

Accounting for Changes in Between-Group Inequality^{*}

Ariel Burstein
UCLA

Eduardo Morales
Princeton University

Jonathan Vogel
Columbia University

January 29, 2015

Abstract

We provide an assignment model to decompose changes in between-group wage inequality into changes in the composition of the workforce, the productivity/demand for tasks, computerization, and labor productivity. The model incorporates comparative advantage between many groups of workers, many types of equipment, and many tasks and yet may be parameterized and estimated in a transparent manner. Our identification of parameters, measurement of shocks, and the equilibrium equation determining wages are all very similar to what have been used in previous reduced-form analyses. We use U.S. data on the allocation of workers to occupations and computer usage as well as changes in average wages across worker groups between 1984 and 2003 to parameterize our model. We find that computerization and changes in task productivity/demand, which are both measured without directly using data on changes in wages, jointly explain the majority of the rise in the skill premium and more disaggregated measures of between-education group inequality as well as roughly half of the rise in the relative wage of women over this time period. We show how to link the strength of these two forces to changes in the extent of international trade.

^{*}We thank Treb Allen, Costas Arkolakis, David Autor, Gadi Barlevy, Lorenzo Caliendo, Davin Chor, Arnaud Costinot, Jonathan Eaton, Pablo Fajgelbaum, Gene Grossman, and Nancy Stokey for helpful comments and David Autor and David Dorn for making their code available. Vogel thanks the Princeton International Economics Section for their support.

1 Introduction

The last few decades in the United States have witnessed pronounced changes in relative average wages across groups of workers with different observable characteristics (*between-group inequality*). For example, the relative wages of more educated to less educated workers and of women relative to men have increased substantially. What explains these and other observed changes in between-group inequality?¹

A voluminous literature has emerged—following [Katz and Murphy \(1992\)](#)—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 to a number of changes in the economic environment. Prominent amongst these changes are computerization (or a reduction in the price of equipment more generally)—see e.g. [Krueger \(1993\)](#), [Krusell et al. \(2000\)](#), [Acemoglu \(2002a\)](#), and [Autor and Dorn \(2013\)](#)—and changes in relative productivity or demand across occupations or sectors, driven by structural transformation, offshoring, and international trade—see e.g. [Berman et al. \(1994\)](#) and [Autor et al. \(2003\)](#). Related to the first hypothesis, Table 1 shows that between 1984 and 2003 computer use

		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

rose dramatically and that computers are used more intensively by educated workers and women.² Consistent with the second hypothesis, Figure 1 shows that over the same time period education- and female-intensive occupations grew relatively quickly.

In this paper we provide an assignment model with many groups of workers, many types of equipment, and many tasks, building on the work of [Eaton and Kortum \(2002\)](#), [Lagakos and Waugh \(2013\)](#), and [Hsieh et al. \(2013\)](#). Changes in relative wages across worker groups are shaped by shocks to (i) the composition of labor supply across groups,

¹The 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.”

²We describe our data sources in depth in Section 3.1.



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)

(ii) a composite of task demand and productivity, which we refer to as “task shifters,” (iii) a composite of equipment cost and productivity, which we refer to as “equipment productivity,” and (iv) a composite of other factors affecting the relative productivity of worker groups, independent of the equipment workers use and tasks in which they are employed,³ which we refer to as “labor productivity.” The model’s aggregate implications for relative wages nest those of workhorse macro models of between-group inequality, e.g. [Katz and Murphy \(1992\)](#) and [Krusell et al. \(2000\)](#). In spite of its high dimensionality—in our baseline empirics we use 10 education-gender groups, 2 types of equipment, and 30 occupations—the model can be parametrized and estimated in a transparent manner. We use the model to perform aggregate counterfactuals to quantify the impact on between-group inequality of the four shocks above.

In our model, the impact of changes in the economic environment on between-group inequality is shaped by comparative advantage between worker groups, equipment types, and tasks. Consider, for example, the potential impact of computerization on a 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 tasks, as we show is the case in the data for more educated workers. Computerization increases the relative wage of such groups. Second, a labor group may have a comparative advantage in the tasks in which computers have a comparative advantage, in which case this group would be allocated disproportionately to tasks in which all workers are relatively likely to use computers, as we show is the case in the data for female workers. Depending on the value of the elasticity of substitution between tasks, com-

³These factors include, for example, discrimination and the quality of education and health systems; see e.g. [Card and Krueger \(1992\)](#) and [Goldin and Katz \(2002\)](#).

puterization may increase or decrease relative wages of such groups and may increase or decrease employment in computer-intensive tasks.⁴ Therefore, measuring comparative advantage between worker groups, equipment types, and tasks is a key ingredient in our quantitative analysis.

Given data on the allocation of workers to equipment type-task pairs and the share of labor payments across tasks, we measure comparative advantage and shocks to task shifters and equipment productivity. In contrast, we measure changes in labor productivity as a residual to match changes in observed average wages across labor groups. We measure worker allocation using the October Current Population Survey (CPS) Computer Use Supplement, which provides information for five years (1984, 1989, 1993, 1999, and 2003) on whether or not a worker has direct or hands on use of a computer at work—be it a personal computer, laptop, mini computer, or mainframe—in addition to information on worker characteristics, hours worked, and occupation; see, e.g., [Krueger \(1993\)](#) and [Autor et al. \(1998\)](#) for previous studies using the October Supplement. This data is not without limitations, as we discuss in more depth in the paper: 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; and it does not detail the share of each worker’s time at work spent using computers.

We find that the combined effects of changes in task shifters and equipment productivity explain the majority of the rise in both aggregated (e.g., the skill premium) and disaggregated (e.g., the relative wage of workers with graduate training relative to high school dropouts) measures of between-education-group inequality between 1984 and 2003. Because we measure each of these (skill-biased) forces using data on changes in worker allocation and without directly using data on changes in average wages, this finding contrasts with previous results that attribute the majority of the rise in the skill premium to unobservable skill-biased technical change, see e.g. [Lee and Wolpin \(2010\)](#), or do not allow for trend growth in skill-biased technical change, see e.g. [Krusell et al. \(2000\)](#) and its discussion in [Acemoglu \(2002b\)](#). On the other hand, we find that changes in labor productivity, the residual, explain roughly half whereas computerization explains only a very small share of the fall in the gender gap. To demonstrate the importance of combining all three forces in a unified framework, we show how our results change if we turn off different sources of comparative advantage.

In spite of the strong assumptions in our structural model, we show that our identi-

⁴Our model is flexible enough so that computerization may increase the relative wage of workers who are relatively productive using computers and may reduce the relative wage of workers employed in tasks in which computers are particularly productive, as described by, e.g., [Autor et al. \(1998\)](#) and [Autor et al. \(2003\)](#).

fication of parameters, measurement of shocks, and equation determining wages are all very similar to what have been used, separately, in previous reduced-form analyses and, therefore, shed new light on and complement this work. We estimate the key elasticity in our model using an estimating equation almost identical to that introduced in [Katz and Murphy \(1992\)](#). In a limiting case of our model, we measure task shifters as changes in occupation shares of labor income exactly as in the literature using shift-share analyses, e.g. [Autor et al. \(1998\)](#). Moreover, our model generates an equation for wage changes in response to changes in task and equipment prices that provides a structural interpretation for the regression that [Acemoglu and Autor \(2011\)](#) offer as a stylized example of how their model might be brought to the data.

Our exercise is intended to shed light on classes of mechanisms through which changes in the economic environment lead to changes in between-group inequality. In our baseline model we treat shocks to task shifters, equipment productivity, and labor productivity as exogenous. Through the lens of our analysis, more primitive mechanisms—including, e.g., international trade, offshoring, structural change, or directed technical change—may play a central role driving changes in between-education inequality in the U.S. by changing the relative importance of occupations—see e.g. [Grossman and Rossi-Hansberg \(2008\)](#), [Autor et al. \(2003\)](#), and [Garicano and Rossi-Hansberg \(2006\)](#) for work on offshoring, automation of routine tasks, and hierarchies, respectively—or changing the relative prices or stocks of different types of equipment—see e.g. [Krusell et al. \(2000\)](#) and [Acemoglu \(2002a\)](#). In Section 6 we take a first theoretical step in this direction and extend our model to incorporate international trade in equipment and sector output as well as task offshoring. We show that changes in import and export shares in each of these markets shapes what we treated, in our baseline closed-economy model, as exogenous primitive shocks to equipment productivity, sector shifters, and within-sector occupation shifters, respectively.

Our paper’s approach is most similar to [Hsieh et al. \(2013\)](#). We follow [Hsieh et al. \(2013\)](#) and [Lagakos and Waugh \(2013\)](#), who build on [Eaton and Kortum \(2002\)](#) and [Dekle et al. \(2008\)](#), and use an assignment model of the labor market parameterized with a Fréchet distribution. Relative to [Hsieh et al. \(2013\)](#), we introduce equipment, quantify the role of computerization for between-group inequality, and analyze changes in between-education group inequality. We also introduce international trade in equipment, sectoral output, and tasks (e.g. offshoring) into our assignment model of the labor market, operationalizing the theoretical insights of [Costinot and Vogel \(2010\)](#) regarding the impact of international trade on inequality in a high-dimensional environment. We show that one can use a similar approach to that introduced by [Dekle et al. \(2008\)](#) in a single-factor

trade model—i.e. replacing a large number of unknown parameters with observable allocations in an initial equilibrium—in closed and open-economy many-factor assignment models to quantify the impact of various shocks on relative wages; [Burstein et al. \(2013\)](#), [Parro \(2013\)](#), and [Costinot and Rodriguez-Clare \(2014\)](#) use similar approaches in models with two labor groups.

Our paper’s objective is most similar to [Lee and Wolpin \(2010\)](#). They use a model of endogenous human capital accumulation in a dynamic framework 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 adopting an assignment approach, our framework is sufficiently tractable to allow for a greater degree of disaggregation (e.g. 30 occupations) and to exploit the detailed observed allocation of workers to equipment type-task pairs, reducing substantially the role of changes in 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, measuring it in each period directly from the data. Extending our model using standard assumptions to endogenize education and labor participation would give rise to the same equilibrium equations determining factor allocations and wages, given labor composition. Hence, our measures of shocks—to task shifters, equipment productivity, and labor productivity—and our estimates of model parameters would remain unchanged. In our counterfactual exercises, we fix labor composition to isolate the direct effect of individual shocks to task shifters, equipment productivity, and labor productivity on labor demand and wages.

Our paper is organized as follows. We provide our framework, characterize its equilibrium, and describe its mechanisms in [Section 2](#). We parameterize the model in [Section 3](#), describe our baseline closed-economy results in [Section 4](#), and consider various robustness exercises and sensitivity analyses in [Section 5](#). Finally, we extend the model to incorporate sectors and international trade in [Section 6](#) and conclude in [Section 7](#).

2 Model

In this section we introduce our baseline version of the model, characterize the equilibrium, and show how to decompose observed changes in relative average wages between any two periods into four channels. Finally, we provide intuition for how each channel operates.

2.1 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 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 tasks, indexed by ω ; in our quantitative analysis we map these tasks to occupations.⁵

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

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

where $\rho > 0$ is the elasticity of substitution across tasks, $Y_t(\omega) \geq 0$ is the endogenous output of task ω , and $\mu_t(\omega) \geq 0$ is an exogenous demand shifter for task ω .⁶ The final good is used to produce consumption, C_t , and equipment, $Y_t(\kappa)$,⁷ according to the resource constraint

$$Y_t = C_t + \sum_{\kappa} q_t(\kappa) Y_t(\kappa), \quad (2)$$

where $q_t(\kappa)$ denotes the cost of a unit of equipment κ in terms of units of the final good.⁸

Task production occurs within perfectly competitive production units. A unit hiring k units of equipment type κ and l efficiency units of labor group λ produces $k^{\alpha} [T_t(\lambda, \kappa, \omega) l]^{1-\alpha}$ units of output, where α denotes the output elasticity of equipment in each task and

⁵In 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 \(2009b\)](#), empirically 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 theoretically we would have to model dynamic worker optimization. We leave these considerations for future work.

⁶We show in Section 6.1 that we can disaggregate $\mu_t(\omega)$ further into sector shifters and within-sector task shifters. We show in Section 6.2 how changes in the extent of international trade/offshoring in sectoral output and task output generate endogenous changes in these sector shifters and a within-sector task shifters. For now, however, we combine sector and within-sector task shifters and treat them as exogenous.

⁷We assume that the sets of equipment types and tasks (as well as labor groups) are disjoint. Hence, the domain of a function such as $Y_t(\cdot)$ may be the union of these sets.

⁸Here we have assumed 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 results hold comparing across two balanced growth paths in which the real interest rate and the growth rate of relative productivity are both constant over time. We show in Section 6 how changes in the extent of international trade in equipment generates endogenous changes in $q_t(\kappa)$. For now, however, we treat the cost of producing equipment as exogenous.

$T_t(\lambda, \kappa, \omega)$ denotes labor λ 's productivity in task ω when using equipment κ .⁹ Comparative advantage (between labor and equipment) is defined as follows: λ' has a comparative advantage (relative to λ) using equipment κ' (relative to κ) in task ω if $T_t(\lambda', \kappa', \omega) / T_t(\lambda', \kappa, \omega) \geq T_t(\lambda, \kappa', \omega) / T_t(\lambda, \kappa, \omega)$. Labor-task and equipment-task comparative advantage are defined symmetrically.

A worker $z \in \mathcal{Z}_t(\lambda)$ supplies $\varepsilon(z, \kappa, \omega)$ efficiency units of labor if teamed with equipment κ in task ω . For each worker $z \in \mathcal{Z}_t(\lambda)$, the vector $\varepsilon(z)$ —which contains one $\varepsilon(z, \kappa, \omega)$ for each (κ, ω) pair—is drawn 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 λ -specific dispersion of efficiency units across (κ, ω) pairs; a higher value of $\tilde{\theta}(\lambda)$ decreases this dispersion. The parameter $0 \leq \nu(\lambda) \leq 1$ governs the λ -specific correlation of each worker's efficiency units across (κ, ω) pairs; a higher value of $\nu(\lambda)$ increases this correlation.¹⁰ We define $\theta(\lambda) \equiv \tilde{\theta}(\lambda) / (1 - \nu(\lambda))$. The assumption that efficiency units are distributed Fréchet is made first for analytical tractability—e.g., it implies that the average wage of a worker group is a CES function of prices and productivities—and second because it provides a reasonable approximation of the observed distribution of individual wages—see e.g. [Saez \(2001\)](#) and Figure 4 in Appendix A.

Total output of task ω , $Y_t(\omega)$, is the sum of output across all units producing task ω using any labor group λ and equipment type κ . All markets are perfectly competitive and all factors are freely mobile.

Relation to alternative frameworks. Whereas our framework imposes strong restrictions on task 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.¹¹

⁹We restrict α to be common across ω because we do not have the data to assign a different value of $\alpha(\omega)$ to each ω .

¹⁰In allowing for correlation across draws for each worker, we follow the approach of [Ramondo and Rodríguez-Clare \(2013\)](#) and [Hsieh et al. \(2013\)](#). See [Lagakos and Waugh \(2013\)](#) for a related approach.

¹¹The 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 task. 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 task and zero in the other

2.2 Equilibrium

We characterize the competitive equilibrium, first in partial equilibrium—taking task and equipment prices as given—and then in general equilibrium.

Partial equilibrium. A task 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 $p_t(\kappa) k + W_t(\lambda, \kappa, \omega) l$, where $W_t(\lambda, \kappa, \omega)$ is the wage per efficiency unit of labor λ when teamed with equipment κ in task ω and where $p_t(\omega)$ and $p_t(\kappa)$ are the prices of task ω output and of equipment type κ . The profit maximizing choice of equipment quantity and the zero profit condition—due to costless entry of production units—yield

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

if there is positive entry in $(\lambda, \kappa, \omega)$, where $\bar{\alpha} \equiv (1 - \alpha) \alpha^{\alpha/(1-\alpha)}$. Facing the wage profile $W_t(\lambda, \kappa, \omega)$, each worker $z \in \mathcal{Z}_t(\lambda)$ chooses the equipment-task pair (κ, ω) that maximizes her wage, $\varepsilon_t(z, \kappa, \omega) W_t(\lambda, \kappa, \omega)$.

The assumption that idiosyncratic productivity is distributed multivariate Fréchet implies that the probability that a randomly sampled worker, $z \in \mathcal{Z}_t(\lambda)$, uses equipment κ in task ω is

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

The higher is $\theta(\lambda)$ —either because efficiency units are less dispersed or more correlated across (κ, ω) for a given worker—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 task satisfies

$$\left(\frac{\pi_t(\lambda', \kappa', \omega)}{\pi_t(\lambda', \kappa, \omega)} \right)^{1/\theta(\lambda')} \bigg/ \left(\frac{\pi_t(\lambda, \kappa', \omega)}{\pi_t(\lambda, \kappa, \omega)} \right)^{1/\theta(\lambda)} = \frac{T_t(\lambda', \kappa', \omega)}{T_t(\lambda', \kappa, \omega)} \bigg/ \frac{T_t(\lambda, \kappa', \omega)}{T_t(\lambda, \kappa, \omega)}, \quad (4)$$

so that if λ' workers (relative to λ) have a comparative advantage using κ' (relative to κ) in task ω , then they are relatively more likely to be allocated to κ' in task ω , adjusted for potentially different values of $\theta(\lambda')$ and $\theta(\lambda)$. Similar conditions hold for the allocation of workers to tasks (within an equipment type) and for the allocation of equipment to

(i.e. $\alpha = 0$ for the latter task).

tasks (within a worker group).

The average wage of workers in group λ teamed with equipment κ in task ω , denoted by $w_t(\lambda, \kappa, \omega)$, is the integral of $\varepsilon_t(z, \kappa, \omega) W_t(\lambda, \kappa, \omega)$ across workers teamed with κ in task ω , divided by the mass of these workers. Using the distributional assumption on idiosyncratic productivity, we obtain

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

where $\gamma(\lambda) \equiv \Gamma\left(1 - \frac{1}{\theta(\lambda)(1-\nu(\lambda))}\right)$ and $\Gamma(\cdot)$ is the Gamma function. An increase in productivity or task 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 of all λ workers in (κ, ω) increases less than proportionately due to self-selection, i.e. $\pi_t(\lambda, \kappa, \omega)$ increases, which lowers the average efficiency units of λ workers using equipment κ in task ω .

Denoting by $w_t(\lambda)$ the average wage of workers in group λ , the previous expression and equation (3) imply $w_t(\lambda) = w_t(\lambda, \kappa, \omega)$ for all (κ, ω) , where¹²

$$w_t(\lambda) = \bar{\alpha} \gamma(\lambda) \left(\sum_{\kappa, \omega} \left(T_t(\lambda, \kappa, \omega) p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} p_t(\omega)^{\frac{1}{1-\alpha}} \right)^{\theta(\lambda)} \right)^{1/\theta(\lambda)}. \quad (5)$$

General equilibrium. We normalize the price of the final good to one. In any period, task prices $p_t(\omega)$ must be such that total expenditure in task ω is equal to total revenue earned by all factors employed in task ω ,

$$\mu_t(\omega) p_t(\omega)^{1-\rho} E_t = \frac{1}{1-\alpha} \zeta_t(\omega), \quad (6)$$

where $E_t \equiv \frac{1}{1-\alpha} \sum_{\lambda} w_t(\lambda) L_t(\lambda)$ 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 task ω . The left-hand side of equation (6) is expenditure on task ω and the right-hand side is total income earned by factors employed in task ω . The price of each type of equipment is

$$p_t(\kappa) = q_t(\kappa) \quad (7)$$

according to equation (2). In equilibrium, the aggregate quantity of the final good is

¹²In Appendix C we incorporate preference shifters for working in different tasks, similar to Heckman and Sedlacek (1985), generating differences in average wages across tasks within a labor group.

$Y_t = E_t$, 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).

2.3 Decomposing changes in relative wages

Our aim is to decompose observed changes in relative average wages of any two labor groups— $\hat{w}(\lambda) / \hat{w}(\lambda_1)$, where we denote changes over time by $\hat{x} \equiv x_{t_1} / x_{t_0}$ for any variable x —between any two periods t_0 and t_1 . In order to separately identify the impact of changes in labor productivity, task shifters, and equipment productivity, in our baseline specification we impose the following restriction on productivity:

$$T_t(\lambda, \kappa, \omega) \equiv T_t(\lambda) T_t(\kappa) T_t(\omega) T(\lambda, \kappa, \omega). \quad (8)$$

Whereas we allow worker group, $T_t(\lambda) \geq 0$, equipment type, $T_t(\kappa) \geq 0$, and task, $T_t(\omega) \geq 0$, productivity to vary over time, we impose that the interaction between worker group, equipment type, and task productivity, $T(\lambda, \kappa, \omega) \geq 0$, is constant across time. That is, we assume that comparative advantage is fixed over time. In Section 5.3 we allow for more general changes in technology, allowing comparative advantage to vary over time.

We group changes in the economic environment into four components: (i) changes in labor composition, $\hat{L}(\lambda)$; (ii) changes in labor productivity, $\hat{T}(\lambda)$; (iii) changes in task shifters, which combine changes in the productivity of workers employed in different tasks, $\hat{T}(\omega)$, and in the demand for these tasks, $\hat{\mu}(\omega)$, as captured by changes in the term $\tilde{\mu}_t(\omega) \equiv \mu_t(\omega) T_t(\omega)^{(1-\alpha)(\rho-1)}$; and (iv) changes in equipment productivity, which combine changes in the productivity, $\hat{T}(\kappa)$, and production cost, $\hat{q}(\kappa)$, of different types of equipment, as captured by changes in the term $\tilde{T}_t(\kappa) \equiv q_t(\kappa)^{\frac{-\alpha}{1-\alpha}} T_t(\kappa)$.¹³

Changes in average wages, using equations (5) and (8), are given by

$$\hat{w}(\lambda) = \hat{T}(\lambda) \left[\sum_{\kappa, \omega} \left(\tilde{p}(\omega) \tilde{T}(\kappa) \right)^{\theta(\lambda)} \pi_{t_0}(\lambda, \kappa, \omega) \right]^{1/\theta(\lambda)}, \quad (9)$$

where we have defined transformed task prices, $\tilde{p}_t(\omega) \equiv p_t(\omega)^{1/(1-\alpha)} T_t(\omega)$, which depend

¹³One could use additional information on changes in equipment or task prices, which are hard to measure in practice, to separate the effects on relative wages of $\hat{T}(\kappa)$ and $\hat{q}(\kappa)$ or $\hat{T}(\omega)$ and $\hat{\mu}(\omega)$, respectively.

both on task prices and productivities. According to equation (9), changes in wages are proportional to changes in labor productivity, $\hat{T}(\lambda)$, and are a CES combination of changes in transformed task prices, $\hat{\tilde{p}}(\omega)$, and equipment productivities, $\hat{\tilde{T}}(\kappa)$, where the weight given to changes in each of these components depends on factor allocations in the initial period t_0 , $\pi_{t_0}(\lambda, \kappa, \omega)$. Taking a first-order approximation of equation (9) yields

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

Hence, changes in average wages in response to given changes in transformed task prices and equipment productivities do not depend—to a first-order approximation—on the value of $\theta(\lambda)$, which determines the extent to which workers reallocate.¹⁴ However, because the value of $\theta(\lambda)$ determines the extent of factor reallocation in response to given changes in transformed task prices and equipment productivities, the value of $\theta(\lambda)$ does affect the first-order response of transformed task prices to shocks.

Changes in transformed task prices are determined by the following system of equations

$$\hat{\pi}(\lambda, \kappa, \omega) = \frac{\left(\hat{\tilde{p}}(\omega) \hat{\tilde{T}}(\kappa) \right)^{\theta(\lambda)}}{\sum_{\kappa', \omega'} \left(\hat{\tilde{p}}(\omega') \hat{\tilde{T}}(\kappa') \right)^{\theta(\lambda)} \pi_{t_0}(\lambda, \kappa', \omega')}, \quad (11)$$

$$\hat{\mu}(\omega) \left(\hat{\tilde{p}}(\omega) \right)^{(1-\alpha)(1-\rho)} \hat{E} = \frac{1}{\zeta_{t_0}(\omega)} \sum_{\lambda, \kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega). \quad (12)$$

Whereas changes in labor composition and task shifters affect wages only indirectly through transformed task prices—equations (11) and (12)—labor and equipment productivity affect wages both directly and indirectly.

We can re-express equations (9), (11), and (12) so that relative wages, $\hat{w}(\lambda) / \hat{w}(\lambda_1)$, and price changes, $\hat{\tilde{p}}(\omega) / \hat{\tilde{p}}(\omega_1)$, depend on relative shocks to labor composition, task shifters, equipment productivity, and labor productivity, as shown in Appendix B. This transformation is useful because we will measure only relative shocks to task shifters, equipment productivity, and labor productivity.

2.4 Intuition

The impact of shocks on relative wages can be understood as follows.

¹⁴Equation (10) provides a micro-foundation for the regression model that Acemoglu and Autor (2011) offer as a stylized example of how their model might be brought to the data.

Consider an increase in the task shifter $\tilde{\mu}_t(\omega)$, i.e. $\tilde{\mu}(\omega) > 1$. By equation (9), $\tilde{\mu}(\omega)$ only impacts average wages through general equilibrium changes in transformed prices $\hat{p}(\omega)$. This shock raises the price of task ω and, therefore, the average wages of worker groups that are disproportionately employed in task ω . Similarly, an increase in labor supply $L_t(\lambda)$ reduces the prices of tasks in which group λ is disproportionately employed. This lowers the relative wage not only of group λ , but also of worker groups employed in similar tasks as λ . An increase in labor productivity $T_t(\lambda)$ directly raises the relative wage of group λ and affects all other labor groups through changes in task prices similarly to an increase in $L_t(\lambda)$.¹⁵ In all cases, the effect through task prices is stronger for lower values of ρ , since task prices are more responsive to shocks in this case.

Finally, consider the impact on relative wages of a change in the productivity of equipment κ . An increase in $\tilde{T}_t(\kappa)$ raises the relative wages of worker groups that use κ intensively. It also reduces transformed task prices, i.e. $\hat{p}(\omega) < 1$, in the tasks in which κ is used intensively, lowering the relative wages of worker groups that tend to be employed in these tasks. The impact on relative wages of changes in equipment productivity depend on whether aggregate patterns of labor allocation across equipment types are generated directly by labor-equipment comparative advantage or indirectly by labor-task and equipment-task 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 $\tilde{T}_t(\kappa)$ does not affect relative task prices. In this case, relative wages are affected only directly through changes in equipment productivity.

On the other hand, if there is no comparative advantage between workers and equipment but there is comparative advantage between workers and tasks and between equipment and tasks, then an increase in $\tilde{T}_t(\kappa)$ directly increases the relative wage of workers employed in κ -intensive tasks and indirectly, through task prices, reduces the relative wage of workers employed in κ -intensive tasks. The relative strength of the direct and indirect channels depends on ρ . The relative wage of workers employed in κ -intensive tasks falls—i.e. the indirect task price effect dominates the direct effect—if and only if $\rho < 1$. Intuitively, an increase in $\tilde{T}_t(\kappa)$ acts like a positive productivity shock to the tasks in which κ has a comparative advantage. If $\rho < 1$ this reduces employment and the relative wages of worker groups disproportionately employed in the tasks in which κ has a

¹⁵Costinot and Vogel (Forthcoming) provide analytic results on the implications for relative wages of changes in labor composition, $L_t(\lambda)$, and task demand, $\mu_t(\omega)$, in an environment to which our model limits when $\alpha = 0$ (i.e., in the absence of equipment) and when $T(\lambda, \omega)$ —i.e. our $T(\lambda, \kappa, \omega)$ in the absence of equipment—is log-supermodular.

comparative advantage.¹⁶

3 Parameterization

In our baseline approach, we impose a common $\theta(\lambda)$ for all λ and denote this value by θ . We relax this restriction in our robustness section.

Using equations (9), (11), and (12) to quantify the impact of shocks between period t_0 and t_1 on relative wages, we require: (i) period t_0 measures of factor allocations, $\pi_{t_0}(\lambda, \kappa, \omega)$, relative average wages, $w_{t_0}(\lambda)$, labor composition, $L_{t_0}(\lambda)$, and the share of labor payments by occupation, $\zeta_{t_0}(\omega)$; (ii) measures of shocks to labor composition, $\hat{L}(\lambda)/\hat{L}(\lambda_1)$, task shifters, $\hat{\mu}(\omega)/\hat{\mu}(\omega_1)$, equipment productivity (to the power θ), $\hat{T}(\kappa)^\theta/\hat{T}(\kappa_1)^\theta$, and labor productivity, $\hat{T}(\lambda)/\hat{T}(\lambda_1)$; and (iii) the parameters α , ρ , and θ .

3.1 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 A. 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 10 labor groups by gender and education (high school dropouts, HSD; high school graduates, HSG; some college, SMC; completed college, CLG; and graduate training, GTC); as a robustness check, we consider 30 labor groups in Section 5.4. We consider two types of equipment: computers and other equipment. We link tasks in the model to occupations in the data and consider thirty occupations, which we list, together with summary statistics, in Table 10 in Appendix A.¹⁷

¹⁶More generally—without imposing Cobb Douglas task-level production functions—in the absence of labor-equipment comparative advantage an increase in equipment κ productivity reduces employment in tasks in which κ has a comparative advantage if the elasticity of substitution between labor and equipment at the level of the production unit, assumed to be one in our model, is greater than the elasticity of substitution between tasks, assumed to be ρ in our model.

¹⁷An alternative approach, discussed in Autor (2013), maps tasks to data by appending a set of standardized job descriptors to each occupation (e.g., routineness). We have taken a step in this direction, estimating comparative advantage between worker *characteristics*, occupation *characteristics*, and equipment types (rather than measuring comparative advantage between worker groups, occupations, and equipment types, as in our baseline approach). Details are available upon request.

We use the MORG CPS to construct total hours worked and average wages by labor group by year.¹⁸ 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 , denoted by $\pi_t(\lambda, \kappa, \omega)$. 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 main-frame. We construct $\pi_t(\lambda, \kappa', \omega)$ as the hours worked in occupation ω by λ workers who report that they use a computer κ' at work relative to 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 $\kappa'' = \text{other equipment}$) relative to total hours worked by labor group λ in year t .¹⁹

Constructing allocations, $\pi_t(\lambda, \kappa, \omega)$, as we do introduces three 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.²⁰

3.2 Factor allocations

In Table 1 we showed that women and more educated workers use computers more intensively than men and less educated workers, respectively, by aggregating $\pi_t(\lambda, \kappa, \omega)$ across ω and λ . To quantify the impact of shocks between t_0 and t_1 , however, we require the disaggregated measures of factor allocations. Here we identify a few key patterns in the $\pi_t(\lambda, \kappa, \omega)$ data.

¹⁸We measure wages using the MORG CPS rather than the March CPS because the March CPS does not directly measure hourly wages of workers paid by the hour and, therefore, introduces substantial measurement error in individual wages; see Lemieux (2006). Both datasets imply similar changes in average wages within a labor group. However, measurement error in individual wages introduce bias in one of our approaches to estimate $\theta(\lambda)$.

¹⁹We observe $\pi_t(\lambda, \kappa, \omega) = 0$ in the data for roughly 12% of $(\lambda, \kappa, \omega)$ observations in any given year.

²⁰The German *Qualification and Working Conditions* survey helps mitigate the second and third issues by providing data on worker usage of multiple types of equipment and the share of hours spent using computers. In ongoing work we have found the following preliminary results using this data. We find in the raw data similar patterns of factor allocations—on computer use by education group and gender within occupation—as in the U.S. data described below. In our counterfactual analysis applied to Germany between 1986 and 2006, in which we include three types of equipment—computers, writing implements, and other equipment—we find that computerization generates a substantial rise in the skill premium but only a small reduction in the gender gap, as we find in the U.S., detailed below. These results are available upon request. We discuss the relationship between our work and DiNardo and Pischke (1997), which uses this dataset, below.

To determine the extent to which college educated workers (λ') compared to workers with high school degrees in the same gender 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 2 genders. 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.




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

The right panel of Figure 2 plots a similar histogram, where λ' and λ denote female and male workers, respectively, in the same education groups. 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).

We can similarly study the extent to which worker 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 worker 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 worker group. These comparisons provide an example of a more general relationship: women tend to be employed in occupations in which all worker groups are relatively more likely to use computers.

3.3 Measuring shocks

Here we describe how we measure shocks to labor composition, $\hat{L}(\lambda)/\hat{L}(\lambda_1)$, equipment productivity (to the power θ), $\hat{T}(\kappa)^\theta/\hat{T}(\kappa_1)^\theta$, task shifters, $\hat{\mu}(\omega)/\hat{\mu}(\omega_1)$, and labor productivity, $\hat{T}(\lambda)/\hat{T}(\lambda_1)$. We measure the change in labor composition directly from the MORG CPS. We measure changes in equipment productivity using data only on changes

in equipment usage over time. We measure changes in task shifters using data on changes in occupation employment over time as well as changes in the share of labor payments by occupation. Note that we do not directly use changes in relative wages to measure changes in either equipment productivity or task shifters. Finally, given measured changes in equipment productivity and task shifters, and an estimate of θ , we measure changes in labor productivity as a residual to match relative wages. Details are provided in Appendix B.

Consider first our measure of changes in equipment productivity (to the power θ). Equations (3) and (8) give,

$$\frac{\hat{T}(\kappa)^\theta}{\hat{T}(\kappa_1)^\theta} = \frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa_1, \omega)} \quad (13)$$

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 task ω will increase. It is important to condition on (λ, ω) pairs when identifying changes in equipment productivity (to the power θ) 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 task shifters that are biased towards the occupations in which computers have a comparative advantage.

Figure 3 shows a large increase between each pair of consecutive years (i.e., 1984-1989, 1989-1993, etc...) in computer usage within most (λ, ω) pairs. This pattern is consistent




Figure 3: Histogram across (λ, ω) pairs of the logarithm of the right-hand-side of equation (13), with κ =computers and κ_1 =other equipment, between all pairs of consecutive years

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.²¹

²¹The decline over time in the U.S. in the price of equipment relative to structures—see e.g. Greenwood et al. (1997)—is mostly driven by a decline in computer prices. For example, between 1984 and 2003: (i) the

The few (λ, ω) pairs, visible in Figure 3, that experience a decrease in the share of hours spent using computers tend to have a small number of observations in the relevant period t_0 or t_1 , suggestive of measurement error. More generally, in our data the right-hand-side of equation (13) is clearly not equal across (λ, ω) pairs for any given pair of years t_0 and t_1 . In our baseline approach we rationalize these differences through measurement error in observed allocations and measure the left-hand-side of equation (13) using the exponential of the sample average over all (λ, ω) pairs of the logarithm of the right-hand side.²² We take a complementary approach to rationalize these differences in Section 5, where we allow for comparative advantage to evolve over time.

To measure changes in task shifters, we first use equations (3) and (8) to obtain measures of changes in transformed task prices (to the power θ) between t_0 and t_1 ,

$$\frac{\hat{p}(\omega)^\theta}{\hat{p}(\omega_1)^\theta} = \frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa, \omega_1)} \quad (14)$$

for any (λ, κ) pair.²³ Hence, if the transformed task price of ω rises relative to ω_1 between t_0 and t_1 , then the share of λ hours spent working with κ in task ω relative to in task ω_1 will increase. As above, it is important to condition on (λ, κ) when identifying changes in transformed task prices; and, as above, we measure the left-hand-side of equation (14) using the exponential of the sample average over all (λ, κ) pairs of the logarithm of the right-hand side. Given values for α , ρ , and θ , we then use the task-market clearing condition and changes in labor payments by occupation to recover changes in task shifters,

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

If $\rho = 1$, measured changes in task shifters depend only on changes in the share of labor payments across occupations, shown in Figure 1, exactly as in, e.g. Autor et al. (1998).

price of industrial equipment and transportation equipment relative to computers and peripheral equipment has risen by a factor 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 industrial equipment and 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). We do not use equipment price or quantity data directly in our procedure in part because of quality-adjustment issues raised by, e.g., Gordon (1990).

²²In robustness exercises we have also used weighted averages as well as dropped outliers in constructing changes in both equipment productivity (to the power θ) and transformed task prices. Our results are robust to these alternatives.

²³Here we use observed changes in allocations between t_0 and t_1 . Of course, in calculating changes in relative wages in response to a subset of shocks, we solve for counterfactual changes in transformed task prices and allocations using equations (11) and (12).

Finally, given an estimate of θ , measures of allocations in period t_0 , and measures of changes in relative transformed task prices and equipment productivities, both to the power θ , we measure changes in labor productivity as a residual to match exactly changes in relative wages,

$$\frac{\hat{T}(\lambda)}{\hat{T}(\lambda_1)} = \frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} \left(\frac{S(\lambda_1)}{S(\lambda)} \right)^{1/\theta} \quad (16)$$

where

$$S(\lambda) = \sum_{\kappa, \omega} \frac{\hat{p}(\omega)^\theta}{\hat{p}(\omega_1)^\theta} \frac{\hat{T}(\kappa)^\theta}{\hat{T}(\kappa_1)^\theta} \pi_{t_0}(\lambda, \kappa, \omega) \quad (17)$$

is a labor-group specific weighted average of changes in relative transformed task prices to the power θ and equipment productivities to the power θ , which are pinned down in the data independently of the value of θ using equations (13) and (14). Note that changes in relative wages only directly affect our measures of changes in labor productivities.

3.4 Parameter estimates

We assign the values of α and ρ as follows. The parameter α determines the share of payments that accrue to equipment. We set $\alpha = 0.24$, consistent with estimates in [Burstein et al. \(2013\)](#). Note that when $\rho = 1$ the value of $\alpha \in (0, 1)$ does not impact any of our quantitative results. The parameter ρ determines the elasticity of substitution between tasks. We set $\rho = 0.9$ as in [Goos et al. \(2014\)](#), who estimate this elasticity using 21 occupations in 16 Western European countries. In our robustness section we conduct sensitivity analyses using values of ρ ranging from 0.1 to 10.

In our baseline approach for estimating θ , we express changes in relative labor productivity as following a linear time trend with deviations around the trend,

$$\log \frac{\hat{T}(\lambda)}{\hat{T}(\lambda_1)} = \gamma(\lambda) \times (t_1 - t_0) + \iota(\lambda, t_0, t_1).$$

Hence, changes in relative average wages can be expressed as

$$\log \frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} = \gamma(\lambda) (t_1 - t_0) + \frac{1}{\theta} \log \frac{S(\lambda)}{S(\lambda_1)} + \iota(\lambda, t_0, t_1). \quad (18)$$

Note that equation (18) is equivalent to the equation that is used to estimate the key elasticity in the canonical model, see e.g. [Katz and Murphy \(1992\)](#), except we replace changes in relative labor-group specific hours worked, $\hat{L}(\lambda) / \hat{L}(\lambda_1)$, with relative labor-group specific weighted averages of measured changes in transformed task prices and equip-

ment productivities to the power θ , $S(\lambda) / S(\lambda_1)$.

We observe changes in relative wages between each pair of consecutive years in our sample and we construct $S(\lambda) / S(\lambda_1)$ between each pair of years using equation (17), period t_0 allocations, and our measures of changes in equipment productivity and transformed task prices to the power θ , all of which are measured independently of θ . Using equation (18), we estimate θ and $\gamma(\lambda)$, a total of 10 parameters, via OLS. This estimator is consistent if deviations from trend in labor productivity are mean-independent of $S(\lambda)$. Of course, these deviations affect task prices, and task prices enter $S(\lambda)$. This bias decreases in the number of labor groups if deviations from trend are uncorrelated across labor groups. This approach yields $\theta = 2.67$.

In an alternative approach, based on Lagakos and Waugh (2013) and Hsieh et al. (2013), we use the empirical distribution of wages within each λ to estimate $\tilde{\theta}(\lambda)$ —separately for each labor group λ —using both maximum likelihood and method of moments, and recover $\theta(\lambda) = \tilde{\theta}(\lambda) / (1 - \nu(\lambda))$ —where $\nu(\lambda)$ governs the λ -specific correlation of each worker’s efficiency units across (κ, ω) pairs—using the value of $\nu(\lambda)$ obtained by Hsieh et al. (2013). This approach, described in Section 5.1, yields an average of θ across λ equal to 2.50 and 3.82 using MLE and method of moments, respectively.

4 Results

In this section we summarize our baseline results.²⁴ Specifically, given our baseline parameter values, initial allocations, and measures of the four shocks, we quantify—using equations (9), (11), and (12)—the implications for relative wages of changes in labor composition, task shifters, equipment productivity, and labor productivity. We report results for each sub-period in our data: 1984-1989 (that is, we set $t_0 = 1984$ and $t_1 = 1989$), 1989-1993, 1993-1997, and 1997-2003. We also report the cumulative change in log relative wages between 1984 and 2003, calculated as the sum of the log change over all sub-periods in our data.²⁵

Skill premium. We begin by decomposing changes in the skill premium between each pair of consecutive years and over the full sample, displayed in Table 2. The first column reports the change in the data, which is also the change predicted by the model when all changes—in labor composition, task shifters, equipment productivity, and labor

²⁴All wage changes reported in the paper are *composition adjusted*, i.e. holding fixed over time the relative employment shares of demographic groups. We use the same procedure in the data and in the model.

²⁵For cumulative changes in log relative wages we obtain very similar results if we set $t_0 = 1984$ and $t_1 = 2003$ instead of adding changes in log relative wages over all sub-periods.

productivity—are simultaneously considered. Between 1984 and 2003 the skill premium increased by 16.1 log points, 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 component is changed over time (i.e. holding the other components at their t_0 level).

	Data	Labor comp.	Task shifters	Equip. prod.	Labor prod.
1984 - 1989	0.062	-0.028	0.043	0.023	0.023
1989 - 1993	0.070	-0.013	0.020	0.022	0.038
1993 - 1997	0.033	-0.023	0.040	0.006	0.010
1997 - 2003	-0.005	-0.044	0.030	0.018	-0.008
1984 - 2003	0.161	-0.108	0.134	0.070	0.063

Table 2: Decomposing changes in the log skill premium

Changes in labor composition decrease the skill premium between every pair of years in response to the large increase in the share of hours of more educated workers over this period. The increase in hours worked by those with college degrees relative to those without of 46 log points between 1984 and 2003 decreases the skill premium by 10.8 log points. Changes in relative demand must, therefore, not only generate the observed rise of the skill premium in the data, but also compensate for the impact of changes in labor composition.

The combination of changes in equipment productivity and task shifters explain the majority of the rise in the skill premium between 1984 and 2003. The equipment productivity component alone accounts for roughly 26% of the sum of the forces pushing the skill premium upwards: $0.26 \simeq 0.070 / (0.134 + 0.070 + 0.063)$. 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.

We obtain the result that computerization has substantially increased the U.S. skill premium because the data indicate that there is: (i) strong education-computer comparative advantage (see Figure 2), (ii) a substantial share of workers using computers (see Table 1), and (iii) large growth in computer usage within worker-task pairs (see Figure 3).²⁶

²⁶DiNardo and Pischke (1997) critique Krueger (1993) by showing that pencils can explain wage premia as well as computers. Their critique does not apply here for two reasons. First, our approach is fundamentally different from Krueger (1993). Instead of using the October Supplement to regress wages on computer

The task shifter component accounts for roughly 50% of the sum of the forces pushing the skill premium upwards over the full sample. We obtain the result that task shifters have substantially increased the U.S. skill premium because we find: (i) large variation in education-intensity across tasks, (ii) a substantial share of workers in expanding and contracting tasks, and (iii) large changes in task shifters. This is related to a key result in Autor et al. (2003)—who, using a shift-share analysis, find that task shifters account for roughly 60% of the relative demand shift favoring college labor between 1970 and 1988—but is not implied by it. Under some assumptions (Cobb-Douglas utility and production functions), shift-share analyses structurally decompose into within and between task shifters changes in wage bill shares, i.e. changes in $w(\lambda) L(\lambda)$ relative to the sum of labor payments across all λ .²⁷ Changes in wage bill shares can be very different from changes in relative wages when changes in labor composition are large, as they are in the data.

Perhaps surprisingly, of the demand-side mechanisms affecting relative wages, the one that we measure as a residual to match observed changes in relative wages, labor productivity, only accounts for roughly 24% of the sum of the effects of the three demand-side mechanisms.

Disaggregated groups. Table 3 decomposes changes in between-education-group wage inequality at a higher level of worker disaggregation, comparing changes in average wages across the five education groups considered in our analysis over the full sample, 1984-2003. The 16.1 log point change in the skill premium aggregates across heterogeneous changes in relative wages between more disaggregated education groups.

	Data	Labor comp.	Task shifters	Equip. prod.	Labor prod.
HSG / HSD	0.112	-0.043	0.059	0.055	0.044
SMC / HSD	0.166	-0.087	0.121	0.103	0.031
CLG / HSD	0.254	-0.147	0.185	0.129	0.088
GTC / HSD	0.305	-0.178	0.238	0.139	0.107

Table 3: Decomposing changes in log relative wages across education groups between 1984 and 2003

usage, we use it to identify comparative advantage and measure shocks. Second, in order for pencils to drive changes in wages (as we find computers do), we would have to find (i) strong education-pencil comparative advantage (identified within tasks), (ii) a large share of workers using pencils, and (iii) extremely large and systematic changes in pencil usage within worker-task pairs over time. Using the German *Qualification and Working Conditions* dataset, we find that pencil use, measured by whether a worker answers that writing implements are his/her most-used tool at work, fails all three criteria, unlike computer use.

²⁷Katz and Autor (1999) show that shift-share analyses provide a local decomposition of relative wages (that is, given arbitrarily small changes in labor composition and wages).

Nevertheless, the results reported in Table 2 are robust: the labor productivity component is not particularly important for explaining the rise in between-education-group inequality.

Gender wage gap. The average wage of men relative to women, the gender gap, declined by 13.2 log points between 1984 and 2003. Table 4 decomposes changes in the gender wage gap between each pair of consecutive years and over the full sample. The increase in hours worked by women relative to men of roughly 12 log points between 1984 and 2003 increases the gender gap by 4.5 log points.

	Data	Labor comp.	Task shifters	Equip. prod.	Labor prod.
1984 - 1989	-0.058	0.012	-0.016	-0.003	-0.050
1989 - 1993	-0.050	0.014	-0.036	-0.003	-0.023
1993 - 1997	-0.001	0.007	0.008	-0.001	-0.015
1997 - 2003	-0.023	0.012	-0.046	-0.003	0.016
1984 - 2003	-0.132	0.045	-0.090	-0.010	-0.072

Table 4: Decomposing changes in the log gender gap

Equipment productivity has almost no effect on the gender wage gap—in spite of the fact that women are substantially more likely to use computers at work than men—because we incorporate comparative advantage between worker groups, equipment types, and tasks. Whereas women have a slight comparative advantage using computers (see Figure 2)—through which computerization decreases the gender gap—they have a comparative advantage in the tasks in which computers have a comparative advantage—through which computerization increases the gender gap with $\rho < 1$. Unlike the skill premium, changes in labor productivity account for almost half (42%) of the impact of the demand-side forces affecting the gender wage gap and plays a central role in each sub-period except for 1997-2003. This suggests that changes in gender discrimination—if they affect labor productivity irrespective of (κ, ω) —may have played a substantial role in reducing the gender wage gap, as discussed in Hsieh et al. (2013), especially early in our sample (in the 1980s and early 1990s).

Job polarization. Autor et al. (2008) document simultaneous growth in the share of employment in high-skill, high-wage occupations and in low-skill, low-wage occupations (“job polarization”) in the U.S. starting in 1989; Goos and Manning (2007) and Goos et al. (2014) document similar patterns in the U.K. and in 15 additional European countries, respectively. Whereas our focus is on relative wages, our framework is also well-suited to analyze the effect of shocks on changes in employment across tasks. Grouping occupations into three categories by average wage across all years—High, Middle, and Low,

as detailed in Appendix A—we observe growth between 1989 and 2003 in the share of employment in High and Low wage occupations relative to Middle wage occupations, as shown in Table 5. Not surprisingly, we find that task shifters are the central force driving this result. However, the combination of changes in labor composition and changes in equipment productivity have a non-negligible role, explaining roughly 25% and 15% of the forces expanding High and Low wage occupations, respectively, relative to Middle wage occupations.

	Data	Labor comp. + Equip. prod.	Task shifters	Labor prod.
High - Middle employment share	0.16	0.04	0.14	-0.02
Low - Middle employment share	0.10	0.02	0.12	0.01

Table 5: Decomposing the difference in the (log) growth rate in the total hours of High versus Middle wage occupations, and Low versus Middle wage occupations between 1989 and 2003

5 Robustness and sensitivity analysis

In this section we consider four types of sensitivity exercises. First, we perform sensitivity to different values of ρ and θ . Second, we illustrate the importance of all three forms of comparative advantage by turning some of them off. Third, we allow for changes in comparative advantage over time. Finally, we consider a greater degree of disaggregation across labor groups.

5.1 Alternative parameter values

We now consider the sensitivity of our results as we vary θ and ρ . For each parameterization, we re-calculate task shifters, equipment shifters, and labor productivity.

Alternative values for θ . As briefly mentioned in Section 3.4, we obtain alternative estimates of $\theta(\lambda)$ using the empirical distribution of wages within each λ . Our assumption on the distribution of idiosyncratic productivity implies that the distribution of wages within a labor group λ is Fréchet with shape parameter $\tilde{\theta}(\lambda)$, where $\theta(\lambda) = \tilde{\theta}(\lambda) / (1 - \nu(\lambda))$. We take two approaches to recover $\tilde{\theta}(\lambda)$ using the empirical distribution of wages within each λ in the MORG CPS. First, we jointly estimate the shape and scale parameter for each λ in each year t using maximum likelihood (MLE). Figure 4 in Appendix A plots the empirical and predicted (using MLE) wage distributions for all ten labor groups in 2003.

Second, we use method of moments (MM) to match the coefficient of variation of wages for each λ within each year t , as in [Hsieh et al. \(2013\)](#). In both cases, we average across years our estimates of the shape parameter to obtain $\tilde{\theta}(\lambda)$ and obtain an estimate of $\theta(\lambda)$ from $\tilde{\theta}(\lambda)$ using [Hsieh et al.'s \(2013\)](#) implied estimate of $\nu \equiv \nu(\lambda) \approx 0.1$. Averaging $\theta(\lambda)$ across λ yields MLE and MM estimates of θ equal to 2.50 and 3.82, respectively.

In the left panel of Table 6 we decompose changes in the skill premium between 1984 and 2003 using these two alternative values of θ . We also use our MLE estimate of λ -specific values of $\theta(\lambda)$. Varying θ yields three conclusions. First, the combination of changes in equipment productivity and task shifters explains the majority of the rise in the skill premium for all alternative values of θ . Second, results are similar using the average value or heterogeneous values of $\theta(\lambda)$. Third, a higher value of θ increases the relative importance of changes in labor productivity. Here we describe the intuition for this third result.

value of θ	Labor comp.	Task shifters	Equip. prod.	Labor prod.	value of ρ	Labor comp.	Task shifters	Equip. prod.	Labor prod.
$\theta(\lambda)$ (MLE)	-0.120	0.140	0.071	0.067	$\rho = 0.1$	-0.162	0.257	0.006	0.052
$\theta = 2.50$ (MLE)	-0.113	0.141	0.074	0.058	$\rho = 0.9$	-0.108	0.134	0.070	0.063
$\theta = 2.67$ (baseline)	-0.108	0.134	0.070	0.063	$\rho = 1$	-0.103	0.123	0.075	0.064
$\theta = 3.82$ (MM)	-0.081	0.100	0.055	0.086	$\rho = 10$	-0.022	-0.066	0.167	0.080

Table 6: Decomposing changes in the log skill premium between 1984 and 2003 for alternative values of θ and ρ

As discussed in Section 2.2, whereas changes in average wages in response to given changes in task prices or productivities do not depend—to a first-order approximation—on the value of θ , the value of θ does affect the first-order response of task prices to shocks. However, the logic for how the value of θ affects our decomposition is more complicated because we are not taking as given changes to primitives. Instead, the value of θ shapes the shocks that we infer.

Recall from equation (16) that we measure changes in relative labor productivities to exactly match changes in relative wages given our measure of $S(\lambda) / S(\lambda_1)$, which is independent of θ , as well as the value of θ . In the extreme in which θ is infinite, changes in wages are explained exclusively by changes in relative labor productivities. More generally, a higher value of θ implies a larger role for changes in relative labor productivities and a smaller role for the remaining shocks in explaining changes in relative wages.

Alternative values for ρ . Recall that ρ is the elasticity of substitution across tasks in the aggregate production function. In the right panel of Table 6 we decompose changes in the skill premium between 1984 and 2003 using values of ρ ranging from 1/10 to 10. Our

main result—that the combination of changes in task shifters and equipment productivity explains the majority of the rise in between-education group inequality—is robust to alternative values of ρ within this range. Specifically, for all values of ρ between 1/10 and 10, changes in labor productivity increase the skill premium by between 5 and 8 log points.

However, the importance of task shifters falls substantially relative to the importance of equipment productivity as we increase ρ from 0.1 to 10. Intuitively, task prices are less responsive to changes in primitives for larger values of ρ . This directly reduces the impact of changes in labor composition, which affect wages only through changes in task prices. Given that educated workers tend to have a comparative advantage in the tasks in which computers do, this also increases the impact of changes in equipment productivity.

5.2 Sources of comparative advantage

To demonstrate the importance of including each of the three forms of comparative advantage, we perform two exercises. First, we abstract from comparative advantage related to tasks, imposing $\pi_t(\lambda, \kappa, \omega_i) = \pi_t(\lambda, \kappa, \omega)$ for all (λ, κ) and all i . This is equivalent, in terms of the model’s implications for changes in relative wages in response to changes in fundamentals, to assuming that there is a single task. Second, we abstract from comparative advantage related to equipment. To do so, we impose that $\pi_t(\lambda, \kappa_1, \omega) = \pi_t(\lambda, \kappa_2, \omega)$ for all (λ, ω) . This is equivalent to assuming that there is a single equipment good. In all cases, we hold the values of α , ρ , and θ fixed.

Table 7 reports our baseline decomposition between 1984-2003 both for the skill premium (in the left panel) and the gender wage 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 tasks.

	Skill premium				Gender gap			
	Labor comp.	Task shifters	Equip. prod.	Labor prod.	Labor comp.	Task shifters	Equip. prod.	Labor prod.
Baseline	-0.108	0.134	0.070	0.063	0.045	-0.090	-0.010	-0.072
Only labor-equipment CA	0	0	0.160	0.000	0	0	-0.070	-0.062
Only labor-task CA	-0.108	0.146	0	0.121	0.045	-0.063	0	-0.111

Table 7: Decomposing changes in the log skill premium and log gender gap between 1984 and 2003 under different assumptions on comparative advantage

Abstracting from any comparative advantage at the level of tasks (i.e. assuming away worker-task and equipment-task comparative advantage) has two effects. First, it im-

plies that the labor composition and task shifters components of our decomposition go to zero (because changes in labor composition and task shifters affect relative wages only through task prices). This affects the labor productivity component, since changes in labor productivity are identified as a residual to match observed changes in relative wages. Second, it implies that worker-equipment comparative advantage is the only force giving rise to the observed allocation of worker 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.

Table 7 shows that if we were to abstract from any comparative advantage at the level of tasks, we would incorrectly conclude that almost 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 in the absence of worker-task or equipment-task comparative advantage, we would incorrectly conclude that changes in equipment productivity played a central role in reducing the gender gap.

Similarly, abstracting from any comparative advantage at the level of equipment implies that the equipment productivity component of our decomposition goes to zero and that the only force giving rise to the allocation of worker groups to tasks is worker-task comparative advantage. Table 7 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 task shifters on the gender gap does not change dramatically, suggesting that the results in Hsieh et al. (2013) are robust to the inclusion of equipment.

In summary, abstracting from comparative advantage at the level of either tasks 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 inferred strength of the remaining source of comparative advantage.

5.3 Evolving comparative advantage

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

in occupation characteristics that differentially affect worker groups (e.g. job flexibility, which women may value relatively more, see e.g. [Goldin 2014](#)), or by changes in the characteristics of equipment. Here we generalize our baseline model to incorporate such potential changes over time and show that our results remain largely unchanged.

Specifically, we consider separately three extensions to 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 tasks in case 1, workers and equipment in case 2, and equipment and tasks in case 3. If we allow for the most general form of changes in comparative advantage in which the triple interaction between λ , κ , and ω evolves over time, then we could only decompose changes in between-group inequality into changes in the composition of the labor force (which would be unchanged relative to the baseline) and changes in overall productivity. In [Appendix D](#) we show how to measure the relevant shocks and how to decompose changes in between-group inequality into labor composition, task shifter, and labor-equipment components in case 2. Details for cases 1 and 3 are similar. [Table 8](#) 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.

Changes in CA	Labor comp.	Task shifters	Equip. prod.	Labor prod.	Labor-task	Labor-equip.	Equip.-task
None (baseline)	-0.108	0.134	0.070	0.063	-	-	-
Worker-task (case 1)	-0.108	-	0.070	-	0.203	-	-
Worker-equipment (case 2)	-0.108	0.128	-	-	-	0.141	-
Equip.-task (case 3)	-0.108	-	-	0.050	-	-	0.219

Table 8: 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 each case 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 (i.e., abstracting from all other shocks) 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, in case 1 our measure of changes in equipment productivity as well as the system of equa-

tions that determines their impact are exactly the same as in our baseline model. Hence, the equipment productivity component is unchanged from the baseline in case 1. In case 2, whereas our measure of changes in transformed task prices is exactly the same as in our baseline model, our measure of changes in task labor payment shares—and, therefore, our measure of task 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 task shifters is the same, our results on task shifters is very similar to the baseline in case 2. Finally, since the sum of all four (in the baseline model) or three (in the extensions considered here) components of our decomposition match the change in relative wages in the data, the change in wages resulting from sum of the labor productivity and task productivity components in our baseline (when fed in one at a time) must closely match the change in wages from the labor-task 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.

5.4 Greater worker disaggregation

In theory we could incorporate as many labor groups, equipment types, and tasks as the data permits without complicating our measurement of shocks or our estimation of θ . In practice, we are constrained by data. Specifically, as we increase the number of labor groups, equipment types, or tasks we increase the share of $(\lambda, \kappa, \omega)$ triplets for which $\pi_t(\lambda, \kappa, \omega) = 0$ and measurement error in factor allocations in general. Here we increase the number of labor groups from 10 to 30 by incorporating three age groups: 17-30, 31-43, and 44 and older. In this case, the share of $(\lambda, \kappa, \omega)$ observations for which $\pi_t(\lambda, \kappa, \omega) = 0$ rises from (roughly) 12% to 27%. Moreover, because we composition-adjust wages not only for gender and education, but also for age, we find slightly different changes in the skill premium, 15.1 instead of 16.1 log points, and the gender wage gap, -13.3 instead of -13.2 log points, between 1984 and 2003.

We re-measure all shocks and report results in Table 9 both using a re-estimated value of θ , $\theta = 2.02$, and using our baseline value of θ , $\theta = 2.67$. Our baseline results are quite robust to increasing worker group disaggregation. For example, when we re-estimate θ , we find that changes in labor productivity only account for roughly 19% of the sum of the effects of the three demand-side mechanisms pushing the skill premium upwards (compared to 24% in our baseline).

	Skill premium				Gender gap			
	Labor comp.	Task shifters	Equip. prod.	Labor prod.	Labor comp.	Task shifters	Equip. prod.	Labor prod.
10 labor groups: baseline	-0.108	0.134	0.070	0.063	0.045	-0.090	-0.010	-0.072
30 labor groups $\theta = 2.67$	-0.126	0.131	0.071	0.073	0.047	-0.091	-0.009	-0.076
30 labor groups $\theta = 2.02$	-0.156	0.163	0.085	0.057	0.058	-0.116	-0.008	-0.063

Table 9: Decomposing changes in the log skill premium and log gender gap between 1984 and 2003 with 30 labor groups

6 International trade

We extend our baseline model to incorporate international trade in equipment, sectoral output, and task output.²⁸ We show how the degree of openness is reflected in what we had treated as exogenous primitive shocks to the cost of producing equipment and to task shifters in our baseline closed-economy model. As a preliminary step, we first introduce sectors in the closed-economy model in Section 6.1 and show that we can decompose the effects of changes in task shifters into changes in (i) sector shifters and (ii) within-sector task shifters.

6.1 Sectors in a closed economy

Sectors are indexed by σ . The final good combines sectoral output, $Y_t(\sigma)$, according to a CES production function,

$$Y_t = \left(\sum_{\sigma} \mu_t(\sigma)^{1/\rho_{\sigma}} Y_t(\sigma)^{(\rho_{\sigma}-1)/\rho_{\sigma}} \right)^{\rho_{\sigma}/(\rho_{\sigma}-1)} \quad (19)$$

where $\rho_{\sigma} > 0$ is the elasticity of substitution across sectors and $\mu_t(\sigma) \geq 0$ is an exogenous demand shifter for sector σ . Sectoral output is itself a CES combination of the output of different tasks,

$$Y_t(\sigma) = \left(\sum_{\omega} \mu_t(\omega, \sigma)^{1/\rho} Y_t(\omega, \sigma)^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)} \quad (20)$$

²⁸See [Burstein et al. \(2013\)](#) and [Parro \(2013\)](#) for quantitative analyses of the impact of trade in capital equipment on the skill premium and [Stokey \(1996\)](#) for a theoretical treatment of trade in skill complementary capital. See e.g. [Grossman and Rossi-Hansberg \(2008\)](#) for a theoretical analysis of task trade and inequality and [Feenstra and Hanson \(1999\)](#) for an empirical treatment of offshoring and relative wages. Our framework, which includes multiple types of equipment and comparative advantage between labor groups and equipment types rationalizes the findings in [Caselli and Wilson \(2004\)](#), that countries with different distributions of education import different mixes of equipment.

where $Y_t(\omega, \sigma) \geq 0$ denotes the absorption of task ω in the production of sector σ , $\mu_t(\omega, \sigma) \geq 0$ is an exogenous demand shifter for task ω in sector σ , and $\rho > 0$ is the elasticity of substitution across tasks within each sector.

Tasks are produced exactly as in our baseline specification: a worker's productivity depends only on her task ω , and not on her sector of employment σ . 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 task according to equation (20).²⁹ Production and absorption must satisfy the resource constraint $Y_t(\omega) = \sum_{\sigma} Y_t(\omega, \sigma)$.

In Appendix E we show how we can use this extension to further decompose the effects of changes in task shifters into changes in (i) sector shifters and (ii) within-sector task shifters. Sector shifters are changes in the relative demand for sectoral output in the production of the final good, which combine changes in $\mu_t(\sigma) / \mu_t(\sigma_1)$ and $\mu_t(\omega_1, \sigma) / \mu_t(\omega_1, \sigma_1)$, so that a proportional change in $\mu_t(\omega, \sigma)$ across all ω in a given σ is treated as a sector shifter. Within-sector task shifters are changes in the relative productivity/demand for tasks within sectors, denoted by $\hat{\mu}(\omega, \sigma) / \hat{\mu}(\omega_1, \sigma)$, where, as in our baseline model, $\tilde{\mu}_t(\omega, \sigma)$ is a transformed variable combining $\mu_t(\omega, \sigma)$ and $T_t(\omega)$.

The system of equations to solve for changes in wages is very similar to equations (9), (11), and (12). The task-market clearing condition (12) must be generalized: the demand shifter across tasks on the left-hand side of equation (12), $\hat{\mu}(\omega) / \hat{\mu}(\omega_1)$, in our baseline model is replaced by a combination of changes in sector shifters, within-sector task shifters, and changes in sectoral prices. All shocks excluding sector shifters and within-sector task shifters are measured as in our baseline model. Given a value of ρ_{σ} , constructing sector shifters and within-sector task shifters is straightforward using readily available data on the share of labor payments (i) across sectors and (ii) across tasks (i.e. occupations in our application) within sectors.

Finally, in the special case in which $\rho = \rho_{\sigma}$, this extension is equivalent to our baseline model where the task shifter in our baseline model, $\mu_t(\omega)$, is given by $\sum_{\sigma} \mu_t(\omega, \sigma) \mu_t(\sigma)$. Therefore, the quantitative results from our decomposition (where task shifters now combine changes in sector shifters and within-sector task shifters) are exactly the same as in our baseline model (given the same value of ρ).

²⁹Alternatively, we could assume that worker productivity depends both on task 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 may become sparse—in the sense that there might be many $(\lambda, \kappa, \omega, \sigma, t)$ for which $\pi_t(\lambda, \kappa, \omega, \sigma) = 0$ —given this alternative assumption.

6.2 International trade in tasks, sectors, and equipment

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. We assume that labor is internationally immobile.

The final good and sector production functions in country n , the open economy counterparts of equations (19) and (20), are given by

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

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

Country n produces task, sector, and equipment output: $Y_n(\omega)$, $Y_n(\sigma)$, and $Y_n(\kappa)$, respectively. Its absorption of task, sector, and equipment goods is itself a CES aggregate of these goods sourced from all countries in the world. For example, absorption of task ω in country n is

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

where $D_{in}(\omega)$ is absorption in country n of task ω sourced from country i , and $\eta(\omega) > 1$ is the elasticity of substitution across source countries for task ω .³⁰ Trade is subject to iceberg transportation costs, where $d_{ni}(x) \geq 1$ denotes units of x output that must be shipped from origin country n in order for one unit to arrive in destination country i . For example, output of task ω in country n satisfies

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

Without loss of generality for our results on relative wages, we abstract from trade in the consumption good. Production and absorption must also satisfy $Y_n = C + \sum_{\kappa} q_n(\kappa) Y_n(\kappa)$ and $D_n(\omega) = \sum_{\sigma} D_n(\omega, \sigma)$. As in our baseline model, $Y_n(\omega)$ is the sum of output across all workers employed in ω . For the exercises we consider below, we do not need to specify conditions on trade balance in each country. We provide additional details in Appendix F.

³⁰We assume an Armington trade model only for expositional simplicity. Our results would also hold in a Ricardian model as in Eaton and Kortum (2002).

We consider two counterfactual exercises quantifying the impact of international trade on country n that do not require solving the full world general equilibrium or assigning parameters in any other country. The first counterfactual quantifies the impact on wages in country n at time t if it were to move to autarky, holding all of country n 's parameters fixed. The second 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 these counterfactuals, define $w_n(\lambda; \Phi_t, \Phi_t^*, d_t)$ to be the average wage of worker 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 (for equipment, sectors, and tasks) are d_t . Define $d_{n,t}^A$ to be an alternative matrix of world trade costs in which country n 's trade costs are infinite ($d_{in,t} = \infty$ for all $i \neq n$). Our first counterfactual calculates $\hat{w}_{n,t}^A(\lambda) \equiv w_n(\lambda; \Phi_t, \Phi_t^*, d_{n,t}^A) / w_n(\lambda; \Phi_t, \Phi_t^*, d_{n,t})$. Our second counterfactual calculates

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

which is simply $\hat{w}_{n,t_0}^A(\lambda) / \hat{w}_{n,t_1}^A(\lambda)$. Because the second counterfactual amounts to conducting the first counterfactual twice, at different points in time, we only explain how to solve the first counterfactual.

In Appendix F we show that $\hat{w}_{n,t}^A(\lambda) / \hat{w}_{n,t}^A(\lambda_1)$ is the solution to a system of equations very similar to the sector version of the closed-economy system. Changes in equipment absorption prices induced by moving to autarky at time t are a simple function of domestic absorption shares for each κ at time t and equipment trade elasticities for each κ , $\eta(\kappa)$. If the domestic absorption share of equipment type κ is low (and trade elasticities are common across equipment goods), then moving to autarky increases the relative absorption price of equipment type κ . This has the same implications for relative wages as an increase in the relative cost of producing equipment κ , $q(\kappa)$, in a closed economy.

Changes in sector shifters induced by moving to autarky at time t are a simple function of domestic absorption shares and export shares across sectors at time t ; sectoral trade elasticities, $\eta(\sigma)$; and the elasticity of substitution across sectors, ρ_σ . For instance, under mild parametric restrictions, if sector σ has a low export share and/or a high import share relative to σ' , then moving to autarky increases the demand for σ relative to σ' . This has the same implications for relative wages as an increase in $\mu(\sigma)$ relative to $\mu(\sigma')$ in a closed economy. Similarly, changes in within-sector task shifters induced by moving to

autarky at time t are a simple function of domestic absorption shares and export shares across tasks at time t ; task trade elasticities, $\eta(\omega)$; and the elasticity of substitution across tasks ρ .³¹

While measures of trade in equipment and sectoral output are typically available for many countries, constructing measures of trade in tasks (or occupations in our application) is more challenging.

7 Conclusions

In this paper we have developed a framework in which changes in workforce composition, task shifters, computerization, and labor productivity shape changes in between-group inequality across many labor groups. We have parameterized the model to match observed factor allocations and wages in the United States between 1984 and 2003. We have shown that the combination of computerization and task shifters (both of which are measured without directly using data on wage changes) explain the majority of the rise in the skill premium and the rise in inequality across more disaggregated education groups as well as roughly half the fall in the gender wage gap.

Our framework and estimation strategy could be used more broadly in any country with sufficiently rich data on worker allocation to equipment types and occupations, sectors, or firms. For instance, in ongoing work, we use this approach to study the evolution of between-group inequality in Germany.

In spite of its high dimensionality, our framework remains tractable, lending itself to a variety of extensions and applications. We have extended our model to incorporate international trade in equipment and sector output as well as offshoring of tasks and have shown that changes in import and export shares in each of these markets shape what we treated, in our baseline closed-economy model, as exogenous primitive shocks. It would be interesting to bring this extended model to the data. One challenge in implementing such an application is the lack of available data on trade in tasks. Another interesting extension would be to model inter- and intra-national trade—and, potentially frictional labor mobility as in, e.g., Redding (2012)—and use the information from regional analyses—see e.g. Autor et al. (2013), Autor and Dorn (2013), and Kovak (2013)—to dis-

³¹Our 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.

cipline the parameters that shape the outcomes of our aggregate counterfactual analyses.

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**, “Directed Technical Change,” *Review of Economic Studies*, October 2002, 69 (4), 781–809.
- , “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,
- 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.**, “The ‘Task Approach’ to Labor Markets: An Overview,” January 2013, (18711).
- **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.
- , —, **and Gordon H. Hanson**, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, October 2013, 103 (6), 2121–68.
- , **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.
- , —, **and Melissa S. Kearney**, “Trends in U.S. Wage Inequality: Revising the Revisionists,” *The Review of Economics and Statistics*, May 2008, 90 (2), 300–323.
- Berman, Eli**, **John Bound**, **and Zvi Griliches**, “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures,” *The Quarterly Journal of Economics*, May 1994, 109 (2), 367–97.
- Burstein, Ariel** **and Jonathan Vogel**, “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.
- Caselli, Francesco** and **Daniel J. Wilson**, “Importing technology,” *Journal of Monetary Economics*, January 2004, 51 (1), 1–32.
- Costinot, Arnaud** and **Andres Rodriguez-Clare**, *Trade Theory with Numbers: Quantifying the Consequences of Globalization*, Vol. 4 of *Handbook of International Economics*, Elsevier,
- 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.
- Eaton, Jonathan** and **Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, September 2002, 70 (5), 1741–1779.
- 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.
- Garicano, Luis** and **Esteban Rossi-Hansberg**, “Organization and Inequality in a Knowledge Economy,” *The Quarterly Journal of Economics*, November 2006, 121 (4), 1383–1435.
- 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.
- Goos, Maarten**, **Alan Manning**, and **Anna Salomons**, “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 2014, 104 (8), 2509 – 2526.
- and —, “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *The Review of Economics and Statistics*, February 2007, 89 (1), 118–133.
- 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 Iouri 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.
- Kovak, Brian**, "Regional Effects of Trade Reform: What is the Correct Measure of Liberalization?," *American Economic Review*, 2013, 103 (5), 1960–1976.
- 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.
- 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 Rodríguez-Clare**, "Trade, Multinational Production, and the Gains from Openness," *Journal of Political Economy*, 2013, 121 (2), 273 – 322.
- Redding, Stephen J.**, "Goods Trade, Factor Mobility and Welfare," April 2012, (18008).
- Saez, Emmanuel**, "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, January 2001, 68 (1), 205–29.
- Stokey, Nancy L**, "Free Trade, Factor Returns, and Factor Accumulation," *Journal of Economic Growth*, December 1996, 1 (4), 421–47.
- Yeaple, Stephen Ross**, "A simple model of firm heterogeneity, international trade, and wages," *Journal of International Economics*, January 2005, 65 (1), 1–20.

A 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—for each worker group—of hours worked and income. 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 observations are multiplied by 1.5. Workers earning below \$1.675/hour in 1982 dollars are dropped, as are workers whose hourly wages exceed 1/35th the top-coded value of weekly earnings (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). Allocated earnings observations are excluded in all years. Our measure of labor composition, $L_t(\lambda)$, is hours worked within each labor group λ (weighted by sample weights).

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

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

When we group occupations into High, Middle, and Low wage bins in Section 4, we do so as follows. We calculate the average wage in each occupation in each year, take a simple average across years, and bin occupations as follows. High wage occupations include executive; management; architect; engineer; life, physical, and social science; computer and mathematical; lawyers; education, training, etc...; health diagnosing; and health assessment and treating. Middle wage occupations include community and social services; arts, design, entertainment, sports, media; technicians; financial sales; retail sales; administrative support; protective services; mechanics and repairers; construction; precision production; machine operators, assemblers, inspectors; and transportation and material moving. Low wage occupations include housekeeping, cleaning, laundry; food preparation and service; health service; building, grounds cleaning, maintenance; miscellaneous; child care, agriculture and mining; and handlers, equipment cleaners, helpers, and laborers.

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

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

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

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

College intensity (Female intensity) indicates hours worked in the occupation by the combination of CLG and GTC (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.

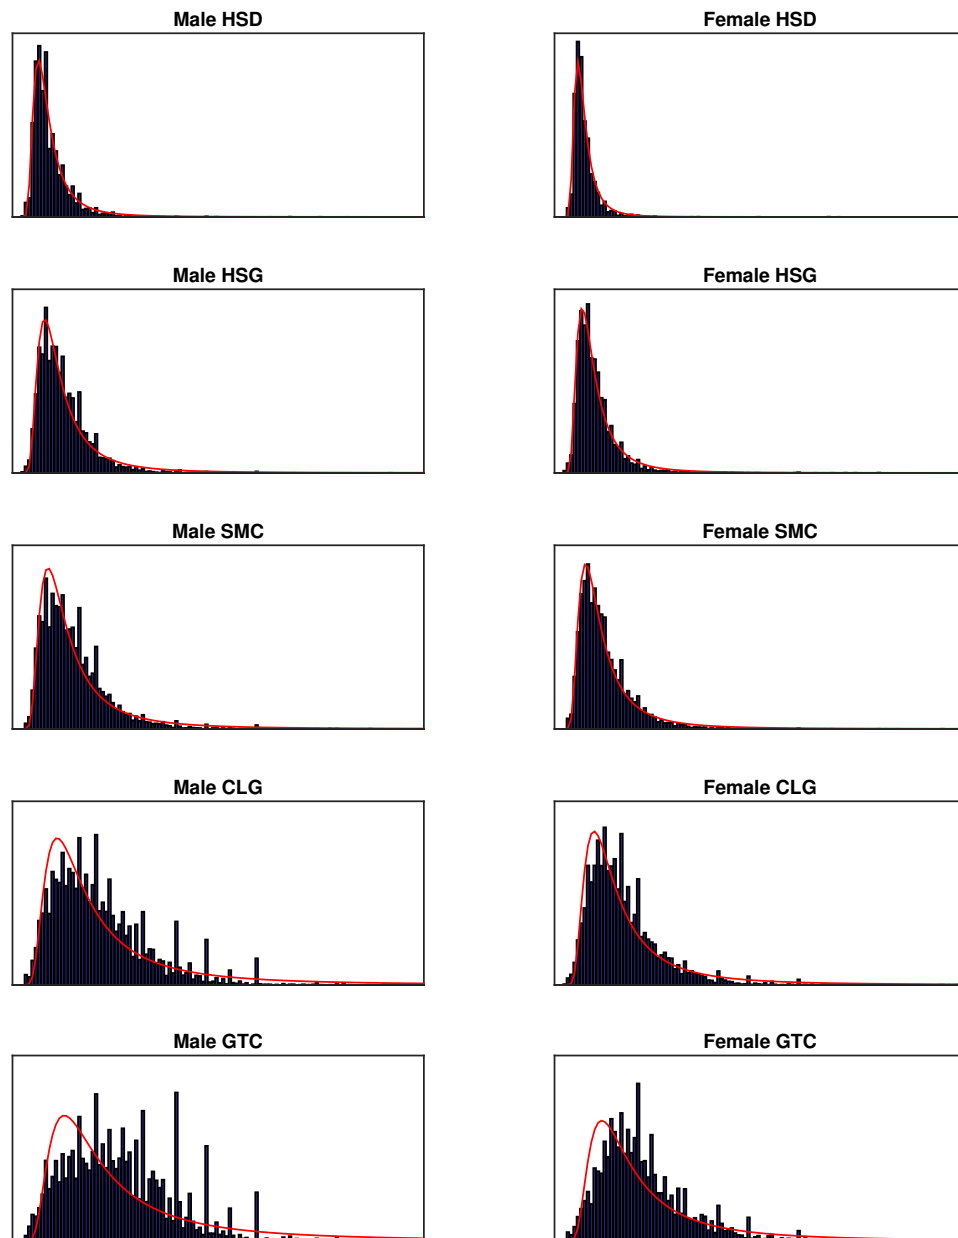


Figure 4: Empirical and predicted (Fréchet distribution estimated using maximum likelihood) wage distributions for all worker groups in 2003

B Measurement of shocks and estimation

B.1 Baseline

Here we describe in depth the steps we take to measure shocks and estimate our model. First, we re-express equations (9), (11), and (12) so that relative wages, $\hat{w}(\lambda) / \hat{w}(\lambda_1)$, and price changes, $\hat{\tilde{p}}(\omega) / \hat{\tilde{p}}(\omega_1)$, depend on relative shocks to labor composition, task shifters, equipment productivity, and labor productivity. Second, we measure changes in equipment and transformed task prices (both to the power θ) using only factor allocations. Third, we use these and changes in relative wages between groups to estimate θ . Finally, we measure changes in task shifters and changes in labor productivity.

Equations (9), (11), and (12) can be written as

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

$$\hat{\pi}(\lambda, \kappa, \omega) = \frac{\left(\frac{\hat{\tilde{p}}(\omega)}{\hat{\tilde{p}}(\omega_1)} \frac{\hat{T}(\kappa)}{\hat{T}(\kappa_1)} \right)^{\theta(\lambda)}}{\sum_{\kappa', \omega'} \left(\frac{\hat{\tilde{p}}(\omega')}{\hat{\tilde{p}}(\omega_1)} \frac{\hat{T}(\kappa')}{\hat{T}(\kappa_1)} \right)^{\theta(\lambda)} \pi_{t_0}(\lambda, \kappa', \omega')}$$

and

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

We measure changes in equipment productivity (to the power θ) using equation (13) as

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

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 zeroes in allocations we have $N(\kappa_1, \kappa_2) = 300$, which is the number of labor groups multiplied by the number of occupations.

We measure changes in transformed task prices relative to task ω_0 (to the power θ) using equation (14) as

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

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 zeroes in allocations we have $N(\omega, \omega_0) = 20$, which is the number of labor groups multiplied by the number of equipment types. Under the assumption that $\hat{\pi}(\lambda, \kappa, \omega)$ are affected by multiplicative error terms, the two expressions above yield consistent estimators.

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, because of measurement error, in the expression above the estimate of changes in relative transformed task prices crucially depends on the choice of ω_0 . In order to avoid this sensitivity to the choice of ω_0 , we compute relative changes in relative transformed task prices using the following average

$$\frac{\hat{p}(\omega)^\theta}{\hat{p}(\omega_1)^\theta} = \frac{1}{20} \sum_{\omega_0} \frac{\hat{p}(\omega)^\theta}{\hat{p}(\omega_0)^\theta} \left(\frac{\hat{p}(\omega_1)^\theta}{\hat{p}(\omega_0)^\theta} \right)^{-1}$$

where $\frac{\hat{p}(\omega)^\theta}{\hat{p}(\omega_0)^\theta}$ for each ω and ω_0 is calculated as described above. This expression yields estimates that do not depend on the choice of ω_1 .

Given our measures of changes in equipment and transformed task prices (both to the power θ), we construct $S(\lambda)$ using equation (17). Given $S(\lambda)$, we estimate θ using equation (18) via OLS. Finally, given our estimate of θ we measure changes in labor productivity using equation (16) and we measure changes in task shifters using equation (15) and values of α and ρ . In constructing task shifters using equation (15), we construct $\frac{\hat{\zeta}(\omega)}{\hat{\zeta}(\omega_1)}$ 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

$$\frac{\zeta_{t_1}(\omega)}{\zeta_{t_1}(\omega_1)} = \frac{\sum_{\lambda, \kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega)}{\sum_{\lambda', \kappa'} w_{t_0}(\lambda') L_{t_0}(\lambda') \pi_{t_0}(\lambda', \kappa', \omega_1) \frac{\hat{w}(\lambda')}{\hat{w}(\lambda'_1)} \hat{L}(\lambda') \hat{\pi}(\lambda', \kappa', \omega_1)}.$$

B.2 Elasticities vary across worker group

In the general case in which $\theta(\lambda)$ varies across λ and we know the value of $\theta(\lambda)$ for each λ , we can measure shocks as follows. Instead of equation (13), we have

$$\frac{\hat{T}(\kappa_2)}{\hat{T}(\kappa_1)} = \left(\frac{\hat{\pi}(\lambda, \kappa_2, \omega)}{\hat{\pi}(\lambda, \kappa_1, \omega)} \right)^{1/\theta(\lambda)}$$

for every (λ, ω) pair. We therefore estimate these relative changes in capital shifters as

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

Similarly, instead of equation (14), we measure

$$\frac{\hat{p}(\omega)}{\hat{p}(\omega_0)} = \left(\frac{\hat{\pi}(\lambda, \kappa, \omega)}{\hat{\pi}(\lambda, \kappa, \omega_0)} \right)^{1/\theta(\lambda)}$$

for every (λ, κ) pair. Hence, we estimate the changes in transformed task prices relative to each occupation ω_0 as

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

We then construct $\frac{\hat{p}(\omega)}{\hat{p}(\omega_1)}$ that does not depend on the choice of ω_1 exactly as in our baseline. Given these measures, constructing changes in task shifters and labor productivity is straightforward.

C Compensating differentials

Our baseline model implies that the average wage of workers in group λ is the same across all equipment-task pairs. Here we extend our model to incorporate preference heterogeneity in addition to productivity heterogeneity. This simple extension implies that the average wage of workers in group λ varies across equipment-task pairs. We show how to use data on average wages across equipment-task pairs to identify the parameters of the extended model and how to use these parameters to conduct an extended decomposition.

C.1 Environment and equilibrium

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

$$U(z, \kappa, \omega) = I_t(z) u_t(\lambda, \kappa, \omega) \quad (21)$$

where $u_t(\lambda, \kappa, \omega) > 0$ is a time-varying preference shifter.³² 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 $(\lambda, \kappa, \omega)$.

A task production unit hiring k units of equipment κ and l efficiency units of labor λ earns profits $p_t(\omega) k^\alpha [T_t(\lambda, \kappa, \omega) l]^{1-\alpha} - p_t(\kappa) k - W_t(\lambda, \kappa, \omega) l$. The profit maximizing choice of equipment quantity and the zero profit condition yield

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

if there is positive entry in $(\lambda, \kappa, \omega)$. Facing the wage profile $W_t(\lambda, \kappa, \omega)$, each worker $z \in \mathcal{Z}(\lambda)$

³²In 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 .

chooses (κ, ω) to maximize her indirect utility, $\varepsilon_t(z, \kappa, \omega) u_t(\lambda, \kappa, \omega) W_t(\lambda, \kappa, \omega)$.

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, as shown by the previous expression. Hence, they also affect worker allocation in the same way: the probability that a randomly sampled worker, $z \in \mathcal{Z}(\lambda)$, uses equipment κ in task ω is

$$\pi_t(\lambda, \kappa, \omega) = \frac{\left[u_t(\lambda, \kappa, \omega) T_t(\lambda, \kappa, \omega) p_t(\kappa)^{\frac{-\alpha}{1-\alpha}} p_t(\omega)^{\frac{1}{1-\alpha}} \right]^{\theta(\lambda)}}{\sum_{\kappa', \omega'} \left[u_t(\lambda, \kappa', \omega') T_t(\lambda, \kappa', \omega') p_t(\kappa')^{\frac{-\alpha}{1-\alpha}} p_t(\omega')^{\frac{1}{1-\alpha}} \right]^{\theta(\lambda)}}. \quad (22)$$

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 task ω is now given by

$$w_t(\lambda, \kappa, \omega) = \frac{\gamma(\lambda)}{u_t(\lambda, \kappa, \omega)} \left(\sum_{\kappa', \omega'} \left[u_t(\lambda, \kappa', \omega') T_t(\lambda, \kappa', \omega') p_t(\kappa')^{\frac{-\alpha}{1-\alpha}} p_t(\omega')^{\frac{1}{1-\alpha}} \right]^{\theta(\lambda)} \right)^{1/\theta(\lambda)} \quad (23)$$

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 conditions are identical to our baseline model and are given by equations (6) and (7), although total labor income in task ω is now given by

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

C.2 Parameterization

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

Preference shifters. From equation (23), we have

$$\frac{w_t(\lambda, \kappa, \omega)}{w_t(\lambda, \kappa_1, \omega_1)} = \frac{1}{u_t(\lambda, \kappa, \omega)}. \quad (24)$$

Hence, we measure preference shifters directly from average wages.

Equipment productivity and task shifters. Equations (8) and (22) give us,

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

which, together with equation (24), gives us

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

Hence, we obtain

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

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 task prices to the power θ similarly and use this to measure changes in task shifters using equation (15), as in our baseline approach. Finally, we have

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

so that

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

Hence, given measures of changes in transformed task prices (to the power θ), changes in equipment productivity (to the power θ), and changes in preference shifters (obtained above) 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 .

D Evolving comparative advantage details

Here we study case 2, where

$$T_t(\lambda, \kappa, \omega) \equiv T_t(\omega) T_t(\lambda, \kappa) T(\lambda, \kappa, \omega). \quad (25)$$

Cases 1 and 3 are similar and available upon request.

The equilibrium conditions are unchanged: equations (3), (5), (6), and (7) all hold as in our baseline model. However, we can re-express the system in changes as follows. Defining

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

equations (9) and (11) become

$$\hat{w}(\lambda) = \left(\sum_{\kappa, \omega} \pi_{t_0}(\lambda, \kappa, \omega) \left(\hat{T}_t(\lambda, \kappa) \hat{p}(\omega) \right)^{\theta(\lambda)} \right)^{1/\theta(\lambda)} \quad (26)$$

$$\hat{\pi}(\lambda, \kappa, \omega) = \frac{\left(\hat{p}(\omega) \hat{T}(\lambda, \kappa) \right)^{\theta(\lambda)}}{\sum_{\kappa', \omega'} \left(\hat{p}(\omega') \hat{T}(\lambda, \kappa') \right)^{\theta(\lambda)} \pi_{t_0}(\lambda, \kappa', \omega')}, \quad (27)$$

whereas equation (12) remains unchanged. Expressing equation (26) in relative terms yields

$$\frac{\hat{w}(\lambda)}{\hat{w}(\lambda_1)} = \frac{\hat{T}(\lambda, \kappa_1)}{\hat{T}(\lambda_1, \kappa_1)} \frac{\left(\sum_{\kappa, \omega} \pi_{t_0}(\lambda, \kappa, \omega) \left(\frac{\hat{T}(\lambda, \kappa)}{\hat{T}(\lambda, \kappa_1)} \frac{\hat{p}(\omega)}{\hat{p}(\omega_1)} \right)^{\theta(\lambda)} \right)^{1/\theta(\lambda)}}{\left(\sum_{\kappa', \omega'} \pi_{t_0}(\lambda_1, \kappa', \omega') \left(\frac{\hat{T}(\lambda_1, \kappa')}{\hat{T}(\lambda_1, \kappa_1)} \frac{\hat{p}(\omega')}{\hat{p}(\omega_1)} \right)^{\theta(\lambda_1)} \right)^{1/\theta(\lambda_1)}} \quad (28)$$

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

Here we provide an overview—similar in structure to that provided in Section 3.3—of how we measure shocks taking as given the parameters α, ρ , and θ . Equations (3) and (25) give us

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

for each λ and ω . Hence, we can measure $\hat{T}(\lambda, \kappa_1)^\theta / \hat{T}(\lambda, \kappa_2)^\theta$ for each λ as the exponential of the average across ω of the log of the right-hand side of the previous expression. This procedure yields consistent estimates of the changes in technology in the presence of multiplicative measurement error in the observed allocations. We can recover changes in transformed task prices to the power θ and use these to measure changes in task shifters exactly as in our baseline. Finally, given these measures, we can recover $\hat{T}(\lambda, \kappa_1)^\theta / \hat{T}(\lambda_1, \kappa_1)^\theta$ to match changes in relative wages using equation (28).

From equation (28) it might appear possible to divide the labor-equipment component of the decomposition into two separate parts: an equipment component, $\frac{\hat{T}(\lambda, \kappa)}{\hat{T}(\lambda, \kappa_1)}$, and a labor component, $\frac{\hat{T}(\lambda, \kappa_1)}{\hat{T}(\lambda_1, \kappa_1)}$. However, whereas the choice of κ_1 is irrelevant for quantitative results on the labor composition, task shifter, and labor-equipment components, the choice of κ_1 has large effects if we attempt to decompose the labor-equipment component into separate equipment productivity and labor productivity components.

Whereas in case 2 the choice of κ_1 is irrelevant for all quantitative results in Table 8, the same is not true of the similar choice of λ_1 in case 1 or of κ_1 in case 3.³³ In practice, however, the dependence of our results on these choices is very small. As in case 2, if we try to decompose

³³Specifically, in case 1 the choice of λ_1 does affect the labor-task component of the decomposition (although it does not affect the labor composition or equipment productivity components). Similarly, in case 3 the choice of κ_1 does affect the equipment-task and the labor productivity components of the decomposition (although it does not affect the labor composition component).

the labor-task component into separate task shifter and labor productivity components in case 1 or the equipment-task component in case 3 into separate equipment productivity and task shifter components in case 3, then the choice of λ_1 and κ_1 , respectively, has large effects on these results. In the results presented in Table 8, we chose λ_1 as the young, male, HSD group in case 1 and we chose κ_1 as computers in case 3.

E Model with sectors: Details

Here we provide additional details on the closed economy extension with sectors. 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 (5), respectively. The only change to the equilibrium equations (in levels) is to the task market clearing condition, which becomes

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

where $E_t(\omega, \sigma)$ denotes income (or expenditure) on task 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

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

We now provide the system of equations in changes, analogous to equations (9), (11), and (12) 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 (9) and (11), respectively. The right hand side of the task market clearing condition in changes is the same as in the baseline model,

$$\hat{\zeta}(\omega) = \frac{1}{\zeta_{t_0}(\omega)} \sum_{\lambda, \kappa} w_{t_0}(\lambda) L_{t_0}(\lambda) \pi_{t_0}(\lambda, \kappa, \omega) \hat{w}(\lambda) \hat{L}(\lambda) \hat{\pi}(\lambda, \kappa, \omega). \quad (29)$$

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

$$\hat{p}(\omega)^{1-\rho} \hat{E} \sum_{\sigma} v_{t_0}(\sigma|\omega) \hat{\mu}(\sigma) \hat{\mu}(\omega, \sigma) \hat{p}(\sigma)^{\rho-\rho\sigma} = \hat{\zeta}(\omega) \quad (30)$$

and

$$\hat{p}(\sigma) = \left[\sum_{\omega} v_{t_0}(\omega|\sigma) \hat{\mu}(\omega, \sigma) \hat{p}(\omega)^{1-\rho} \right]^{\frac{1}{1-\rho}} \quad (31)$$

Here, $v_t(\sigma|\omega) \equiv \frac{E_t(\omega, \sigma)}{\sum_{\sigma'} E_t(\omega, \sigma')}$ denotes the share of expenditure on task ω across all sectors that occurs within sector σ and $v_t(\omega|\sigma) = \frac{E_t(\omega, \sigma)}{\sum_{\omega'} E_t(\omega', \sigma)}$ denotes the share of expenditure on task ω across all tasks employed within sector σ . Defining, as in the baseline model, changes in transformed task shifters $\hat{\mu}(\omega, \sigma) = \hat{\mu}(\omega) \hat{T}(\omega)^{(1-\alpha)(\rho-1)}$ and changes in transformed task prices $\hat{\tilde{p}}(\omega) = \hat{p}(\omega)^{1/(1-\alpha)} \hat{T}(\omega)$ we can write (30) and (31) as

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

and

$$\hat{p}(\sigma) = \hat{\mu}(\omega_1, \sigma)^{\frac{1}{1-\rho}} \left[\sum_{\omega} v_{t_0}(\omega|\sigma) \frac{\hat{\mu}(\omega, \sigma)}{\hat{\mu}(\omega_1, \sigma)} \hat{\tilde{p}}(\omega)^{(1-\rho)(1-\alpha)} \right]^{\frac{1}{1-\rho}}$$

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

$$\hat{\tilde{p}}(\omega)^{(1-\alpha)(1-\rho)} \hat{E} \sum_{\sigma} v_{t_0}(\sigma|\omega) \hat{\tilde{\mu}}(\sigma) \frac{\hat{\mu}(\omega, \sigma)}{\hat{\mu}(\omega_1, \sigma)} \hat{\tilde{p}}(\sigma)^{\rho-\rho\sigma} = \hat{\zeta}(\omega) \quad (32)$$

$$\hat{\tilde{p}}(\sigma) = \left[\sum_{\omega} v_{t_0}(\omega|\sigma) \frac{\hat{\mu}(\omega, \sigma)}{\hat{\mu}(\omega_1, \sigma)} \hat{\tilde{p}}(\omega)^{(1-\rho)(1-\alpha)} \right]^{\frac{1}{1-\rho}} \quad (33)$$

Therefore we can solve for changes in relative wages $\hat{w}(\lambda) / \hat{w}(\lambda_1)$, relative task prices $\hat{\tilde{p}}(\omega) / \hat{\tilde{p}}(\omega_1)$ and relative sectoral prices $\hat{\tilde{p}}(\sigma) / \hat{\tilde{p}}(\sigma_1)$ using equations (9), (11), (29), (32) and (33), given shocks $\hat{L}(\lambda) / \hat{L}(\lambda_1)$, $\hat{T}(\lambda) / \hat{T}(\lambda_1)$, $\hat{T}(\kappa) / \hat{T}(\kappa_1)$, $\hat{\mu}(\omega, \sigma) / \hat{\mu}(\omega_1, \sigma)$ and $\hat{\mu}(\sigma) / \hat{\mu}(\sigma_1)$. Note that when $\rho = \rho_{\sigma}$, $\hat{\tilde{p}}(\sigma)$ drops out from equation (32) and the extension of our baseline model that accounts for sectors is equivalent to our baseline model where the task shifter in our baseline model, $\hat{\mu}(\omega)$, is replaced by $\sum_{\sigma} v_{t_0}(\sigma|\omega) \hat{\tilde{\mu}}(\sigma) \frac{\hat{\mu}(\omega, \sigma)}{\hat{\mu}(\omega_1, \sigma)}$.

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

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

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

$$\frac{\hat{E}(\omega, \sigma)}{\hat{E}(\omega_1, \sigma)} = \frac{\hat{\mu}(\omega, \sigma)}{\hat{\mu}(\omega_1, \sigma)} \left(\frac{\hat{p}(\omega)}{\hat{p}(\omega_1)} \right)^{(1-\alpha)(1-\rho)}. \quad (34)$$

We use equation (34) to back-out $\hat{\mu}(\omega, \sigma) / \hat{\mu}(\omega_1, \sigma)$. To estimate $\frac{\hat{\mu}(\sigma)}{\hat{\mu}(\sigma_1)}$, we start from the equilibrium relationship

$$\frac{\hat{E}(\sigma)}{\hat{E}(\sigma_1)} = \frac{\hat{\mu}(\sigma)}{\hat{\mu}(\sigma_1)} \left(\frac{\hat{p}(\sigma)}{\hat{p}(\sigma_1)} \right)^{1-\rho_\sigma} \quad (35)$$

or in terms of transformed variables

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

The previous expression and equation (33) yield

$$\frac{\hat{E}(\sigma)}{\hat{E}(\sigma_1)} = \frac{\hat{\mu}(\sigma)}{\hat{\mu}(\sigma_1)} \left(\frac{\sum_{\omega} \nu_{t_0}(\omega|\sigma) \frac{\hat{\mu}(\omega, \sigma)}{\hat{\mu}(\omega_1, \sigma)} \hat{p}(\omega)^{(1-\rho)(1-\alpha)}}{\sum_{\omega'} \nu_{t_0}(\omega'|\sigma_1) \frac{\hat{\mu}(\omega', \sigma_1)}{\hat{\mu}(\omega_1, \sigma_1)} \hat{p}(\omega')^{(1-\rho)(1-\alpha)}} \right)^{\frac{1-\rho_\sigma}{1-\rho}} \quad (36)$$

We use equation (36) to back-out $\hat{\mu}(\sigma) / \hat{\mu}(\sigma_1)$.

F Model with international trade: Details

Here we derive the system of equations (analogous to equations (9), (11), and (12) 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 presented in Section 6.1 and as detailed in Appendix E. Here, however, changes in equipment productivity, sector shifters, and within-sector task 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 task 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 : $\hat{y} = y^A / y_{t_0}$. For simplicity, in this section we omit time indices.

In an open economy, we must distinguish between production prices and absorption prices. For example, we denote by $p_{in}(\omega)$ the price of country i 's output of task ω in country n (inclusive

of trade costs) and by $p_n(\omega)$ the absorption price of task ω in country n , given by

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

output prices $p_{in}(\sigma)$, $p_{in}(\kappa)$, and absorption prices $p_n(\sigma)$, and $p_n(\kappa)$ are defined analogously.

Changes in relative wages and in allocations depend on changes in absorption prices for equipment (since it is an input in production) and production prices for tasks (since tasks are produced in each country). Since productivities are assumed constant when moving to autarky, we set $\hat{T}(\kappa) = \hat{p}_n(\kappa)^{\frac{-\alpha}{1-\alpha}}$ and $\hat{p}(\omega) = \hat{p}_{nn}(\omega)^{\frac{1}{1-\alpha}}$, and equations (9) and (11) (moving to autarky) become

$$\hat{w}_n(\lambda) = \left\{ \sum_{\kappa, \omega} \left[\hat{p}_{nn}(\omega)^{\frac{1}{1-\alpha}} \hat{p}_n(\kappa)^{\frac{-\alpha}{1-\alpha}} \right]^{\theta(\lambda)} \pi_n(\lambda, \kappa, \omega) \right\}^{1/\theta(\lambda)} \quad (37)$$

$$\hat{\pi}_n(\lambda, \kappa, \omega) = \frac{\left[\hat{p}_{nn}(\omega)^{\frac{1}{1-\alpha}} \hat{p}_n(\kappa)^{\frac{-\alpha}{1-\alpha}} \right]^{\theta(\lambda)}}{\sum_{\omega', \kappa'} \left[\hat{p}_{nn}(\omega')^{\frac{1}{1-\alpha}} \hat{p}_n(\kappa')^{\frac{-\alpha}{1-\alpha}} \right]^{\theta(\lambda)} \pi_n(\lambda, \kappa', \omega')} \quad (38)$$

The remaining equations are the open-economy versions of the task-market clearing condition (12) (the analog of equation (30)) and the sectoral price index (the analog of equation (31)).

The right-hand side of equation (12) remains unchanged, so we focus on the left-hand side only. The level of worldwide absorption expenditure on country n 's produced task ω is $\sum_i E_{ni}(\omega)$, where $E_{ni}(\omega)$ denotes country i 's absorption expenditure of task 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 (12)) is

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

where $f_{nn}(\omega) = \frac{E_{nn}(\omega)}{\sum_i E_{ni}(\omega)}$ denotes the share of domestic sales of task 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 task ω in country n . The level of $E_n(\omega)$ is given by

$$E_n(\omega) = \sum_{\sigma} E_n(\omega, \sigma) \quad (39)$$

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

$$E_n(\omega, \sigma) = \mu_n(\omega, \sigma) \left(\frac{p_n(\omega)}{p_{nn}(\sigma)} \right)^{1-\rho} Y_n(\sigma) p_{nn}(\sigma) \quad (40)$$

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

$$Y_n(\sigma) p_{nn}(\sigma) = \sum_i Y_{ni}(\sigma) d_{in}(\sigma) p_{nn}(\sigma) = \sum_i E_{ni}(\sigma) \quad (41)$$

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

$$E_{ni}(\sigma) = \mu_i(\sigma) \left(\frac{p_{ni}(\sigma)}{p_i(\sigma)} \right)^{1-\eta(\sigma)} \left(\frac{p_i(\sigma)}{p_i} \right)^{1-\rho_\sigma} E_i \quad (42)$$

and E_i denotes total expenditures on the final good in country i . Combining equations (40), (41), and (42) 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 is the defined analogously to $f_{nn}(\omega)$. Equation (39) therefore yields

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

Combining these results, we have

$$\frac{\sum_i E_{ni}^A(\omega)}{\sum_i E_{ni}(\omega)} = f_{nn}(\omega) \hat{p}_n(\omega)^{1-\rho} \left(\frac{\hat{p}_{nn}(\omega)}{\hat{p}_n(\omega)} \right)^{1-\eta(\omega)} \sum_{\sigma} v_n(\sigma|\omega) \hat{p}_{nn}(\sigma)^{\rho-\rho_\sigma} \left(\frac{\hat{p}_n(\sigma)}{\hat{p}_{nn}(\sigma)} \right)^{\eta(\sigma)-\rho_\sigma} f_{nn}(\sigma) \hat{E}_n \quad (43)$$

where the change in the production sectoral price index is

$$\hat{p}_{nn}(\sigma) = \left(\sum_{\omega} v_n(\omega|\sigma) \hat{p}_n(\omega)^{1-\rho} \right)^{1/(1-\rho)}. \quad (44)$$

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 task ω , for example, is

$$\hat{p}_n(\omega) = \frac{p_{nn}^A(\omega) / p_{nn}(\omega)}{\left(\sum_i (p_{in}(\omega) / p_{nn}(\omega))^{1-\eta(\omega)}\right)^{\frac{1}{1-\eta(\omega)}}} = \frac{\hat{p}_{nn}(\omega)}{s_{nn}(\omega)^{\frac{1}{\eta(\omega)-1}}} \quad (45)$$

where $s_{nn}(\omega)$ denotes expenditure on domestic task ω relative to total expenditure on task 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 (45) 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

$$\hat{p}_n(\sigma) = \frac{\hat{p}_{nn}(\sigma)}{s_{nn}(\sigma)^{\frac{1}{\eta(\sigma)-1}}} \quad (46)$$

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

$$\hat{p}_n(\kappa) = s_{nn}(\kappa)^{\frac{-1}{\eta(\kappa)-1}}, \quad (47)$$

where $s_{nn}(\kappa)$ is defined analogously to $s_{nn}(\omega)$ and where we have used the fact that $\hat{p}_{nn}(\kappa) = 1$ given our choice of numeraire.

We can substitute equation (47) directly into equations (37) and (38). Similarly, substituting (45) and (46) into (43) and (44) we have

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

and

$$\hat{p}_{nn}(\sigma) = \left(\sum_{\omega} v_n(\omega|\sigma) s_{nn}(\omega)^{\frac{\rho-1}{\eta(\omega)-1}} \hat{p}_{nn}(\omega)^{1-\rho} \right)^{1/(1-\rho)}. \quad (48)$$

In sum, the system of equations to solve for changes in factor allocations and relative prices

when moving to autarky is given by

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

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

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

and

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

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 task shifters, $\hat{\mu}_n(\omega, \sigma)$, and between-sector shifters, $\hat{\mu}_n(\sigma)$, are equal to

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

and changes in equipment costs are equal to

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