

Online Appendix

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

Ariel Burstein
UCLA

Eduardo Morales
Princeton

Jonathan Vogel
UCLA

E Robustness and sensitivity analyses

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

E.1 Alternative parameter values

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

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

Our main results are robust for all these alternative estimates of the parameter vector (θ, ρ) . First, computerization is the most important force accounting for the rise in between-education-group inequality between 1984 and 2003. Alone, it accounts for between roughly 45% and 94% of the demand-side forces raising the skill premium. Second, residual labor productivity accounts for no more than one-third of the demand-side forces raising the skill premium. Third, changes in each of the demand-side forces play an im-

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

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

portant role in accounting for the reduction in the gender gap between 1984 and 2003. Specifically, changes in labor productivity play a larger role in accounting for the reduction in the gender gap than in accounting for the rise in the skill premium. Alone, they account for between roughly 20% and 46% of the demand-side forces reducing the gender gap.

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

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

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

As is evident from rows two and three of Table 12, a higher value of θ implies a larger role for changes in labor productivity and a smaller role for the other shocks in accounting for changes in the skill premium and the gender gap. The intuition is straightforward. According to equation (17), the elasticity of changes in average wages of workers in labor

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

group λ , $\hat{w}(\lambda)$, to changes in the measured labor-group-specific average of equipment productivities and transformed occupation prices (both to the power θ), $\hat{s}(\lambda)$, is $1/\theta$. Because our measure of $\hat{s}(\lambda)$ is independent of θ , a higher value of θ reduces the impact on wages of changes in the labor-group-specific average of equipment productivities and transformed occupation prices and, therefore, increases the impact of changes in labor productivity, identified as a residual to match observed changes in average wages.

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

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

E.2 Sources of comparative advantage

To demonstrate the importance of including each of the three forms of comparative advantage, we perform two exercises. We first assume there is no comparative advantage related to occupations and then we redo the decomposition under the assumption that there is no comparative advantage related to equipment. In all cases, we hold the values of α , ρ , and θ fixed to the same values employed in Section 4.4.

Table 13 reports our baseline decomposition between 1984 and 2003 both for the skill premium (in the left panel) and the gender gap (in the right panel) as well as decompo-

sitions under the restriction that there is comparative advantage only between labor and equipment or only between labor and occupations.

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

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

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

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

Similarly, assuming there is no comparative advantage at the level of equipment implies that the equipment productivity component of our decomposition is zero and that the only force giving rise to the allocation of labor groups to occupations is worker-occupation comparative advantage. Row 3 of Table 13 shows that abstracting from any comparative advantage at the level of equipment magnifies the importance of labor productivity in explaining the rise of the skill premium and the fall in the gender gap. The impact of occupation shifters on the gender gap does not change significantly.

In summary, abstracting from comparative advantage at the level of either occupations or equipment has a large impact on the decomposition of changes in between-group in-

equality. It does so by forcing changes in labor productivity to absorb the impact of the missing component(s) and by changing the importance of the remaining source of comparative advantage.

E.3 Evolving comparative advantage

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

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

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

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

Our results are largely unchanged and the intuition for why is straightforward in cases 1 and 2. In all three cases, our measures of initial factor allocations and changes in labor composition as well as the system of equations that determines the impact of changes in labor composition on relative wages are exactly the same as in our baseline model. Hence, the labor composition component of our baseline decomposition is unchanged if we incorporate time-varying comparative advantage. Similarly, our measure of changes

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

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

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

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

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

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

The equilibrium conditions are unchanged: equations (3), (4), and (5) hold as in our base-

line model. However, we can re-express the system in changes as follows. Defining

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

equations (11) and (12) become

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

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

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

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

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

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

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

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

E.4 Letting θ vary by λ

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

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

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

and

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

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

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

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

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

F Additional exercises

F.1 Occupation characteristics

As documented in Figure 1 and in Table 10 in Appendix B, certain occupations—including, for example, executive, administrative, managerial as well as health assessment and treating—have grown disproportionately over the last three decades. As discussed in, e.g., Autor et al. (2003), these changes have been systematically related to the task content of each occupation; for example, there has been an expansion of occupations intensive in non-routine cognitive analytical and non-routine cognitive interpersonal tasks and a corresponding contraction in occupations intensive in routine manual and non-routine manual physical tasks, as defined following Acemoglu and Autor (2011).

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

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

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

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

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

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

F.2 Worker aggregation

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

Our objective here is to understand the extent to which our particular disaggregation may be driving our results. To do so, we decrease the number of labor groups from 30 to 10 by dropping age as a characteristic. In this case, the share of $(\lambda, \kappa, \omega)$ observations for which $\pi_t(\lambda, \kappa, \omega) = 0$ falls from (roughly) 26 percent to 12 percent. Because, in our baseline, we composition adjust the skill premium and the gender gap using gender, education, and age, whereas here we only use gender and education, we find slightly different changes in the skill premium, 16.1 instead of 15.1 percent, and the gender gap, -13.2 instead of -13.3 percent, between 1984 and 2003.

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

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

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

robust to decreasing the number of labor groups.

G Decompositions

G.1 Computer use decomposition

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

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

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

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

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

occupations at common (average) computer usage within each occupation. We refer to the second term of the above decomposition as the within component, since it reflects the share of the difference in aggregate computer usage between λ' and λ that reflects differences in computer usage within each occupation at common (average) allocations across occupations.

G.2 Between-within wage decomposition

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

The following accounting identity must hold at each t ,

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

and, therefore, we can write

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

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

H International trade in sectoral goods

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Absorption of sector σ in country n is

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

where the change in the output price of sector σ is

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

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

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

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

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

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

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

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

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

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

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

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

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

(36), then

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

which coincides with expression (60).