069	-	-
Worker-equip. (case 2)	-0.114	0.046	-	-	-	0.223	-
Equip.-occ. (case 3)	-0.114	-	-	0.013	-	-	0.251

Table 10: Decomposing changes in the log skill premium between 1984 and 2003 allowing comparative advantage to evolve over time

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

7 International trade in equipment

In this section we extend our model to incorporate international trade in equipment, show theoretically how the degree of openness is reflected in what we have treated—in our closed-economy model—as exogenous primitive shocks to the cost of producing equipment, and quantify the impact of equipment trade on between-group inequality. Here we sketch the main elements of the extended model. Additional details are provided in Appendix I, where we also extend the model to allow for trade in occupations and in sectoral output.[34]

Setup. All variables are indexed by country, n, and we omit time subscripts for simplicity. We use Y to indicate output and D to indicate absorption; this distinction is required in the open economy but not in the closed economy. Absorption of equipment of type κ in country n is a CES aggregate of equipment sourced from all countries in the world,

$$D_n(\kappa) = \left(\sum_i D_{in}(\kappa)^{\frac{\eta(\kappa)-1}{\eta(\kappa)}} \right)^{\eta(\kappa)/(\eta(\kappa)-1)}, \tag{20}$$

where $D_{in}(\kappa)$ is absorption in country n of equipment κ sourced from country i, and $\eta(\kappa) > 1$ is the elasticity of substitution across source countries for equipment κ.[35] Trade is subject to iceberg transportation costs, where $d_{ni}(\kappa) \geq 1$ denotes the units of equipment κ output that must be shipped from origin country n in order for one unit to arrive in destination country i; we impose $d_{nn}(\kappa) = 1$ for all n. Output of equipment κ in country n satisfies

$$Y_n(\kappa) = \sum_i d_{ni}(\kappa) D_{ni}(\kappa). \tag{21}$$

Note that the previous expression limits to the closed economy version, $Y_n(\kappa) = D_{nn}(\kappa)$, when $d_{ni}(\kappa) \to \infty$ for all $n \neq i$. For simplicity, but without loss of generality for our results on relative wages, we abstract from trade in the consumption good. Therefore, the resource constraint for the final good still satisfies equation (2). For the exercises we consider below, we do not need to specify conditions on trade balance in each country.

The equations determining the allocations $\pi_n(\lambda, \kappa, \omega)$ and the average wage $w_n(\lambda)$, as well as the (non-traded) occupation market clearing conditions determining occupation

[34] See e.g. Grossman and Rossi-Hansberg (2008) for a theoretical analysis of occupation trade and inequality and Feenstra and Hanson (1999) for an empirical treatment of offshoring and relative wages. See Galle et al. (2015), Lee (2015), and Adao (2015) for an analysis of the impact of sectoral trade on between-group inequality.

[35] 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).

prices are the same as in the baseline model and are given by (3), (4) and (5), respectively, where every variable now includes an n country-subscript. Therefore, given equipment prices in country n, the equations to solve for allocations and wages in country n are the same as in the closed economy. In an open economy, however, we must distinguish between production prices and absorption prices (which include the price of imported equipment). Since domestic and imported equipment are inputs in production, absorption prices are the relevant equipment prices shaping $\pi_n(\lambda, \kappa, \omega)$ and $w_n(\lambda)$ in equations (3) and (4).

We denote by $p_{in}(\kappa)$ the price of country i's output of equipment κ in country n inclusive of trade costs, given by $p_{in}(\kappa) = p_{ii}(\kappa)\, d_{in}(\kappa)$, where $p_{ii}(\kappa)$ denotes the technological cost in terms of units of the final good of producing one unit of equipment κ in country i. We denote by $p_n(\kappa)$ the absorption price of equipment κ in country n (the relevant one in equations (3) and (4)), and it is given by

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

Note that the previous expression limits to the closed economy version, $p_n(\kappa) = p_{nn}(\kappa)$, when $d_{ni}(\kappa) \to \infty$ for all $n \neq i$. Taking the ratio between any two time periods t_0 and t_1, we have

$$\hat{p}_n(\kappa) = \left[\sum_i s_{in}(\kappa)\left(\hat{d}_{in}(\kappa)\,\hat{p}_{ii}(\kappa)\right)^{1-\eta(\kappa)}\right]^{\frac{1}{1-\eta(\kappa)}}, \tag{22}$$

where $s_{in}(\kappa)$ denotes the fraction of expenditures on equipment κ in country n purchased from country i at time t_0,

$$s_{in}(\kappa) = \frac{p_{in}(\kappa)\, D_{in}(\kappa)}{\sum_j p_{jn}(\kappa)\, D_{jn}(\kappa)}.$$

The change in country n's transformed equipment prices in equations (7) and (8) is given by $\hat{q}_n(\kappa) \equiv \hat{p}_n(\kappa)^{\frac{-\alpha}{1-\alpha}}\,\hat{T}_n(\kappa)$. Clearly, in this open economy model, changes in equipment absorption prices in country n depend on changes in equipment producer prices in all foreign countries and on changes in trade costs between country n and all of its trading partners.

Counterfactual exercise. Ideally, one might like to use the framework described above to quantify the impact of changes in trade costs, $\hat{d}_{in}(\kappa)$, and equipment producer prices in all foreign countries, $\hat{p}_{ii}(\kappa)$, between 1984 and 2003 on wage inequality in the U.S. In practice, measuring these changes abroad would require similar data to that described in Section 4.1 but for foreign countries (as well as equipment trade data). Instead, with the aim of quantifying the impact of international trade in equipment on country n, we consider a counterfactual exercise that does not require 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 of changes in primitives (i.e. worldwide technologies, labor compositions, and trade costs) between periods t_0 and t_1 on wages in country n, relative to the effects of the same changes in primitives if country n were a closed economy?

To understand this counterfactuals, define $w_n(\lambda; \Phi_t, \Phi_t^*, d_t)$ to be the average wage of labor group λ in country n given that country n fundamentals are Φ_t, fundamentals in the rest of the world are Φ_t^*, and the full matrix of world trade costs are d_t. Define $d_{n,t}^A$ to be an alternative matrix of world trade costs in which country n's international trade costs are infinite ($d_{in,t} = \infty$ for all $i \neq n$). Our counterfactual calculates

$$\frac{w_n\left(\lambda; \Phi_{t_1}, \Phi_{t_1}^*, d_{t_1}\right)}{w_n\left(\lambda; \Phi_{t_0}, \Phi_{t_0}^*, d_{t_0}\right)} \Bigg/ \frac{w_n\left(\lambda; \Phi_{t_1}, \Phi_{t_1}^*, d_{t_1}^A\right)}{w_n\left(\lambda; \Phi_{t_0}, \Phi_{t_0}^*, d_{t_0}^A\right)}$$

Defining the impact on the wage of group λ of moving country n to autarky at time t as $\hat{w}_{n,t}^A(\lambda) \equiv w_n\left(\lambda; \Phi_t, \Phi_t^*, d_{n,t}^A\right) / w_n\left(\lambda; \Phi_t, \Phi_t^*, d_{n,t}\right)$, our counterfactual can be expressed more simply as $\hat{w}_{n,t_0}^A(\lambda) \Big/ \hat{w}_{n,t_1}^A(\lambda)$. Because this counterfactual amounts to moving to autarky twice at different points in time, we only describe how to calculate the counterfactual change in wages when we move to autarky.

The system of equations to calculate changes in relative wages in some country n at time t_0 when this country moves to autarky corresponds to the same system of equations in the closed economy version of the model—(7), (8) and (9)—setting $\hat{T}_n(\lambda) = \hat{L}_n(\lambda) = \hat{a}_n(\omega) = 1$, and setting changes in transformed equipment prices to

$$\hat{q}_n(\kappa) = s_{nn}(\kappa)^{\frac{1}{1-\eta(\kappa)} \frac{-\alpha}{1-\alpha}} \tag{23}$$

where $s_{nn}(\kappa)$ (one minus the import share) is evaluated at time t_0. Expression (23) is obtained from the definition of transformed equipment prices, $\hat{q}_n(\kappa) = \hat{p}_n(\kappa)^{\frac{-\alpha}{1-\alpha}} \hat{T}_n(\kappa)$ with $\hat{T}_n(\kappa) = 1$, combined with the change in equipment κ's absorption price when moving to autarky, $\hat{p}_n(\kappa) = s_{nn}(\kappa)^{\frac{1}{1-\eta(\kappa)}}$ calculated using equation (22) setting $\hat{d}_{in}(\kappa) = \infty$ for all $i \neq n$ and $\hat{p}_{nn}(\kappa) = 1$.

The mapping between import shares at time t_0 and the corresponding closed-economy shocks is intuitive. If the import share of equipment type κ is relatively high and trade elasticities are common across equipment goods, then moving to autarky is equivalent to decreasing equipment κ productivity (or increasing its domestic cost of production) in the closed economy. Similarly, for given equipment import shares, if the trade elasticity is lower, then moving to autarky is equivalent to decreasing equipment productivity by more in the closed economy.

Results. We conduct our counterfactuals for the U.S. for our sample time period, 1984 and 2003. We need to assign values to $s_{nn}(\kappa)$ in 1984 and 2003 and to $\eta(\kappa)$. We calculate $s_{nn}(\kappa)$ for the U.S. as $s_{nn}(\kappa) = 1 - \frac{Import_{n,t}(\kappa)}{Production_{n,t}(\kappa) - Export_{n,t}(\kappa) + Import_{n,t}(\kappa)}$. We obtain Production, Export, and Import data using 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, κ_1, in the model to industry 30 (Office, Accounting, and Computing Machinery) and non-computer equipment, κ_2, in the model to industries 29-33 less 30 (Machinery and Equipment less Office, Accounting, and Computing Machinery) and 34-35 (Transport Equipment). We observe that $s_{nn}(\kappa_1)$ fell from 0.80 in 1984 to 0.26 in 2003, while $s_{nn}(\kappa_2)$ fell from 0.83 in 1984 to 0.65 in 2003; that is, the U.S. import share in computers rose significantly more than in non-computer equipment goods. We assume $\eta(\kappa_1) = \eta(\kappa_2)$ and consider values of the trade elasticity, $\eta(\kappa) - 1$, ranging between 1.5 and 5.5.

Table 11 reports the results from the counterfactual that quantifies how important was trade in equipment in generating relative wage changes between 1984 and 2003. The rise in the U.S. skill premium between 1984 and 2003 was 2 percentage points higher (with $\eta(\kappa) - 1 = 3.5$) than the counterfactual rise in the skill premium had the U.S. been in autarky over this time period for the same changes in worldwide primitives. The increase in the returns to college and graduate training (relative to a high school dropouts) accounted for by trade in equipment was roughly 3.5 percentage points, and the decrease in the gender gap was 0.5 percentage points. These numbers roughly double if we consider a lower trade elasticity of $\eta(\kappa) - 1 = 1.5$.

To summarize, the large rise in the import share of computers relative to non-computer equipment reduces the relative price of computers to non-computer equipment, and this change in the relative price of computers to non-computer equipment affects relative wages in exactly the same way as an increase in computer productivity in the closed economy.

8 Conclusions

In this paper we study how changes in workforce composition, occupation shifters, computerization, and labor productivity shape the evolution of between-group inequality

	Value of η $(\kappa) - 1$		
	1.5	3.5	5.5
Skill premium	0.050	0.021	0.013
HS grad / HS dropout	0.026	0.012	0.008
Some college / HS dropout	0.047	0.021	0.013
College / HS dropout	0.078	0.034	0.021
Grad training / HS dropout	0.080	0.034	0.022
Gender Gap	-0.011	-0.005	-0.003

Table 11: Impact of trade in equipment between 1984 and 2003 calculated for a range of equipment trade elasticities.

across many labor groups. We parameterize and estimate an assignment model to match observed factor allocations and wages in the United States between 1984 and 2003. We show that computerization alone accounts for the majority of the observed rise in between-education-group inequality over this period (e.g. 60% of the rise in the skill premium). The combination of observables—computerization and occupation shifters—explain roughly 80% 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.

Our framework remains tractable in spite of its high dimensionality, lending itself to a variety of extensions and applications. This tractability is attained by using strong parametric assumptions (e.g. Fréchet distribution of workers' productivity across equipment types and occupations), which would be interesting to relax in future work. We have extended our model to incorporate international trade in equipment and sector output as well as offshoring of occupations and have shown that changes in import and export shares in each of these markets shape what we treat, in our baseline closed-economy model, as exogenous primitive shocks. One challenge in bringing this general model to the data is the lack of available data on trade in occupation-specific export and import shares.

Finally, 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. (2012).

References

Acemoglu, Daron, "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, March 2002, *40* (1), 7–72.

_ **and David Autor**, *Skills, Tasks and Technologies: Implications for Employment and Earnings*, Vol. 4 of *Handbook of Labor Economics*, Elsevier,

Adao, Rodrigo, "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil," *mimeo*, 2015.

Autor, David, "Skills, Education, and the Rise of Earnings Inequality Among the 'Other 99 Percent'," *Science*, May 2014, *344* (6186), 843–851.

Autor, David H. and David Dorn, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, August 2013, *103* (5), 1553–97.

_ **, Frank Levy, and Richard J. Murnane**, "The Skill Content Of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, November 2003, *118* (4), 1279–1333.

_ **, Lawrence F. Katz, and Alan B. Krueger**, "Computing Inequality: Have Computers Changed The Labor Market?," *The Quarterly Journal of Economics*, November 1998, *113* (4), 1169–1213.

Beaudry, Paul and Ethan Lewis, "Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980-2000," *American Economic Journal: Applied Economics*, April 2014, *6* (2), 178–94.

Buera, Fracisco J., Joseph P. Kaboski, and Richard Rogerson, "Skill-Biased Structural Change and the Skill-Premium," *Working Paper*, 2015.

Burstein, Ariel and Jonathan Vogel, "Factor Prices and International Trade: A Unifying Perspective," March 2011, (16904).

_ **and** _ , "International Trade, Technology, and the Skill Premium," *mimeo Columbia University*, 2012.

_ **, Javier Cravino, and Jonathan Vogel**, "Importing Skill-Biased Technology," *American Economic Journal: Macroeconomics*, April 2013, *5* (2), 32–71.

Bustos, Paula, "The Impact of Trade Liberalization on Skill Upgrading Evidence from Argentina," April 2011, (559).

Card, David and Alan B Krueger, "School Quality and Black-White Relative Earnings: A Direct Assessment," *The Quarterly Journal of Economics*, February 1992, *107* (1), 151–200.

Costinot, Arnaud and Jonathan Vogel, "Matching and Inequality in the World Economy," *Journal of Political Economy*, 08 2010, *118* (4), 747–786.

_ **and** _ , "Beyond Ricardo: Assignment Models in International Trade," *Annual Review of Economics*, Forthcoming.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, July 2008, *55* (3), 511–540.

DiNardo, John E and Jorn-Steffen Pischke, "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?," *The Quarterly Journal of Economics*, February 1997, *112* (1), 291–303.

Dustmann, Christian, Johannes Ludsteck, and Uta Schonberg, "Revisiting the German Wage Structure," *The Quarterly Journal of Economics*, May 2009, *124* (2), 843–881.

Eaton, Jonathan and Samuel Kortum, "Trade in capital goods," *European Economic Review*, 2001, *45* (7), 1195–1235.

_ and _ , "Technology, Geography, and Trade," *Econometrica*, September 2002, *70* (5), 1741–1779.

Entorf, Horst, Michel Gollac, and Francis Kramarz, "New Technologies, Wages, and Worker Selection," *Journal of Labor Economics*, July 1999, *17* (3), 464–491.

Feenstra, Robert C. and Gordon H. Hanson, "The Impact Of Outsourcing And High-Technology Capital On Wages: Estimates For The United States, 1979-1990," *The Quarterly Journal of Economics*, August 1999, *114* (3), 907–940.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux, "Occupational Tasks and Changes in the Wage Structure," February 2011, (5542).

Galle, Simon, Andres Rodriguez-Clare, and Moises Yi, "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," *Working Paper*, 2015.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, April 2014, *104* (4), 1091–1119.

_ and Lawrence F. Katz, "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions," *Journal of Political Economy*, August 2002, *110* (4), 730–770.

Gordon, Robert J., *The Measurement of Durable Goods Prices*, National Bureau of Economic Research Research Monograph Series, 1990.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, June 1997, *87* (3), 342–62.

Grossman, Gene M. and Esteban Rossi-Hansberg, "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, December 2008, *98* (5), 1978–97.

Heckman, James J and Guilherme Sedlacek, "Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-selection in the Labor Market," *Journal of Political Economy*, December 1985, *93* (6), 1077–1125.

Helpman, Elhanan, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen J. Redding, "Trade and Inequality: From Theory to Estimation," NBER Working Papers 17991, National Bureau of Economic Research, Inc April 2012.

Hornstein, Andreas, Per Krusell, and Giovanni L. Violante, "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *American Economic Review*, December 2011, *101* (7), 2873–98.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and U.S. Economic Growth," NBER Working Papers 18693, National Bureau of Economic Research, Inc January 2013.

Huggett, Mark, Gustavo Ventura, and Amir Yaron, "Sources of Lifetime Inequality," *American Economic Review*, December 2011, *101* (7), 2923–54.

Kambourov, Gueorgui and Iourii Manovskii, "Occupational Mobility and Wage Inequality," *Review of Economic Studies*, 2009, *76* (2), 731–759.

_ and _ , "Occupational Specificity Of Human Capital," *International Economic Review*, 02 2009, *50* (1), 63–115.

Karabarbounis, Loukas and Brent Neiman, "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

Katz, Lawrence F and David H. Autor, "Changes in the Wage Structure and Earnings Inequality," in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 3 of *Handbook of Labor Economics*, Elsevier, 1999, chapter 26, pp. 1463–1555.

_ and Kevin M Murphy, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, February 1992, *107* (1), 35–78.

Klenow, Peter and Andrés Rodriǵuez-Clare, "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?," in "NBER Macroeconomics Annual 1997, Volume 12" NBER Chapters, National Bureau of Economic Research, Inc, september 1997, pp. 73–114.

Krueger, Alan B, "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989," *The Quarterly Journal of Economics*, February 1993, *108* (1), 33–60.

Krusell, Per, Lee E. Ohanian, Jose-Victor Rios-Rull, and Giovanni L. Violante, "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, September 2000, *68* (5), 1029–1054.

Lagakos, David and Michael E. Waugh, "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, April 2013, *103* (2), 948–80.

Lee, Donghoon and Kenneth I. Wolpin, "Accounting for wage and employment changes in the US from 1968-2000: A dynamic model of labor market equilibrium," *Journal of Econometrics*, May 2010, *156* (1), 68–85.

Lee, Eunhee, "Trade, Inequality, and the Endogenous Sorting of Heterogeneous Workers," *mimeo*, 2015.

Lemieux, Thomas, "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?," *American Economic Review*, June 2006, *96* (3), 461–498.

Oberfield, Ezra and Devesh Raval, "Micro Data and Macro Technology," NBER Working Papers 20452, National Bureau of Economic Research, Inc September 2014.

Parro, Fernando, "Capital-Skill Complementarity and the Skill Premium in a Quantitative Model of Trade," *American Economic Journal: Macroeconomics*, April 2013, *5* (2), 72–117.

Ramondo, Natalia and Andrés Rodriǵuez-Clare, "Trade, Multinational Production, and the Gains from Openness," *Journal of Political Economy*, 2013, *121* (2), 273 – 322.

Saez, Emmanuel, "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, January 2001, *68* (1), 205–29.

Yeaple, Stephen Ross, "A simple model of firm heterogeneity, international trade, and wages," *Journal of International Economics*, January 2005, *65* (1), 1–20.

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 $p_t(\omega)\,k^{\alpha}\left[T_t(\lambda,\kappa,\omega)\,l\right]^{1-\alpha}$ 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 (λ,κ,ω) 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 $p_t(\omega)^{\frac{1}{1-\alpha}}\left(\alpha p_t(\kappa)^{-1}\right)^{\frac{\alpha}{1-\alpha}} T_t(\lambda,\kappa,\omega)\,l$ and its cost as $\left[p_t(\kappa)^{\frac{-\alpha}{1-\alpha}}\left(\alpha p_t(\omega)\right)^{\frac{1}{1-\alpha}} T_t(\lambda,\kappa,\omega) + v_t(\lambda,\kappa,\omega)\right] l$. 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

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

which is independent of $\epsilon(z)$. The probability that a randomly sampled worker in group

λ chooses φ is

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

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

$$\begin{aligned}
\pi_t(\lambda,\varphi) &= -\int_0^\infty \exp\left(-\varepsilon^{-\theta(\lambda)} n_t(\lambda,\varphi)\right)(-\theta(\lambda))\,\varepsilon^{-1-\theta(\lambda)} d\varepsilon \\
&= \frac{1}{n_t(\lambda,\varphi)}\exp\left(-\varepsilon^{-\theta(\lambda)} n_t(\lambda,\varphi)\right)\Big|_{\varepsilon=0}^{\infty} \\
&= \frac{1}{n_t(\lambda,\varphi)}
\end{aligned}$$

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

$$\pi_t(\lambda,\varphi) = \frac{v_t(\lambda,\varphi)^{\theta(\lambda)}}{\sum_{\varphi'} v_t(\lambda,\varphi')^{\theta(\lambda)}} \tag{24}$$

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

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

where the first two equalities follow from both the assumption that $\epsilon(z)$ 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 $\epsilon(z)$ and the third equality follows from normalizing $\mathbb{E}[\epsilon(z)|z\in\mathcal{Z}_t(\lambda)]$ to be equal to 1.[36] 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

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

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

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

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

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

[36]This is a normalization because, for any value of $T_t'(\lambda)$ and $\mathbb{E}[\epsilon_t'(z)|z\in\mathcal{Z}_t(\lambda)]\neq 1$, we can always define an alternative $T_t(\lambda)$ and $\epsilon(z)$ such that $T_t(\lambda)=T_t'(\lambda)\times\mathbb{E}[\epsilon_t'(z)|z\in\mathcal{Z}_t(\lambda)]$ and $\mathbb{E}[\epsilon(z)|z\in\mathcal{Z}_t(\lambda)]=1$.

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

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

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

$$w_t(\lambda, \varphi) = \gamma(\lambda)\, v_t(\lambda, \varphi)\, \pi_t(\lambda, \varphi)^{-1/\theta(\lambda)}. \tag{25}$$

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 (24) into equation (25), we obtain

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

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 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 12, 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 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 six composite measures of O*NET Work Activities and Work Context Importance scales: (i) non-routine cognitive analytical, (ii) non-routine cognitive interpersonal, (iii) routine cognitive, (iv) routine manual, (v) non-routine manual physical, and (vi) social perceptiveness. We aggregate their measures up to our thirty occupations and standardize each to have mean zero and standard deviation one.

Occupations	College intensity 1984	College intensity 2003	Female intensity 1984	Female intensity 2003	Income share 1984	Income share 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	0.01
Agriculture and mining	0.05	0.06	0.10	0.16	0.01	0.00
Mechanics and repairers	0.04	0.07	0.03	0.04	0.05	0.01
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 12: 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.

C Measurement

Equations (7), (8), and (9) 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

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

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 (11) 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 (13) 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 (15). Given $\hat{s}(\lambda)$, we estimate θ using equation (16) 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 (11) and (13), 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 (14).

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 (11) and (13), 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 (12). 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 (11) and (13), 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 (8) as

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

where we define

$$\hat{q}(\lambda) \equiv \sum_{\kappa',\omega'} \hat{q}(\omega')^{\theta}\, \hat{q}(\kappa')^{\theta}\, \pi_{t_0}(\lambda,\kappa',\omega')$$

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 (11) and (13), 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 $\frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega)}$ and $\frac{\hat{\zeta}(\kappa)}{\hat{\zeta}(\kappa)}$. 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 $\frac{\hat{q}(\kappa)^{\theta}}{\hat{q}(1)^{\theta}}$, 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.

C.4 Estimation of θ and ρ using equations in levels

In Section 4.3, we describe approaches to estimate θ and ρ that derive moment conditions from equilibrium equations of our model expressed in time differences. One could also estimate θ and ρ using moment conditions derived from the same equilibrium equations expressed in levels instead of in time differences. In order to derive an estimating equation in levels that is analogous to equation 16, note that we can express wages as

$$w_t(\lambda) = \bar{\alpha}\gamma \times T_t(\lambda) \times S_t(\lambda)^{1/\theta} \tag{26}$$

where $S_t(\lambda)$ is a labor-group-specific average of equipment productivities and transformed occupation prices (both to the power θ),

$$S_t(\lambda) \equiv \sum_{\kappa,\omega} \left(T(\lambda,\kappa,\omega) q_t(\kappa) q_t(\omega)\right)^{\theta}. \tag{27}$$

To derive an estimating equation from expressions (26) and (27), we decompose $\log T_t(\lambda)$ into a labor group effect, a time effect, and labor-group-time-specific deviations and express equation (26) as

$$\log w_t(\lambda) = \varsigma_{\theta 2}(t) + \beta_{\theta 2}(\lambda) + \beta_{\theta} \log s_t(\lambda) + \iota_{\theta 2}(\lambda, t), \tag{28}$$

where $s_t(\lambda) \equiv S_t(\lambda) q_t(\omega_1)^{-\theta} q_t(\kappa_1)^{-\theta}$. In order to use this expression to estimate θ, we require an instrument for $\log s_t(\lambda)$ for the same reason we require an instrument for $\log \hat{s}(\lambda, t)$ in our baseline approach. We use a similar instrument,

$$\chi_{\theta 2}(\lambda, t) \equiv \log \sum_{\kappa} \left(\frac{q_t(\kappa)}{q_t(\kappa_1)}\right)^{\theta} \sum_{\omega} \pi_{1984}(\lambda, \kappa, \omega),$$

which is a labor-group-time-specific productivity shifter generated by the level rather than change in equipment productivity, $q_t(\kappa)^{\theta}/q_t(\kappa_1)^{\theta}$. A higher value of equipment κ productivity in period t raises the wage of group λ relatively more if a larger share of λ workers use equipment κ.

Similarly, we can derive an estimating equation in levels that is analogous to equation (17) by decomposing $\log a_t(\omega)$ into an occupation effect, a time effect, and an occupation-time-specific deviation and expressing equation (5) as

$$\log \zeta_t(\omega) = \varsigma_{\rho 2}(t) + \beta_{\rho 2}(\omega) + \beta_{\rho} \log \frac{q_t(\omega)^{\theta}}{q_t(\omega_1)^{\theta}} + \iota_{\rho 2}(\omega, t). \tag{29}$$

In order to use this expression to identify θ and ρ, we require an instrument for $q_t(\omega)^{\theta} / q_t(\omega_1)^{\theta}$ for the same reason we require an instrument for $\hat{q}(\omega, t)^{\theta} / \hat{q}(\omega_1, t)^{\theta}$ in our baseline approach. We use a similar instrument

$$\chi_{\rho 2}\left(\omega,t\right) \equiv \log \sum_{\lambda,\kappa} \frac{q_t(\kappa)^{\theta}}{q_t(\kappa_1)^{\theta}} \frac{L_{1984}\left(\lambda\right)\pi_{1984}\left(\lambda,\kappa,\omega\right)}{\sum_{\lambda',\kappa'} L_{1984}\left(\lambda'\right)\pi_{1984}\left(\lambda',\kappa',\omega\right)},$$

which is an occupation-time-specific productivity shifter generated by the level, rather than change, in $q_t(\kappa)^{\theta}/q_t\left(\kappa_1\right)^{\theta}$. A higher value of equipment κ productivity in period t lowers the price of occupation ω relatively more if a larger share of workers use equipment κ in occupation ω.

Using the estimating equations (28) and (29) and the instruments $\chi_{\theta 2}\left(\lambda,t\right)$ and $\chi_{\rho 2}\left(\omega,t\right)$ to estimate θ and ρ requires measures of (i) $q_t\left(\omega\right)^{\theta}/q_t\left(\omega_1\right)^{\theta}$ for all ω; (ii) $q_t\left(\kappa\right)^{\theta}/q_t\left(\kappa_1\right)^{\theta}$ for all κ, where $q_t\left(\kappa\right) \equiv p_t(\kappa)^{\frac{-\alpha}{1-\alpha}}T_t(\kappa)$; and (iii) $T\left(\lambda,\kappa,\omega\right)^{\theta}$ for all $\left(\lambda,\kappa,\omega\right)$. In order to construct these measures, note that equation (3) can be expressed as

$$\log \pi_t\left(\lambda,\kappa,\omega\right) = \log q_t\left(\kappa\right)^{\theta} + \log q_t\left(\omega\right)^{\theta} + \log q_t\left(\lambda\right) + \iota_t^L\left(\lambda,\kappa,\omega\right) \tag{30}$$

where

$$q_t\left(\lambda\right) \equiv \left(\sum_{\kappa',\omega'} T\left(\lambda,\kappa',\omega'\right) q_t\left(\kappa'\right)^{\theta} q_t\left(\omega'\right)^{\theta}\right)^{-1}$$

and

$$\iota_t^L\left(\lambda,\kappa,\omega\right) = \log T\left(\lambda,\kappa,\omega\right)^{\theta}.$$

Hence, regressing observed values of $\log \pi_t(\lambda,\kappa,\omega)$ on labor group, equipment, and occupation effects and exponentiating the resulting occupation and equipment fixed effects, we obtain estimates of $q_t\left(\omega\right)^{\theta}/q_t\left(\omega_1\right)^{\theta}$ and $q_t\left(\kappa\right)^{\theta}/q_t\left(\kappa_1\right)^{\theta}$. Finally, exponentiating the average across time of the resulting residual within $\left(\lambda,\kappa,\omega\right)$, we obtain an estimate of $T\left(\lambda,\kappa,\omega\right)^{\theta}$.

Given equations (28) and (29) and measures of both all their covariates and the instruments $\chi_{\theta 2}\left(\lambda,t\right)$ and $\chi_{\rho 2}\left(\omega,t\right)$, we estimate θ and ρ using a GMM estimator. We derive the two moments needed for identification of these two parameters by assuming that $\mathbb{E}_{\lambda,t}\left[\iota_{\theta 2}\left(\lambda,t\right) \times \chi_{\theta 2}\left(\lambda,t\right)\right] = 0$ and $\mathbb{E}_{\omega,t}\left[\iota_{\rho 2}\left(\omega,t\right) \times \chi_{\rho 2}\left(\omega,t\right)\right] = 0$. The resulting estimates are θ =1.58, with a standard error of 0.14, and $\rho = 3.27$, with a standard error of 1.34.

D Factor allocation in Germany

Constructing factor allocations, $\pi_t(\lambda, \kappa, \omega)$, using U.S. data from the October Supplement faces certain limitations. For example, (*i*) our view of computerization is narrow, (*ii*) our computer-use variable is zero-one at the individual level, (*iii*) we are not using any information on the allocation of non-computer equipment, and (*iv*) the computer use question was discontinued after 2003. Here, we use data on the allocation of German workers to different types of equipment in order to address possible concerns raised by limitations (*ii*) and (*iii*).

We use the 1986, 1992, 1999, and 2006 waves of the German *Qualification and Working Conditions* survey, which asks detailed questions about usage of different types of equipment (i.e. tools) at work. Specifically, respondents are asked which tool, out of many, they use most frequently at work. In 1986, 1992, and 1999, respondents are also asked whether or not they use each tool, regardless of whether it is the tool they use most frequently, whereas in 2006 respondents are asked about the share of time they spend using computers. The list of tools changes over time (discussed below) and is extensive. For instance, workers are asked if they use simple transportation tools such as wheel barrows or fork lifts, computers, and writing implements such as pencils. After cleaning, there are between 10,700 and 21,150 observations, depending on the year.

We group workers into twelve labor groups using three education groups (low education workers who do not have post-secondary education or an apprenticeship degree, medium education workers who have either post-secondary education or an apprenticeship degree, and high education workers who have a university degree), two age groups (20-39 and 40-up), and two genders. We consider twelve occupations. We drop workers who do not report using any tool most frequently. Because the list of tools changes over time, we allocate workers to computer usage (using the question about most used equipment type or the questions about whether a worker uses a type of equipment at all) as follows. In 1986 and 1992, we allocate a worker to computer usage if she reports using a computer terminal, computer-controlled medical instrument, electronic lists or forms, personal computer, computer, screen operated system, or CAD graphics systems. In 1999 we allocate a worker to computer usage if she reports using computerized control or measure tools, personal computers, computers with connection to intranet or internet, laptops, computers to control machines, or other computers. In 2006 we allocate a worker to computer usage if she reports using computerized control or measure tools, computers, personal computers, laptops, peripheries, or computers to control machines. We have considered a range of alternative allocations and obtained similar results.

According to our baseline definition, the share of workers for whom computers are the most-used tool rose from roughly 5% to 50% between 1986 and 2006. Clearly, no other equipment type reported in the data either grew or shrank at a similar pace.

Computer-education and computer-gender comparative advantage. We first use the question about the most used equipment type to study comparative advantage in Germany. This question helps address limitation (*iii*) in the U.S. data, since here we are using information on the allocation of non-computer equipment, both by dropping workers who do not report a most used equipment type and by including in the group of computer users only those workers who report that they use computers more than any other type of equipment. Specifically, we construct histograms in Figure 3 for Germany—analogous to Figure 2 for the U.S.—detailing education $\times$ computer and gender $\times$ computer comparative advantage. The left panel of Figure 3 shows that German workers with a high level of education have a strong comparative advantage using computers relative to workers with a low level of education, the middle panel shows that German workers with a high level of education have a mild comparative advantage using computers relative to workers with a medium level of education, and the right panel shows that women have no discernible comparative advantage with computers relative to men, where each of these patterns is identified within occupation and holding all other worker characteristics fixed. These patterns in Germany resemble the patterns we document in the U.S. in Figure 2.

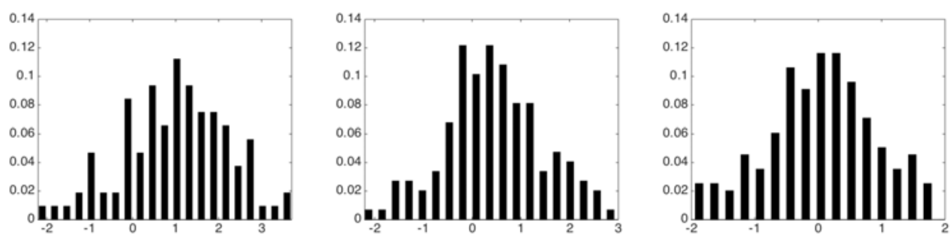

Figure 3: Computer relative to non-computer usage (constructed using the question about whether computers are the most used tool) for high relative to low education, high relative to medium education, and female relative to male workers, respectively, in Germany. Outliers have been truncated.

Second, we study the extent to which allocating workers to computers using the most-used or the used-at-all-question matters for measuring comparative advantage, since we only have access to the second type of question in the U.S. In the three years with available data (1986, 1992, and 1999) we construct allocations separately using these two questions. We then construct the share of hours worked with computers within each (λ, ω), i.e. $\pi_t^{Comp}(\lambda, \omega) \equiv \frac{\pi_t(\lambda, \kappa_{Comp}, \omega)}{\pi_t(\lambda, \kappa_{Comp}, \omega) + \pi_t(\lambda, \kappa_{Non-comp}, \omega)}$, separately using each type of question.

The correlation of $\pi_t^{Comp}(\lambda, \omega)$ constructed using the two different questions is high: 0.86, 0.78, and 0.55 in 1999, 1992, and 1986, respectively. To further understand the similarities and differences in measures of comparative advantage constructed using these questions, in Figure 4 we replicate the histograms in Figure 3 using the question on whether computers are used at all. The patterns of comparative advantage of education groups with computers, the left and middle panels of Figure 4, replicate the patterns in Figure 3: high education German workers have a strong and mild comparative advantage using computers relative to low and medium education German workers, respectively. However, constructing allocations using the used-at-all-question we measure a mild comparative advantage between men and computers in the right panel in Figure 4, unlike what we observe in Figure 3 in Germany or in Figure 2 in the U.S.

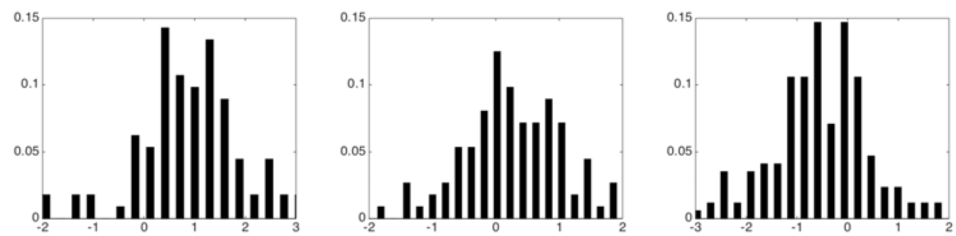

Figure 4: Computer relative to non-computer usage (constructed using questions about whether computers are used at all) for high relative to low education, high relative to medium education, and female relative to male workers, respectively, in Germany. Outliers have been truncated.

Finally, we use the question asked only in 2006 about the share of a worker's time spent using a computer. When constructing allocations using this question, we allocate the share of each worker's hours to computer or non-computer accordingly, whereas in our baseline approach we must allocate all of each worker's hours either to computers or non-computer equipment. Hence, this question helps address limitation (ii) in the U.S. data. Since figures like 2, 3, and 4 are noisy when constructed using a single year of data (using any question to determine allocations), here we focus instead on the correlation in the share of hours worked with computers within each (λ, ω), $\pi_t^{Comp}(\lambda, \omega)$, constructed using the most-used equipment type question and the share of hours worked with computers question. This correlation is above 0.9.

E Multivariate Fréchet

Recall that in our baseline approach, 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. Hence, in spite of the fact that each worker $z \in \mathcal{Z}_t(\lambda)$ draws $\varepsilon(z,\kappa,\omega)$ across (κ,ω) pairs from a Fréchet distribution with CDF $G(\varepsilon) = \exp\left(\varepsilon^{-\theta}\right)$, 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 approach to allow for correlation—see e.g. Ramondo and Rodríguez-Clare (2013) and Hsieh et al. (2013)—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 $\{\varepsilon(z,\kappa,\omega)\}_{\kappa,\omega}$ from a multivariate Fréchet distribution,

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

The parameter $\tilde{\theta}(\lambda) > 1$ governs the $\lambda-$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 $\lambda-$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.

Imposing a common $\theta(\lambda)$ across λ. It is straightforward to show that our baseline equations, parameterization strategy, and results hold exactly in the case in which a worker $z \in \mathcal{Z}_t(\lambda)$ supplies $\epsilon(z) \times \varepsilon(z,\kappa,\omega)$ efficiency units of labor if $\{\varepsilon(z,\kappa,\omega)\}_{\kappa,\omega}$ is drawn from a multivariate Fréchet distribution and if $\theta(\lambda)$ is constant across λ. Hence, given $\theta = \theta(\lambda)$ for all λ, all of our results are independent of the values of $\tilde{\theta}(\lambda)$ and $\nu(\lambda)$. Our baseline assumption that $\nu(\lambda) = 0$ is, therefore, without loss of generality under the common assumption, see e.g. Hsieh et al. (2013), that $\theta(\lambda)$ is constant across λ. Note that under the assumption that $\theta(\lambda)$ is common across λ, we can incorporate both a multivariate Fréchet distribution of $\{\varepsilon(z,\kappa,\omega)\}_{\kappa,\omega}$ and an arbitrary distribution of $\epsilon(z)$.

Allowing $\theta(\lambda)$ to vary across λ. In what follows, we describe the model, parameterization, and results allowing $\theta(\lambda)$ to vary across λ. In this section, we assume away $\epsilon(z)$ (that is, we assume that its distribution is degenerate).

Our baseline equilibrium equations in levels—(3), (4), and (5)—and in changes—(7), (8), and (9)—are unchanged except for $\theta(\lambda)$ replacing θ. The key distinction between our baseline and extended model is the parameterization.

When $\theta(\lambda)$ varies across λ, we parameterize $\theta(\lambda)$ first, following the approaches in Lagakos and Waugh (2013) and Hsieh et al. (2013). 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, therefore, use the empirical distribution of wages within each λ to estimate $\tilde{\theta}(\lambda)$, separately for each labor group λ, using maximum likelihood. Specifically, we jointly estimate the shape and scale parameter for each λ in each year t using maximum likelihood (MLE). Figure 5 plots the empirical and predicted wage distributions for all middle-aged work-

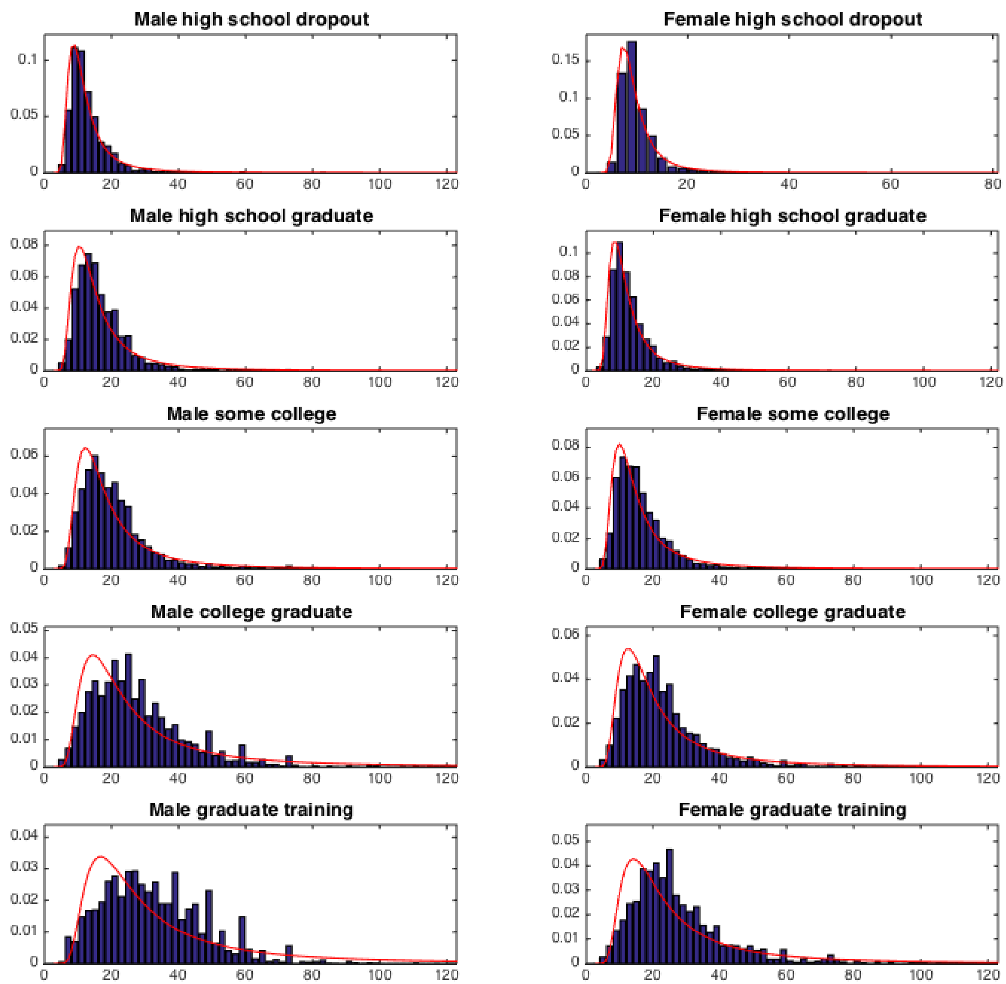

Figure 5: Empirical and predicted (Fréchet distribution estimated using maximum likelihood) wage distributions for all middle-aged labor groups in 2003

ers in 2003. We 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 (2013) 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 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 2.21 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—and for older than younger workers—averaging within each age group across education and gender we obtain estimates that fall monotonically with age from 2.96 to 2.37. The average across λ varies non-monotonically across years from a low of 2.56 to a high of 2.69. Finally, averaging across all groups and years yields an estimates of $\theta = 2.62$.

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 (11) and (13)

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

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 (12). Finally, we could also estimate ρ using the following variant of equation (17)

$$\log\hat{\zeta}(\omega) = \beta^{\rho}(t) + \beta_1^{\rho}\log\frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} + \iota^{\rho}(\omega),$$

where $\beta_1^{\rho} \equiv (1-\alpha)(1-\rho)$ is the coefficient of interest, using the same instrument as in our baseline approach.[37]

Table 13 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. Finally the final row reports results using distinct values of $\theta(\lambda)$ across each λ. The key message of Table 13 is that our results are robust. This is particularly true comparing between the second and the third rows of Table 13, in which the average value of $\theta(\lambda)$ is constant by construction.

[37]In practice, we will impose the same ρ as in our baseline in our exercises below.

	Skill premium				Gender gap			
	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Baseline	-0.114	0.049	0.159	0.056	0.042	-0.067	-0.047	-0.061
$\theta = 2.62$	-0.094	0.058	0.108	0.078	0.035	-0.058	-0.030	-0.079
$\theta(\lambda)$	-0.098	0.046	0.110	0.093	0.036	-0.061	-0.036	-0.070

Table 13: 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 Occupation characteristics

Here we use the standardized 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 14 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 six 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, non-routine cognitive interpersonal, and socially perceptive 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. Because $\rho \neq 1$, shocks other than occupation shifters generate changes in occupation income shares. We find that these other shocks play a significant role in accounting for the observed systematic evolution of occupation income shares over the years 1984-2003.

F.2 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

	Non-routine cogn. anlyt.	Non-routine cogn. inter.	Routine cogn.	Routine man.	Non-routine man. phys.	Social perc.
Data	0.173***	0.257***	-0.035	-0.273***	-0.214***	0.244***
Labor composition	0.038***	0.052***	0.010	-0.042***	-0.042***	0.032***
Occupation shifters	-0.044	0.061	-0.072	-0.093	-0.030	0.147***
Equipment prod.	0.068***	0.059***	0.042**	-0.063***	-0.073***	0.026
Labor productivity	0.002**	0.004***	-0.001	-0.004***	-0.005***	0.004***

Table 14: 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 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; Routine cogn. refers to Routine cognitive; Non-routine man. phys. refers to Non-routine manual physical; and Social perc. refers to Social perceptiveness

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

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) 27% to 12%. Because, in our baseline, we composition adjust the skill premium and the gender gap using gender, 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 log points, and the gender gap, -13.2 instead of -13.3 log points, between 1984 and 2003.

	Skill premium				Gender gap			
	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.	Labor comp.	Occ. shifters	Equip. prod.	Labor prod.
Baseline	-0.114	0.049	0.159	0.056	0.042	-0.067	-0.047	-0.061
10 groups baseline θ, ρ	-0.097	0.064	0.157	0.037	0.041	-0.068	-0.048	-0.055
10 groups re-estimate θ, ρ	-0.116	0.073	0.190	0.013	0.048	-0.081	-0.057	-0.040

Table 15: 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, $\theta = 1.78$ and $\rho = 1.78$, whereas in the other approach we re-estimate these parameters with 10 labor groups using our baseline estimation approach, yielding $\theta = 1.39$ and $\rho = 1.63$. We report results for both approaches and our baseline in Table 15. Our baseline results are robust to decreasing the number of labor groups.

G Average wage variation within a labor group

Our baseline model implies that the average wage of workers in group λ is the same across all equipment-occupation pairs. This implication is rejected by the data. In Section G.1, we argue that these differences in average wages across (κ, ω) do not drive our results. In Section G.2, we show that incorporating preference shifters for working in different occupations makes our model consistent with differences in average wages across occupations within a labor group and indicate how to decompose changes in wages in this case.

G.1 Between-within decomposition

Here, 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

$$\Delta \frac{w_t(\lambda)}{w_t} = \sum_{\omega} \Delta \frac{w_t(\lambda, \omega)}{w_t} \bar{\pi}_t(\lambda, \omega) + \sum_{\omega} \overline{\frac{w_t(\lambda, \omega)}{w_t}} \Delta \pi_t(\lambda, \omega), \tag{31}$$

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 (31) 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% for each labor group.[38] We conduct this decomposition using the MORG CPS data between 1984 and 2003 for each of 30 labor groups and find that the

[38] Of course, this does not mean that changes in occupation shifters do not drive changes in wages in our model.

median contribution across labor groups of the within component is above 86%. 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.

G.2 Compensating differentials

Here we extend our model to incorporate heterogeneity across labor groups in workers' preferences for working in each equipment-occupation pair. This simple extension implies that the average wage of workers in group λ varies across equipment-occupation pairs. We show how to use data on average wages across equipment-occupation pairs to identify the parameters of the extended model.

Environment and equilibrium. The indirect utility function of a worker $z \in \mathcal{Z}(\lambda)$ earning income $I_t(z)$ and employed in occupation ω with equipment κ is

$$U(z,\kappa,\omega) = I_t(z)\, u_t(\lambda,\kappa,\omega) \tag{32}$$

where $u_t(\lambda,\kappa,\omega) > 0$ is a time-varying preference shifter.[39] We have normalized the price index to one. We normalize $u_t(\lambda,\kappa_1,\omega_1) = 1$ for all λ and t. This model limits to our baseline model when $u_t(\lambda,\omega,\kappa) = 1$ for all t and (λ,κ,ω).

A occupation production unit hiring k units of equipment κ and l efficiency units of labor λ earns profits $p_t(\omega)\, k^{\alpha}\left[T_t(\lambda,\kappa,\omega)\, l\right]^{1-\alpha} - p_t(\kappa)\, k - v_t(\lambda,\kappa,\omega)\, l$. Conditional on positive entry in (λ,κ,ω), the profit maximizing choice of equipment quantity and the zero profit condition yield

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

Facing the wage profile $v_t(\lambda,\kappa,\omega)$, each worker $z \in \mathcal{Z}(\lambda)$ chooses (κ,ω) to maximize her indirect utility, $\varepsilon_t(z,\kappa,\omega)\, u_t(\lambda,\kappa,\omega)\, v_t(\lambda,\kappa,\omega)$.

In our extended model, preference parameters $u_t(\lambda,\kappa,\omega)$ and productivity parameters, $T_t(\lambda,\kappa,\omega)$, affect worker utility in the same way. Hence, they also affect worker allocation in the same way: the probability that a randomly sampled worker, $z \in \mathcal{Z}(\lambda)$,

[39] In this extended environment it is straightforward to allow for $u_t(\lambda,\kappa,\omega) = 0$, in which case no workers in group λ would choose (κ,ω) in period t.

uses equipment κ in occupation ω is

$$\pi_t(\lambda,\kappa,\omega) = \frac{\left[u_t(\lambda,\kappa,\omega)\, T_t(\lambda,\kappa,\omega) p_t(\kappa)^{\frac{-\alpha}{1-\alpha}}\, p_t(\omega)^{\frac{1}{1-\alpha}}\right]^{\theta(\lambda)}}{\sum_{\kappa',\omega'}\left[u_t(\lambda,\kappa',\omega')\, T_t(\lambda,\kappa',\omega') p_t(\kappa')^{\frac{-\alpha}{1-\alpha}}\, p_t(\omega')^{\frac{1}{1-\alpha}}\right]^{\theta(\lambda)}}. \tag{33}$$

On the other hand, preferences and productivities affect wages differently. The average wage of workers $z \in \mathcal{Z}(\lambda)$ teamed with equipment κ in occupation ω is now given by

$$w_t(\lambda,\kappa,\omega) = \frac{\gamma(\lambda)}{u_t(\lambda,\kappa,\omega)}\left(\sum_{\kappa',\omega'}\left[u_t(\lambda,\kappa',\omega')\, T_t(\lambda,\kappa',\omega') p_t(\kappa')^{\frac{-\alpha}{1-\alpha}}\, p_t(\omega')^{\frac{1}{1-\alpha}}\right]^{\theta(\lambda)}\right)^{1/\theta(\lambda)} \tag{34}$$

If $u_t(\lambda,\kappa,\omega) > u_t(\lambda,\kappa',\omega')$, then the average wage of group λ is lower in (κ,ω) than in (κ',ω') in period t.

The general equilibrium condition is identical to our baseline model and is given by equation (5), although total labor income in occupation ω is now given by

$$\zeta_t(\omega) \equiv \sum_{\lambda,\kappa} w_t(\lambda,\kappa,\omega)\, L_t(\lambda)\, \pi_t(\lambda,\kappa,\omega).$$

Parameterization. Here, we focus on measuring preference shifters and shocks under the restriction that $\theta(\lambda) = \theta$ for all λ, taking θ as given.

From equation (34), we have

$$\frac{w_t(\lambda,\kappa,\omega)}{w_t(\lambda,\kappa_1,\omega_1)} = \frac{1}{u_t(\lambda,\kappa,\omega)}. \tag{35}$$

Hence, we measure preference shifters directly from average wages.

Equations (6) and (33) give us,

$$\frac{\pi_t(\lambda,\kappa_2,\omega)}{\pi_t(\lambda,\kappa_1,\omega)} = \frac{u_t(\lambda,\kappa_2,\omega)^{\theta}}{u_t(\lambda,\kappa_1,\omega)^{\theta}}\frac{q_t(\kappa_2)^{\theta}}{q_t(\kappa_1)^{\theta}}$$

which, together with equation (35), gives us

$$\frac{q_t(\kappa_2)^{\theta}}{q_t(\kappa_1)^{\theta}} = \frac{\pi_t(\lambda,\kappa_2,\omega)}{\pi_t(\lambda,\kappa_1,\omega)}\frac{w_t(\lambda,\kappa_2,\omega)^{\theta}}{w_t(\lambda,\kappa_1,\omega)^{\theta}}$$

Hence, we obtain

$$\begin{aligned}\log \frac{\hat{q}(\kappa_2)^{\theta}}{\hat{q}_{\kappa}(\kappa_1)^{\theta}} &= \log \frac{\pi_{t_1}(\lambda,\kappa_2,\omega)}{\pi_{t_1}(\lambda,\kappa_1,\omega)} - \log \frac{\pi_{t_0}(\lambda,\kappa_2,\omega)}{\pi_{t_0}(\lambda,\kappa_1,\omega)} \\ &\quad + \theta \log \frac{w_{t_1}(\lambda,\kappa_2,\omega)}{w_{t_1}(\lambda,\kappa_1,\omega)} - \theta \log \frac{w_{t_0}(\lambda,\kappa_2,\omega)}{w_{t_0}(\lambda,\kappa_1,\omega)}\end{aligned}$$

We can then average over all (λ, ω) and then exponentiate, exactly as in our baseline, to obtain a measure of changes in equipment productivity (to the power θ). We obtain a measure of changes in transformed occupation prices to the power θ similarly and use this to measure changes in occupation shifters using equation (12), as in our baseline approach. Finally, we have

$$w_t(\lambda,\kappa_1,\omega_1) = T_t(\lambda)\,\gamma(\lambda)\left(\sum_{\kappa,\omega}\left[u_t(\lambda,\kappa,\omega)\,q_t(\kappa)\,q_t(\omega)\right]^{\theta}\right)^{1/\theta}$$

so that

$$\hat{w}(\lambda,\kappa_1,\omega_1) = \hat{T}(\lambda)\left(\sum_{\kappa,\omega}\left(\hat{u}(\lambda,\kappa,\omega)\,\hat{q}(\kappa)\,\hat{q}(\omega)\right)^{\theta}\pi_{t_0}(\lambda,\kappa,\omega)\right)^{1/\theta}$$

Hence, given measures of changes in transformed occupation prices (to the power θ), changes in equipment productivity (to the power θ), and changes in preference shifters as well as observed changes in wages and observed allocations in period t_0, we can measure changes in relative labor productivities using the previous expression for group λ relative to group λ_1.

H Evolving comparative advantage: Details

Here we study case 2 described in Section 6.3, where

$$T_t(\lambda,\kappa,\omega) \equiv T_t(\omega)\,T_t(\lambda,\kappa)\,T(\lambda,\kappa,\omega)\,. \tag{36}$$

Cases 1 and 3 are similar and available upon request.

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 (7) and (8) become

$$\hat{w}(\lambda) = \left(\sum_{\kappa,\omega} \pi_{t_0}(\lambda,\kappa,\omega) \left(\hat{q}(\lambda,\kappa)\hat{q}(\omega)\right)^{\theta(\lambda)} \right)^{1/\theta(\lambda)}, \tag{37}$$

$$\hat{\pi}(\lambda,\kappa,\omega) = \frac{\left(\hat{q}(\omega)\hat{q}(\lambda,\kappa)\right)^{\theta(\lambda)}}{\sum_{\kappa',\omega'} \left(\hat{q}(\omega')\hat{q}(\lambda,\kappa')\right)^{\theta(\lambda)} \pi_{t_0}(\lambda,\kappa',\omega')}, \tag{38}$$

whereas equation (9) remains unchanged. Expressing equation (37) in relative terms yields

$$\frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} = \frac{\hat{q}(\lambda,\kappa_1)}{\hat{q}(\lambda_1,\kappa_1)} \frac{\left(\sum_{\kappa,\omega} \pi_{t_0}(\lambda,\kappa,\omega) \left(\frac{\hat{q}(\lambda,\kappa)}{\hat{q}(\lambda,\kappa_1)} \frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} \right)^{\theta(\lambda)} \right)^{1/\theta(\lambda)}}{\left(\sum_{\kappa',\omega'} \pi_{t_0}(\lambda_1,\kappa',\omega') \left(\frac{\hat{q}(\lambda_1,\kappa')}{\hat{q}(\lambda_1,\kappa_1)} \frac{\hat{q}(\omega')}{\hat{q}(\omega_1)} \right)^{\theta(\lambda_1)} \right)^{1/\theta(\lambda_1)}}. \tag{39}$$

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 (36) 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 (39).

I Model with international trade: details

In Section 7 we extended our closed economy model to allow for international trade in equipment goods. We now additionally extend the model to allow for trade in occupations and in sectoral output. To do so, we must first include sectors in our model. We do so in a closed economy version of our model first. We then consider the fully extended open economy model. We show how to compute the two types of counterfactuals discussed in Section 7 (which do not require solving for the full world general equilibrium or estimating parameters in any country other than the U.S) in this model.

I.1 Sectors in a closed economy

Whereas in our baseline model the final good was produced using a CES combination of different occupations, here we assume that the final good is produced using a CES combination of different sectors (indexed by σ) with elasticity ρ_σ and that sectoral output is produced using a CES combination of different occupations with elasticity ρ. Specifically, the final good combines sectoral output, $Y_t(\sigma)$, according to a CES production function,

$$Y_t = \left(\sum_\sigma \mu_t(\sigma)^{1/\rho_\sigma} Y_t(\sigma)^{(\rho_\sigma - 1)/\rho_\sigma} \right)^{\rho_\sigma/(\rho_\sigma - 1)}. \tag{40}$$

Sectoral output is itself a CES combination of the output of different occupations,

$$Y_t(\sigma) = \left(\sum_\omega \mu_t(\omega, \sigma)^{1/\rho} Y_t(\omega, \sigma)^{(\rho - 1)/\rho} \right)^{\rho/(\rho - 1)}, \tag{41}$$

where $Y_t(\omega, \sigma) \geq 0$ denotes absorption (equal to output in a closed economy) of occupation ω in the production of sector σ, $\mu_t(\omega, \sigma) \geq 0$ is an exogenous demand shifter for occupation ω in sector σ, and $\rho > 0$ is the elasticity of substitution across occupations within each sector. Total absorption and output of occupation ω is equal to its demand across sectors, $Y_t(\omega) = \sum_\sigma Y_t(\omega, \sigma)$. 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.[40] Total output of occupation ω, $Y_n(\omega)$, is equal to the sum of output across all workers employed in ω.

Because the partial equilibrium is the same as in our baseline model, the equations in levels determining the allocations $\pi_t(\lambda, \kappa, \omega)$ and the average wage $w_t(\lambda)$ are the same as in the baseline model and are given by (3) and (4), respectively. That is, for given productivities and occupation prices, allocations and wages are the same as in our baseline model. The occupation market clearing conditions, which pin down occupation prices,

[40] Accordingly, 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 $\pi_t(\lambda, \kappa, \omega, \sigma) = 0$.

become

$$\sum_{\sigma} E_t(\omega,\sigma) = \frac{1}{1-\alpha}\zeta_t(\omega), \tag{42}$$

where $E_t(\omega,\sigma)$ denotes income (or expenditure) on occupation ω in sector σ,

$$E_t(\omega,\sigma) = \mu_t(\sigma)\,\mu_t(\omega,\sigma)\,p_t(\omega)^{1-\rho}\,p_t(\sigma)^{\rho-\rho_\sigma}\,E_t,$$

and $p_t(\sigma)$ denotes the price index of sector σ

$$p_t(\sigma) = \left(\sum_{\omega}\mu_t(\omega,\sigma)\,p_t(\omega)^{1-\rho}\right)^{\frac{1}{1-\rho}}. \tag{43}$$

We now provide the system of equations in changes, analogous to equations (7)-(9) with which to calculate wage changes that result from changes in the primitives between periods t_0 and t_1. The expressions for changes in wages and in allocations are given, as in the baseline model, by equations (7) and (8), respectively. The right hand side of the occupation market clearing condition in changes is the same as in the baseline model,

$$\hat{\zeta}(\omega) = \frac{1}{\zeta_{t_0}(\omega)}\sum_{\lambda,\kappa} w_{t_0}(\lambda)\,L_{t_0}(\lambda)\,\pi_{t_0}(\lambda,\kappa,\omega)\,\hat{w}(\lambda)\,\hat{L}(\lambda)\,\hat{\pi}(\lambda,\kappa,\omega). \tag{44}$$

The left hand side of the occupation market clearing condition in changes and the change in the sectoral price index are, respectively,

$$\hat{p}(\omega)^{1-\rho}\,\hat{E}\sum_{\sigma}\nu_{t_0}(\sigma|\omega)\,\hat{\mu}(\sigma)\,\hat{\mu}(\omega,\sigma)\,\hat{p}(\sigma)^{\rho-\rho_\sigma} = \hat{\zeta}(\omega) \tag{45}$$

and

$$\hat{p}(\sigma) = \left[\sum_{\omega}\nu_{t_0}(\omega|\sigma)\,\hat{\mu}(\omega,\sigma)\,\hat{p}(\omega)^{1-\rho}\right]^{\frac{1}{1-\rho}}. \tag{46}$$

Here, $\nu_t(\sigma|\omega) \equiv \frac{E_t(\omega,\sigma)}{\sum_{\sigma'} E_t(\omega,\sigma')}$ denotes the share of expenditure on occupation ω that occurs within sector σ and $\nu_t(\omega|\sigma) = \frac{E_t(\omega,\sigma)}{\sum_{\omega'} E_t(\omega',\sigma)}$ denotes the share of expenditure on occupation ω employed within sector σ. Defining, as in the baseline model, changes in transformed within-sector occupation shifters $\hat{a}(\omega,\sigma) = \hat{\mu}(\omega)\,\hat{T}(\omega)^{(1-\alpha)(\rho-1)}$ and changes in transformed occupation prices $\hat{q}(\omega) = \hat{p}(\omega)^{1/(1-\alpha)}\hat{T}(\omega)$ we can write (45) and (46) as

$$\hat{q}(\omega)^{(1-\alpha)(1-\rho)}\,\hat{E}\sum_{\sigma}\nu_{t_0}(\sigma|\omega)\,\hat{\mu}(\sigma)\,\frac{\hat{a}(\omega,\sigma)}{\hat{a}(\omega_1,\sigma)}\hat{a}(\omega_1,\sigma)\,\hat{p}(\sigma)^{\rho-\rho_\sigma} = \hat{\zeta}(\omega),$$

and

$$\hat{p}(\sigma) = \hat{a}(\omega_1, \sigma)^{\frac{1}{1-\rho}} \left[\sum_{\omega} \nu_{t_0}(\omega|\sigma) \frac{\hat{a}(\omega, \sigma)}{\hat{a}(\omega_1, \sigma)} \hat{q}(\omega)^{(1-\rho)(1-\alpha)} \right]^{\frac{1}{1-\rho}}.$$

Furthermore, defining changes in transformed sector prices $\hat{q}(\sigma) = \hat{p}(\sigma)\,\hat{a}(\omega_1, \sigma)^{\frac{-1}{1-\rho}}$ and changes in transformed sector shifters $\hat{a}(\sigma) = \hat{a}(\omega_1, \sigma)^{\frac{1-\rho_\sigma}{1-\rho}}\,\hat{\mu}(\sigma)$ we can re-write these two equations as

$$\hat{q}(\omega)^{(1-\alpha)(1-\rho)}\,\hat{E} \sum_{\sigma} \nu_{t_0}(\sigma|\omega)\,\hat{a}(\sigma) \frac{\hat{a}(\omega, \sigma)}{\hat{a}(\omega_1, \sigma)} \hat{q}(\sigma)^{\rho-\rho_\sigma} = \hat{\zeta}(\omega), \tag{47}$$

$$\hat{q}(\sigma) = \left[\sum_{\omega} \nu_{t_0}(\omega|\sigma) \frac{\hat{a}(\omega, \sigma)}{\hat{a}(\omega_1, \sigma)} \hat{q}(\omega)^{(1-\rho)(1-\alpha)} \right]^{\frac{1}{1-\rho}}. \tag{48}$$

Therefore, we can solve for changes in relative wages $\hat{w}(\lambda)/\hat{w}(\lambda_1)$, relative occupation prices $\hat{q}(\omega)/\hat{q}(\omega_1)$ and relative sectoral prices $\hat{q}(\sigma)/\hat{q}(\sigma_1)$ using equations (7), (8), (44), (47) and (48), given shocks $\hat{L}(\lambda)/\hat{L}(\lambda_1)$, $\hat{T}(\lambda)/\hat{T}(\lambda_1)$, $\hat{q}(\kappa)/\hat{q}(\kappa_1)$, $\hat{a}(\omega, \sigma)/\hat{a}(\omega_1, \sigma)$ and $\hat{a}(\sigma)/\hat{a}(\sigma_1)$.

Two special cases of our model are as follows. First, if $\rho = \rho_\sigma$, then $\hat{q}(\sigma)$ drops out from equation (47) and the model with sectors is equivalent to the baseline model where the occupation shifter, $\hat{a}(\omega)$, is replaced by $\sum_\sigma \nu_{t_0}(\sigma|\omega)\,\hat{a}(\sigma)\,(\hat{a}(\omega, \sigma)/\hat{a}(\omega_1, \sigma))$. Second, if $\nu_{t_0}(\omega|\sigma) = \nu_{t_0}(\omega)$ and $\hat{a}(\omega, \sigma)/\hat{a}(\omega_1, \sigma) = \hat{a}(\omega)/\hat{a}(\omega_1)$, 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 $\hat{a}(\omega)/\hat{a}(\omega_1)$.

We can measure $\hat{T}(\lambda)/\hat{T}(\lambda_1)$, $\hat{q}(\kappa)/\hat{q}(\kappa_1)$ and $\hat{q}(\omega)/\hat{q}(\omega_1)$ between any two time periods using the same procedure and data as in our baseline model. To measure changes in transformed within-sector occupation shifters and transformed sector shifters, and to construct $\nu_t(\sigma|\omega)$ and $\nu_t(\omega|\sigma)$ we need data on $E_t(\omega, \sigma)$ in t_0 and t_1. To measure within-sector occupation shifters, $\hat{a}(\omega, \sigma)/\hat{a}(\omega_1, \sigma)$, we start from the equilibrium relationship

$$\frac{\hat{E}(\omega, \sigma)}{\hat{E}(\omega_1, \sigma)} = \frac{\hat{a}(\omega, \sigma)}{\hat{a}(\omega_1, \sigma)} \left(\frac{\hat{p}(\omega)}{\hat{p}(\omega_1)} \right)^{1-\rho},$$

which can be re-expressed in terms of transformed shifters as

$$\frac{\hat{E}(\omega, \sigma)}{\hat{E}(\omega_1, \sigma)} = \frac{\hat{a}(\omega, \sigma)}{\hat{a}(\omega_1, \sigma)} \left(\frac{\hat{q}(\omega)}{\hat{q}(\omega_1)} \right)^{(1-\alpha)(1-\rho)}. \tag{49}$$

We use equation (49) to back out $\hat{a}(\omega,\sigma)/\hat{a}(\omega_1,\sigma)$. To estimate $\frac{\hat{\mu}(\sigma)}{\hat{\mu}(\sigma_1)}$, we start from the equilibrium relationship

$$\frac{\hat{E}(\sigma)}{\hat{E}(\sigma_1)} = \frac{\hat{\mu}(\sigma)}{\hat{\mu}(\sigma_1)} \left(\frac{\hat{p}(\sigma)}{\hat{p}(\sigma_1)}\right)^{1-\rho_\sigma}, \tag{50}$$

or, in terms of transformed variables,

$$\frac{\hat{E}(\sigma)}{\hat{E}(\sigma_1)} = \frac{\hat{a}(\sigma)}{\hat{a}(\sigma_1)} \left(\frac{\hat{q}(\sigma)}{\hat{q}(\sigma_1)}\right)^{1-\rho_\sigma}.$$

The previous expression and equation (48) yield

$$\frac{\hat{E}(\sigma)}{\hat{E}(\sigma_1)} = \frac{\hat{\mu}(\sigma)}{\hat{\mu}(\sigma_1)} \left(\frac{\sum_\omega \nu_{t_0}(\omega|\sigma)\frac{\hat{a}(\omega,\sigma)}{\hat{a}(\omega_1,\sigma)}\hat{q}(\omega)^{(1-\rho)(1-\alpha)}}{\sum_{\omega'} \nu_{t_0}(\omega'|\sigma_1)\frac{\hat{a}(\omega',\sigma_1)}{\hat{a}(\omega_1,\sigma_1)}\hat{q}(\omega')^{(1-\rho)(1-\alpha)}}\right)^{\frac{1-\rho_\sigma}{1-\rho}}. \tag{51}$$

We use equation (51) to back out $\hat{a}(\sigma)/\hat{a}(\sigma_1)$.

I.2 Full trade model

We now consider a world economy with many countries that trade occupational, sectoral, and equipment output subject to iceberg costs. As in Section 7, we omit time subscripts. The final good and sector production functions in country n, the open economy counterparts of equations (40) and (41), are given by

$$\begin{aligned} Y_n &= \left(\sum_\sigma \mu_n(\sigma)^{1/\rho_\sigma} D_n(\sigma)^{(\rho_\sigma-1)/\rho_\sigma}\right)^{\rho_\sigma/(\rho_\sigma-1)} \\ Y_n(\sigma) &= \left(\sum_\omega \mu_n(\omega,\sigma)^{1/\rho} D_n(\omega,\sigma)^{(\rho-1)/\rho}\right)^{\rho/(\rho-1)}. \end{aligned} \tag{52}$$

Total absorption of occupation ω is given by $D_n(\omega) = \sum_\sigma D_n(\omega,\sigma)$. Each country produces equipment, sector, and occupation output: $Y_n(\kappa)$, $Y_n(\sigma)$, and $Y_n(\omega)$, respectively. Absorption of each of these goods is a CES aggregate of these goods sourced from all countries in the world, as in expression (20), which now also applies to σ and ω (with elasticities $\eta(\sigma)$ and $\eta(\kappa)$, respectively). Each of these goods is subject to iceberg transportation costs, $d_{ni}(x) \geq 1$ (for $x = \kappa,\sigma,\omega$), and faces a resource constraint analogous to expression (21), with σ and ω in place of κ. The resource constraint for the final good still satisfies equation (2). For the exercises we consider below, we do not need to specify conditions on trade balance in each country.

The equations determining the allocations $\pi_n(\lambda,\kappa,\omega)$ and the average wage $w_n(\lambda)$ are the same as in the baseline model and are given by (3), (4) and (5), respectively. Therefore, given equipment prices in country n, the equations to solve for allocations and wages in country n are the same as in the closed economy. As discussed in Section 7, in an open economy we must distinguish between production prices and absorption prices. We use the analogous notation for prices that we used in Section 7. For example, we denote by $p_{in}(\omega)$ the price of country i's output of occupation ω in country n (inclusive of trade costs) and by $p_n(\omega)$ the absorption price of occupation ω in country n, given by

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

Production and absorption prices of individual sectors σ are defined analogously. 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). In the expressions for the change in the average wage and allocations in country n, given by (7) and (8) where all variables contain a subscript n, the relevant expression for the change in transformed equipment prices is $\hat{q}_n(\kappa) \equiv \hat{p}_n(\kappa)^{\frac{-\alpha}{1-\alpha}}\hat{T}_n(\kappa)$ and the expression for the change in transformed occupation prices is $\hat{q}_n(\omega) \equiv \hat{p}_{nn}(\omega)^{1/(1-\alpha)}\hat{T}_n(\omega)$. Occupation prices can be solved for using the open-economy versions of the occupation-market clearing condition (42) and the sectoral price index (the analog of equation 43), which we display below. In general, solving for the level of occupation prices requires solving for prices in the full world equilibrium (which are functions of worldwide technologies, endowments, and trade costs).

We consider the two type of counterfactual exercises described in Section 7, which do not require solving the full world general equilibrium or assigning parameters in any country other than n.

We derive the system of equations (analogous to equations (7), (8), and (9) in the closed economy model without sectors) that can be used to calculate changes in relative wages in some country n at time t_0 when this country moves to autarky ($d_{ni}(x)$ becomes infinite for all $i \neq n$ and all other primitives remain constant). We show that, moving to autarky, the equilibrium system of equations in an open economy is equivalent to the system that characterizes a closed economy with sectors, as detailed in Appendix I.1. Here, however, changes in equipment productivity, sector shifters, and within-sector occupation shifters are induced by moving to autarky.

Variables with the superscript A refer to counterfactual autarky values (holding all other parameters fixed—such as productivities, primitive occupation and sector shifters, and labor composition—at their time t_0 levels) and variables without the superscript A refer to factual values in period t_0. Variables with hats denote the ratio of the value of this variable in autarky relative to the value of this variable at time t_0: $\hat{y} = y^A_{t_0}/y_{t_0}$. For simplicity, in this section we omit time indices.

Wages and allocations in autarky relative to time t_0 (the counterpart to equations (7) and (8)) are

$$\hat{w}_n(\lambda) = \left\{ \sum_{\kappa,\omega} \left[\hat{p}_{nn}(\omega)^{\frac{1}{1-\alpha}} \hat{p}_n(\kappa)^{\frac{-\alpha}{1-\alpha}} \right]^{\theta(\lambda)} \pi_n(\lambda,\kappa,\omega) \right\}^{1/\theta(\lambda)}, \tag{53}$$

$$\hat{\pi}_n(\lambda,\kappa,\omega) = \frac{\left[\hat{p}_{nn}(\omega)^{\frac{1}{1-\alpha}} \hat{p}_n(\kappa)^{\frac{-\alpha}{1-\alpha}} \right]^{\theta(\lambda)}}{\sum_{\omega',\kappa'} \left[\hat{p}_{nn}(\omega')^{\frac{1}{1-\alpha}} \hat{p}_n(\kappa')^{\frac{-\alpha}{1-\alpha}} \right]^{\theta(\lambda)} \pi_n(\lambda,\kappa',\omega')}. \tag{54}$$

We now derive an expression for the occupation-market clearing in changes, the counterpart to equation (9) in our baseline closed-economy model (or equations (45) and (46) in the model with sectors).

The right-hand side of equation (9) remains unchanged, so we focus on the left-hand side only. The level of worldwide absorption expenditure on country n's produced occupation ω is $\sum_i E_{ni}(\omega)$, where $E_{ni}(\omega)$ denotes country i's absorption expenditure of occupation ω from country n,

$$E_{ni}(\omega) = p_{ni}(\omega) D_{ni}(\omega) = \left(\frac{p_{ni}(\omega)}{p_i(\omega)} \right)^{1-\eta(\omega)} E_i(\omega).$$

In autarky, $E_{ni}(\omega) = 0$ for $i \neq n$. Hence, the ratio of $\sum_i E_{ni}(\omega)$ between autarky and t_0 (the left hand side of equation (9)) is

$$\frac{\sum_i E^A_{ni}(\omega)}{\sum_i E_{ni}(\omega)} = \frac{E_{nn}(\omega)}{\sum_i E_{ni}(\omega)} \frac{E^A_{nn}(\omega)}{E_{nn}(\omega)} = f_{nn}(\omega) \left(\frac{\hat{p}_{nn}(\omega)}{\hat{p}_n(\omega)} \right)^{1-\eta(\omega)} \hat{E}_n(\omega),$$

where $f_{nn}(\omega) = E_{nn}(\omega) / \sum_i E_{ni}(\omega)$ denotes the share of domestic sales of occupation ω relative to its total sales (one minus the export share).

We now calculate an expression for $\hat{E}_n(\omega)$, the change in total expenditure on absorption of occupation ω in country n. The level of $E_n(\omega)$ is given by

$$E_n(\omega) = \sum_\sigma E_n(\omega, \sigma), \tag{55}$$

where $E_n(\omega, \sigma)$ denotes country n's absorption expenditures on occupation ω in sector σ and is given by

$$E_n(\omega, \sigma) = \mu_n(\omega, \sigma) \left(\frac{p_n(\omega)}{p_{nn}(\sigma)} \right)^{1-\rho} Y_n(\sigma) \, p_{nn}(\sigma), \tag{56}$$

where $p_{nn}(\sigma) = \left(\sum_\omega \mu_n(\omega, \sigma) \, p_n(\omega)^{1-\rho} \right)^{1/(1-\rho)}$. The value of sector σ production in country n is

$$Y_n(\sigma) \, p_{nn}(\sigma) = \sum_i Y_{ni}(\sigma) \, d_{in}(\sigma) \, p_{nn}(\sigma) = \sum_i E_{ni}(\sigma), \tag{57}$$

where $E_{ni}(\sigma)$ denotes expenditures on absorption in country i of country n's sector σ output, given by

$$E_{ni}(\sigma) \;\; = \;\; \mu_i(\sigma) \left(\frac{p_{ni}(\sigma)}{p_i(\sigma)} \right)^{1-\eta(\sigma)} \left(\frac{p_i(\sigma)}{p_i} \right)^{1-\rho_\sigma} E_i, \tag{58}$$

and E_i denotes total expenditures on the final good in country i. Combining equations (56), (57), and (58) yields

$$E_n(\omega, \sigma) = \mu_n(\omega, \sigma) \left(\frac{p_n(\omega)}{p_{nn}(\sigma)} \right)^{1-\rho} \sum_i \mu_i(\sigma) \, p_{ni}(\sigma)^{1-\eta(\sigma)} \, p_i(\sigma)^{\eta(\sigma)-\rho_\sigma} \, p_i^{1-\rho_\sigma} E_i.$$

The ratio of $E_n(\omega, \sigma)$ in autarky relative to its level at time t_0 is then

$$\hat{E}_n(\omega, \sigma) \;\; = \;\; \hat{p}_n(\omega)^{1-\rho} \, \hat{p}_{nn}(\sigma)^{\rho-\rho_\sigma} \left(\frac{\hat{p}_n(\sigma)}{\hat{p}_{nn}(\sigma)} \right)^{\eta(\sigma)-\rho_\sigma} f_{nn}(\sigma) \, \hat{E}_n,$$

where $f_{nn}(\sigma)$ denotes the share of domestic sales of sector σ relative to its total sales, defined analogously to $f_{nn}(\omega)$. Equation (55) therefore yields

$$\hat{E}_n(\omega) \;\; = \;\; \sum_\sigma \nu_n(\sigma|\omega) \, \hat{E}_n(\omega, \sigma) = \sum_\sigma \nu_n(\sigma|\omega) \, \hat{p}_n(\omega)^{1-\rho} \, \hat{p}_{nn}(\sigma)^{\rho-\rho_\sigma} \left(\frac{\hat{p}_n(\sigma)}{\hat{p}_{nn}(\sigma)} \right)^{\eta(\sigma)-\rho_\sigma} f_{nn}(\sigma) \, \hat{E}_n.$$

Combining these results, we have

$$\frac{\sum_i E^A_{ni}(\omega)}{\sum_i E_{ni}(\omega)} = f_{nn}(\omega)\,\hat{p}_n(\omega)^{1-\rho}\left(\frac{\hat{p}_{nn}(\omega)}{\hat{p}_n(\omega)}\right)^{1-\eta(\omega)} \sum_{\sigma} \nu_n(\sigma|\omega)\,\hat{p}_{nn}(\sigma)^{\rho-\rho_\sigma}\left(\frac{\hat{p}_n(\sigma)}{\hat{p}_{nn}(\sigma)}\right)^{\eta(\sigma)-\rho_\sigma} f_{nn}(\sigma)\,\hat{E}_n \tag{59}$$

where the change in the production sectoral price index is

$$\hat{p}_{nn}(\sigma) = \left(\sum_{\omega} \nu_n(\omega|\sigma)\,\hat{p}_n(\omega)^{1-\rho}\right)^{1/(1-\rho)}. \tag{60}$$

Finally, we calculate the differential change in absorption and production prices, $\hat{p}_n(\omega)/\hat{p}_{nn}(\omega)$, $\hat{p}_n(\kappa)/\hat{p}_{nn}(\kappa)$, and $\hat{p}_n(\sigma)/\hat{p}_{nn}(\sigma)$. When moving to autarky at time t_0, the change in import prices is infinite. The change in the absorption price of occupation ω, for example, is

$$\hat{p}_n(\omega) = \frac{p^A_{nn}(\omega)/p_{nn}(\omega)}{\left(\sum_i (p_{in}(\omega)/p_{nn}(\omega))^{1-\eta(\omega)}\right)^{\frac{1}{1-\eta(\omega)}}} = \frac{\hat{p}_{nn}(\omega)}{s_{nn}(\omega)^{\frac{1}{\eta(\omega)-1}}} \tag{61}$$

where $s_{nn}(\omega)$ denotes expenditure on domestic occupation ω relative to total expenditure on occupation ω in country n (one minus the import share),

$$s_{nn}(\omega) = \frac{p_{nn}(\omega)\,D_{nn}(\omega)}{\sum_i p_{in}(\omega)\,D_{in}(\omega)}.$$

The second equality in (61) uses the following relationship between the prices of domestic and imported goods

$$(p_{in}(\omega)/p_{nn}(\omega))^{1-\eta(\omega)} = (p_{in}(\omega)\,D_{in}(\omega))\,/\,(p_{nn}(\omega)\,D_{nn}(\omega))\,.$$

Similarly, changes in absorption prices of sector σ are

$$\hat{p}_n(\sigma) = \frac{\hat{p}_{nn}(\sigma)}{s_{nn}(\sigma)^{\frac{1}{\eta(\sigma)-1}}} \tag{62}$$

where $s_{nn}(\sigma)$ is defined analogously to $s_{nn}(\omega)$. Finally, the change in the absorption price of κ is simply

$$\hat{p}_n(\kappa) = s_{nn}(\kappa)^{\frac{1}{1-\eta(\kappa)}}, \tag{63}$$

as derived in Section 7, where $s_{nn}(\kappa)$ was defined analogously to $s_{nn}(\omega)$ and where have used the fact that $\hat{p}_{nn}(\kappa) = 1$ given our choice of numeraire.

We can substitute equation (63) directly into equations (53) and (54). Similarly, substituting (61) and (62) into (59) and (60) we have

$$\frac{\sum_i E^A_{ni}(\omega)}{\sum_i E_{ni}(\omega)} = \hat{p}_{nn}(\omega)^{1-\rho} \frac{f_{nn}(\omega)}{s_{nn}(\omega)^{\frac{\eta(\omega)-\rho}{\eta(\omega)-1}}} \sum_\sigma \nu_n(\sigma|\omega) \frac{f_{nn}(\sigma)}{s_{nn}(\sigma)^{\frac{\eta(\sigma)-\rho_\sigma}{\eta(\sigma)-1}}} \hat{p}_{nn}(\sigma)^{\rho-\rho_\sigma} \hat{E}_n$$

and

$$\hat{p}_{nn}(\sigma) = \left(\sum_\omega \nu_n(\omega|\sigma) s_{nn}(\omega)^{\frac{\rho-1}{\eta(\omega)-1}} \hat{p}_{nn}(\omega)^{1-\rho}\right)^{1/(1-\rho)}. \tag{64}$$

In sum, the system of equation to solve for changes in factor allocations and relative prices when moving to autarky is given by

$$\hat{\pi}_n(\lambda,\kappa,\omega) = \frac{\left[(\hat{p}_{nn}(\omega))^{\frac{1}{1-\alpha}} s_{nnt_0}(\kappa)^{\frac{\alpha}{(1-\alpha)(\eta(\kappa)-1)}}\right]^{\theta(\lambda)}}{\sum_{\omega',\kappa'} \left[(\hat{p}_{nn}(\omega'))^{\frac{1}{1-\alpha}} s_{nnt_0}(\kappa')^{\frac{\alpha}{(1-\alpha)(\eta(\kappa)-1)}}\right]^{\theta(\lambda)} \pi_{nt_0}(\lambda,\kappa',\omega')}$$

$$\hat{w}_n(\lambda) = \left\{\sum_{\kappa,\omega} \left[(\hat{p}_{nn}(\omega))^{\frac{1}{1-\alpha}} s_{nnt_0}(\kappa)^{\frac{\alpha}{(1-\alpha)(\eta(\kappa)-1)}}\right]^{\theta(\lambda)} \pi_{nt_0}(\lambda,\kappa,\omega)\right\}^{1/\theta(\lambda)}$$

$$\hat{p}_{nn}(\sigma) = \left(\sum_\omega \nu_{nt_0}(\omega|\sigma) s_{nnt_0}(\omega)^{\frac{\rho-1}{\eta(\omega)-1}} (\hat{p}_{nn}(\omega))^{1-\rho}\right)^{1/(1-\rho)}$$

and the occupation market clearing condition

$$\begin{aligned}(\hat{p}_{nn}(\omega))^{1-\rho} \frac{f_{nnt_0}(\omega)}{s_{nnt_0}(\omega)^{\frac{\eta(\omega)-\rho}{\eta(\omega)-1}}} \sum_\sigma \nu_n(\sigma|\omega) \frac{f_{nnt_0}(\sigma)}{s_{nnt_0}(\sigma)^{\frac{\eta(\sigma)-\rho_\sigma}{\eta(\sigma)-1}}} (\hat{p}_{nn}(\sigma))^{\rho-\rho_\sigma} \quad \hat{E} \\ = \frac{1}{\zeta_{t_0}(\omega)} \sum_{\lambda,\kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda,\kappa,\omega) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda,\kappa,\omega)\end{aligned}$$

All variables in the previous four equations that are indexed by t_0 represent either their observed level or are constructed based on estimates. Note that this system of equations corresponds to the system of equations in the closed economy version of the model with sectors, where within-sector occupation shifters, $\hat{a}_n(\omega,\sigma)$, and between-sector shifters,

$\hat{a}_n(\sigma)$, are equal to

$$\begin{aligned} \hat{a}_n(\omega, \sigma) &\equiv f_{nnt_0}(\omega) s_{nnt_0}(\omega)^{\frac{\rho-\eta(\omega)}{\eta(\omega)-1}}, \quad (65) \\ \hat{a}_n(\sigma) &\equiv f_{nnt_0}(\sigma) s_{nnt_0}(\sigma)^{\frac{\rho\sigma-\eta(\sigma)}{\eta(\sigma)-1}}, \quad (66) \end{aligned}$$

and changes in equipment costs are equal to

$$\hat{q}_n(\kappa) = s_{nnt_0}(\kappa)^{\frac{1}{\eta(\kappa)-1}\frac{\alpha}{1-\alpha}}.$$

The mapping between import and export shares at time t_0 and the corresponding closed-economy shocks is intuitive. First, as discussed in Section 7, if the import share of equipment type κ is relatively high and trade elasticities are common across equipment goods, then moving to autarky is equivalent to decreasing equipment κ productivity in the closed economy. Second, if sector σ has a relatively low export share and/or a high import share relative to sector σ' then, under mild parametric restrictions, moving to autarky is equivalent to increasing the sector shifter for σ relative to σ' in the closed economy. Finally, if occupation ω has a relatively low export share and/or a high import share relative to occupation ω' then, under mild parametric restrictions, moving to autarky is equivalent to increasing the within-sector occupation shifter for ω relative to ω' in the closed economy.[41]

Factor content of trade. In a special case of the model, the impact of moving to autarky on between-group inequality is captured by measures of the factor content of trade. It is useful to study this special case to understand the role of model assumptions and parameter values.

In a general accounting framework—see Burstein and Vogel (2011) and Burstein and Vogel (2012) for details—we can express the wage in country n of a given factor as

$$w_n(\lambda) = \frac{1}{L_n(\lambda)} \times \left(1 - \frac{FCT_n(\lambda)}{L_n(\lambda)}\right)^{-1} \times FPD_n(\lambda)$$

where $FCT_n(\lambda)$ and $FPD_n(\lambda)$ are the factor content of trade and the factor payments for domestic absorption of factor λ, both defined below. Letting $X_{nj}(\sigma)$ and $X_{nj}(\omega)$ denote

[41]Our extended model does not capture some of the mechanisms that have been studied in the literature linking international trade to between-group inequality. For example, as studied in Yeaple (2005), Bustos (2011), and Burstein and Vogel (2012), trade liberalization increases the measured skill bias of technology by reallocating resources from less to more skill-intensive firms within industries and/or inducing firms to increase their skill intensity. Extending the model to capture these mechanisms and mapping them into the components of our decomposition is a promising area for future work.

sales from country n to country j in sector σ and occupation ω and $\alpha_{nj}(\lambda,\sigma)$ and $\alpha_{nj}(\lambda,\omega)$ denote the share of this revenue that is paid to factor λ, we define $FCT_n(\lambda)$ and $FPD_n(\lambda)$ as follows. The factor content of trade is

$$FCT_n(\lambda) \equiv \frac{1}{w_n(\lambda)} \sum_{z\in\Sigma,\Omega} \sum_{j\neq n} \alpha_{nj}(\lambda,z) X_{nj}(z) - \frac{1}{w_n(\lambda)} \sum_{z\in\Sigma,\Omega} \sum_{j\neq n} \alpha_{nn}(\lambda,z) X_{jn}(z),$$

where the first term is the factor content of exports, which is the amount of λ embodied in country n's exports across all sectors and occupations, and the second term is the factor content of imports, which is the counterfactual amount of λ that would be used in country n to produce for itself the value of imports across all sectors and occupations. The factor payments for domestic absorption is

$$FPD_n(\lambda) \equiv \sum_{z\in\Sigma,\Omega} \sum_{j} \alpha_{nn}(\lambda,z) X_{jn}(z),$$

which is the counterfactual payments to factor λ if domestic absorption in all sectors and occupations were produced domestically.

In the trade model in this Appendix, if we impose no trade in equipment goods, $d_{nj}(\kappa) = \infty$ for all $n \neq j$; Cobb-Douglas sector and aggregate production functions, $\rho_\sigma = \rho = 1$; and $\theta \to 1$, then only the factor content of trade is impacted by foreign shocks (e.g., moving to autarky). In this special case, the impact on the average wage of λ of moving to autarky at period t is simply

$$\frac{\hat{w}^A_{n,t}(\lambda)}{\hat{w}^A_{n,t}(\lambda')} = \left(1 - \frac{FCT_{n,t}(\lambda')}{L_{n,t}(\lambda')}\right) \Big/ \left(1 - \frac{FCT_{n,t}(\lambda)}{L_{n,t}(\lambda)}\right)$$

since neither factor supply nor the factor payments for domestic absorption of λ relative to λ' are affected by moving to autarky.

Here we describe the role of each of the three restrictions above. If there is no trade in equipment goods, then moving to autarky affects wages in country n only through sector shifters and within-sector occupation shifters. If $\rho_\sigma = \rho = 1$, then the share of expenditure in country n on each sector and each occupation is fixed and determined purely by demand parameters $\mu_n(\sigma)$ and $\mu_n(\omega,\sigma)$, which are unaffected by moving to autarky. Finally, if $\theta \to 1$, then the share of revenue earned in each sector and occupation that is paid to each factor is also a function of parameters that is unaffected by moving to autarky. Hence, the factor payments for domestic absorption for any one factor relative to any other are pinned down by domestic parameters and are unaffected by foreign shocks.

Of course, when $\theta > 1$, the share of revenue in each occupation and sector paid to λ is a function of foreign shocks. Moreover, when $\rho \neq 1$ or $\rho_\sigma \neq 1$, the share of country n's income that accrues to each occupation or sector is affected by foreign shocks. In this case, the factor content of trade is not a sufficient statistic for calculating the impact of foreign shocks on domestic wages.