Structural Transformation by Cohort

Bart Hobijn* Todd Schoellman† Alberto Vindas Q.‡

November 16, 2018

Preliminary version – Comments welcome

Abstract

More than half of labor reallocation during structural transformation can be attributed to new cohorts of workers disproportionately entering growing industries. This finding suggests substantial costs to reallocating workers across industries. We integrate an overlapping generations model of life-cycle career choice under switching costs with a canonical model of structural transformation. Switching costs accelerate structural transformation because they cause forward-looking workers to enter growing industries in anticipation of future productivity and wage growth. Most of the impact of switching costs, however, is on the trends in sectoral relative wages. An unanticipated acceleration of structural transformation makes more young workers line up in the service sector and reduces service sector wages. This comes disproportionately at the cost of expected future career earnings of older service sector workers.

*Arizona State University. E-mail: bhobijn@asu.edu

†Opportunity and Inclusive Growth Institute, Federal Reserve Bank of Minneapolis. E-mail: todd.schoellman@gmail.com (corresponding author)

‡Arizona State University. E-mail: avindasq@asu.edu

1 Introduction

One of the key stylized facts of economic growth is that it involves structural transformation: the reallocation of economic activity in predictable ways among the broad industries of an economy. Whereas poor countries typically produce and consume a high share of agriculture, growth entails a shift towards first manufacturing and then services.[1] A recent literature has explored different forms of preferences and technological progress that can generate this predictable reallocation of economic activity as a consequence of growth (Kongsamut et al., 2001; Ngai and Pissarides, 2007).

Although the existing literature has advanced our understanding of structural transformation along many dimensions, it is largely silent about the interaction between structural transformation and labor markets, for two reasons. First, there are few stylized facts about structural transformation and labor markets.[2] Most papers use aggregate data from national accounts, which does little to clarify which workers are responsible for labor reallocation. Second, most papers focus on the special case of frictionless labor markets, which often allows for elegant analytical solutions but abstracts from the interactions we are interested in. Our goal is to make progress on both fronts: we document new stylized facts of what sorts of workers reallocate during structural transformation; we develop a model consistent with these findings and use it to help understand the relationship between labor markets and structural transformation.

The starting point for our empirical contribution is to document stylized facts about which workers reallocate across sectors during structural transformation. To do so we utilize nationally representative, repeated cross sections spanning a long time series for the United States and shorter time series for 60 other countries. By using repeated cross sections we can track reallocation based on observable characteristics such as education. Figure 1 gives a visual representation of our approach for the case of the United States, 1870–2010. The three panels plot the employment shares by birth cohort for each decadal census in which these workers were between the ages of 20 and 70. Each line plots the time series for the workers in a particular birth cohort with the dot at the beginning of the line showing their employment share when they enter the workforce.

[1] See for example Schultz (1953) and Echevarria (1997) for early references, or Herrendorf et al. (2014) for a recent overview. Herrendorf et al. (2014) shows that structural transformation is a predictable function of PPP GDP p.c.

[2] There is however a large and related literature on gross worker flows between industries. See for example Kambourov and Manovskii (2008) for the U.S. and Carrillo-Tudela et al. (2016) for the United Kingdom.

The overall pattern in the figure clearly shows the decline in the employment share of agriculture, the rise and then fall of the manufacturing share, and the increasing share of services over time. The second important point of this figure is that the lines for the individual cohorts do not overlap. Within a given year, newer and older cohorts have different employment patterns. In particular, the lines for each cohort appear to be "flatter" than the pattern for the overall economy. This is a very informal way of saying that within-cohort shifts in employment shares tend to be smaller than those in the overall economy and that differences in employment shares between cohorts are an important part of the sectoral reallocation of labor that has occurred in the United States. We formalize this finding using an accounting decomposition, which shows that 53 percent of reallocation happens between cohorts, both in the United States and in our international sample.[3] We also show that much of the within cohort (life-cycle) reallocation happens at earlier ages.

These findings suggest that new cohorts play a central role in the process of structural transformation. They lead us to formulate a heterogeneous agent overlapping generations model of life-cycle career choice under switching costs and integrate it into a canonical model of structural transformation.[4] The idea is that switching costs prevent older workers from moving between industries and hence give a prominent role to new cohorts for generating structural transformation. In doing so, we deviate from the common assumption of a single frictionless labor market that is made in many growth models, including those with structural transformation. In such a labor market, wages are equated across sectors, which runs counter to empirical evidence (Kim and Topel, 1995; Herrendorf and Schoellman, 2018).

An important contribution of our modeling strategy is to show how to formulate this problem in a tractable way. This is challenging at three levels. First, we need a tractable way to characterize the life-cycle career path of each individual agent as a function of wages. Second, wages themselves depend on the labor supply of past and future cohorts, which requires us to find a way to iterate on the entire path of labor supply and wages. Finally, our labor markets are part of structural transformation, which implies that the economy is experiencing unbalanced growth. We show how to overcome these challenges in our

[3]Earlier authors have documented similar patterns for specific cases: Kim and Topel (1995) in Korea and Perez (2016) for Argentina. Porzio and Santangelo (2017) document similar facts for a large set of countries similar to ours for reallocation out of agriculture. We adopt the format of figure 1 from their work.

[4]We use the life-cycle career choice to mean the sequence of industries of employment. Duernecker and Herrendorf (2017) document a close linkage between reallocation across industries and reallocation across occupations during structural transformation.

quantitative implementation.

At the aggregate we formulate structural transformation as in Ngai and Pissarides (2007). Differential technology growth across sectors and a low elasticity of substitution in the utility function generate trends in both relative prices of goods produced in different sectors as well as the relative levels of labor demand across sectors. This formulation of structural transformation is useful for our purposes because it relies on homothetic preferences, which implies that we can solve for relative consumption as a function of only relative prices and not the entire distribution of income.

Our main theoretical contribution is to integrate the life-cycle career choice of workers. Workers decide on their labor supply taking into account their idiosyncratic sector-specific skills, as in Lagakos and Waugh (2013) and in Bárány and Siegel (2018), the current and future wages in each of the sectors, and the current and expected future retraining costs, as in Caselli and Coleman (2001), associated with changing sectors of employment.[5] We deviate from previous authors by formulating the career choice problem as a dynamic discrete choice problem. Doing so allows us to utilize known closed form solutions for life cycle sector choice and labor supply and avoid numerical integration, which greatly reduces the scale of the problem. Finally, we show how to adopt the extended path method of Fair and Taylor (1983) to this environment with unbalanced growth and solve for the equilibrium path of our model.

Our model provides four important insights. The first is that retraining costs for workers *accelerate* structural transformation. The reason for this counterintuitive result is the following. Given an initial sectoral allocation of labor, retraining costs slow sectoral reallocation down. However, forward-looking workers who face training and retraining costs change their initial labor allocation and shift their labor supply towards growing industries in anticipation of the future productivity and wage growth. This shift in initial sectoral choice more than compensates for lower life cycle sectoral reallocation.

The second result is that retraining costs need to increase with age to match cohort career profiles. To understand this, note that the model embeds an option value to working in a growing industry because of expected future relative wage growth. However, this option value declines as workers get older and have a lower expected length of their future career. As a result workers with a high idiosyncratic opportunity in a declining industry are more

[5] Cociuba and MacGee (2018) also consider sectoral adjustment costs of workers, but do so in a stationary model with search frictions that is suitable at business cycle frequencies but does not allow for the analysis of the long-run trends in structural transformation that we consider here.

inclined to switch industries as they grow older, contrary to the data. In order to prevent these switches, retraining costs need to be increasing in age to offset the decline in the option value of being in a growing industry.

The third insight is that most of the impact of retraining costs is the on the trends in relative wages across sectors and not on the shares of workers employed. This is due to the aggregate technology being parameterized as near-Leontief to be consistent with historical trends in value added shares and relative prices in the U.S. (Ngai and Pissarides, 2008).

The final insight is that because more workers will line up in the service sector when structural transformation accelerates unexpectedly, such an acceleration reduces wages in the service sector in the decades directly after. This reduction disproportionately affects the career earnings outlook of older workers in the service sector when the shock hits. In the longer-run the shock reduces relative wages in agriculture and manufacturing.

2 Cohort Effects and Structural Transformation

In this section we document stylized facts of worker reallocation across industries that motivate our model in Section 3. We focus our attention on a classic three-industry view of the U.S. time series, with some additional results from a large international sample presented for comparison. Details of the data construction and results from alternative industry decompositions are reserved to the appendix.

Our baseline analysis uses the United States census microdata spanning 1870–2010, taken from IPUMS (Ruggles et al., 2010). We study the structural transformation of employment, which is constructed using the reported industry of employed workers with valid responses. IPUMS has devoted substantial effort to harmonizing responses to these and other key variables over time and across countries. We aggregate detailed industry classifications to the standard three broad industry groups: agriculture, manufacturing, and services.[6] We impose no other sample restrictions, because we want the results derived from microdata to be consistent with aggregate trends.

Our main empirical finding is that much of structural transformation can be accounted for by new workers who enter growing sectors disproportionately. Figure 1 provides a visual representation of this finding for the United States. It plots sectoral employment shares against

[6] Agriculture includes all of farming, forestry, and fishing. Manufacturing includes also mining and construction. Services includes utilities.

time, with each individual line representing a distinct decade-of-birth cohort followed over their working life. Ignoring for a moment the distinct lines, the general employment patterns are clear: the decline of agriculture; the inverse-U shape of manufacturing; and the rise of services. The individual lines show that within a particular year younger cohorts had different employment patterns than older ones. For example, in 1900 the younger cohorts had about 15 percent lower employment shares in agriculture and correspondingly higher employment shares in manufacturing and services. The between-cohort gaps within a year provide visual evidence of the importance of cohort effects for structural transformation.

To document this pattern more carefully we utilize a within-between accounting decomposition. Denote by $e_t(i)$ sector i's employment share at time t and by $\Delta e_t(i)$ the change in sector i's employment share between $t-1$ and t. We decompose this total change into two pieces: the portion accounted for by changes in the employment share of the cohort who is age h at time t, $n_t(h)$; and the portion that is accounted for by the changes in the employment patterns of each cohort, $e_t(i;h)$. The usual decomposition holds:

$$\underbrace{\Delta e_t(i)}_{\text{total}} = \underbrace{\sum_h \overline{n_t(h)} \Delta e_t(i;h)}_{\text{within-cohort}} + \underbrace{\sum_h \Delta n_t(h) \overline{e_t(i;h)}}_{\text{between-cohort}}, \tag{1}$$

where Δ denotes differences between $t-1$ and t and bar denotes averages between $t-1$ and t.

We use variance-covariance accounting to perform the decomposition. The within and between shares are simply the covariance of the within and between terms with the total, relative to the variance of the total. This accounting procedure is identical to classical ANOVA. It can also be implemented in a straightforward way by taking the estimated coefficients from regressing the within and between components on the total component without a constant.

Table 1 shows the result of this decomposition. In the first row we show the results for the United States, where we study reallocation between each consecutive pair of censuses, usually taken a decade apart. In the first column we show the results from pooling all three industries. In this case, the between cohort share of structural transformation is just over half, meaning that a little more than half of structural transformation is accounted for by the propensity of new cohorts of workers to work in growing sectors. The remaining columns show the results separately for agriculture, manufacturing, and services. The between share is highest for agriculture and somewhat lower for manufacturing and services. These findings are consistent with the work of Kim and Topel (1995), who showed that between-cohort reallocation was central to the decline of agriculture in South Korea, and support the focus of Porzio and Santangelo (2017) on the role of cohorts for agriculture versus non-agriculture.

Although we focus on the United States, the underlying patterns are quite similar for the international sample, which includes 201 nationally representative surveys from 59 other countries, allowing us to decompose structural transformation across 142 consecutive survey pairs around the world. The results are shown in the second row of Table 1. The overall share of 53 percent is almost identical to the share for the United States. The shares by industry are also quite similar, with again a much larger role of the between share in agriculture.

Just over half of structural transformation is driven by the replacement of old cohorts by new ones. Further, much of the within-cohort reallocation happens early in the life cycle. To document this point, we exploit the fact that our accounting equation is additive in age. We then decompose the within share into the portion that happens within a cohort for those aged 20–29 at time $t-1$; those aged 30–39 at time $t-1$; and so on. The results are shown as solid circles for the United States and the international sample in Figure 2. A further 15–25 percent of all structural transformation happens from the 20s, with the pace of reallocation slowing with age, somewhat more rapidly for the international sample.

Examination of the within component in equation (1) shows that it can decline for two reasons: because the employment share of a cohort falls with age (falling $\overline{n_t(h)}$); or because cohorts are less likely to switch sectors as they age (falling $\Delta e_t(i;h)$). It is useful for our purposes to distinguish between the effects of the employment share versus the unweighted reallocation. The diamonds in Figure 2 plot the average employment share by age. The squares, plotted against the right axis, show the unweighted reallocation, which is the result of doing the same variance-covariance decomposition using only $\Delta e_t(i;h)$ as our measure of within. Although this no longer decomposes total reallocation, it does isolate the pure behavioral response. Indeed, we can see that both for the United States and the international sample the declining within share is driven primarily by a falling employment share by age. The unweighted reallocation effect is mixed: it falls in importance until age 40 or 50 before rebounding and becoming more important at older ages.

To summarize, the main contribution of our empirical work is to show that half of structural transformation seems to be accounted for by the fact that new cohorts disproportionately enter growing sectors. Much of the rest happens early in the life cycle, although this is driven more by demographics than by the behavioral responses of workers. These facts motivate us to write down a model of structural change where demographics and the employment choices of new workers play the central role.

One possible concern with our approach is that our between-cohort effects may proxy for other slow-moving trends that are the fundamental driving forces of structural transformation. For example, recent work has stressed the role of education and female labor force participation for structural transformation (Caselli and Coleman, 2001; Rendall, 2017; Buera et al., forthcoming; Ngai and Petrongolo, 2017). We test the importance of these factors by examining the share of structural transformation that happens within and between $gender \times marital\ status \times education$ groups. We use binary gender and marital status categories and four education bins, which produce a total of 16 possible cells. Table 2 shows the corresponding accounting results. 16 percent of structural transformation happens between demographic cells, which is a much lower share than our cohort results above.[7] This finding suggests to us that cohort is not simply a proxy for trends in education, female labor force participation, and so on. This suggests that new cohorts inherently account for much of structural transformation. We turn now to a model that captures this idea.

3 Structural Transformation with Career Choices

Our empirical findings suggest that new cohorts play a central role in the process of structural transformation. We now formulate a heterogeneous agent overlapping generations model of life-cycle career choice under switching costs and integrate it into a canonical model of structural transformation. In the next section we use the model to infer the nature of the adjustment costs and to perform several counterfactual exercises that highlight how structural transformation is a race between demographics and technology.

3.1 Households

Demographics and cohorts

Because this paper is about the interaction between structural transformation and demographics, we start by defining the demographic structure of our model economy. The economy consists of a unit measure of households, that are made up of members indexed

[7] Hendricks (2010) documents a similar facts for more detailed educational categories.

by age $h = 0, ..., H$. Each year, a new cohort of size $N_t(0)$ is born into each household. The growth rate of these new cohorts is $n > 0$; cohorts aged H die with certainty, and younger ones die with probability $0 \leq \delta < 1$. The resulting law of motion for cohort size by age is given by:

$$N_t(h) = \begin{cases} (1+n) N_{t-1}(0) & h = 0 \\ (1-\delta) N_{t-1}(h-1) & h = 1, \ldots, H \end{cases}. \quad (2)$$

The total size of the household (equivalently, total size of the population) is given by:

$$N_t = \sum_{h=0}^{H} N_t(h). \quad (3)$$

It also grows at rate n.

Preferences, consumption, and labor supply

The members of the household pool their income risk and maximize the present discounted value of the household's log consumption flow. The factor at which future consumption is discounted is β and this present discounted value equals

$$\sum_{t=0}^{\infty} \beta^t \ln C_t. \quad (4)$$

Here, following Ngai and Pissarides (2007), the aggregate consumption level C_t is a CES aggregate of consumption $C_{a,t}$, $C_{m,t}$ and $C_{s,t}$ from the agriculture, manufacturing, and services industries, with C_t given by:

$$C_t = \left(\sum_{i \in \{a,m,s\}} \lambda_i C_{i,t}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}, \text{ where } \varepsilon < 1. \quad (5)$$

Here ε is the elasticity of substitution between the goods and services produced by the three main sectors. It determines how quickly households change their consumption patterns in response to trends in relative prices between sectors due to structural transformation; Ngai and Pissarides (2007) show that $\varepsilon < 1$ generates trends in expenditure shares consistent with the data. The preference weights satisfy $\lambda_a + \lambda_m + \lambda_s = 1$.

We use log preferences here such that the real interest rate implied by the household's intertemporal choice, $r_t = \frac{1}{\beta}\frac{C_{t+1}}{C_t} - 1$, does not depend on population growth. Therefore, household's intertemporal choices are not affected by demographic factors. This allows us to isolate the effect of demographics on the transitional dynamics related to the (re-)allocation and training of workers that is the result of structural transformation.

Let $p_{i,t}$ for $i \in \{a, m, s\}$ be the price of goods and services of sector i, expressed in terms of units of the consumption aggregate C_t, which we use as our numeraire good throughout. The demand for each type of good $i \in \{a, m, s\}$ implied by the CES preferences is

$$C_{i,t} = \lambda_i^{\varepsilon} \left(\frac{1}{p_{i,t}}\right)^{\varepsilon} C_t. \tag{6}$$

The associated expenditure shares can be written as

$$s_{i,t} = \frac{p_{i,t} C_{i,t}}{C_t} = \lambda_i^{\varepsilon} p_{i,t}^{1-\varepsilon}. \tag{7}$$

The assumption that $\varepsilon < 1$ is sufficient to ensure that $s_{i,t}$ is increasing in $p_{i,t}$, which is consistent with cross-country evidence (Ngai and Pissarides, 2007).

Households do not incur any disutility from working and so supply their labor inelastically (as in Ngai and Pissarides, 2007; Herrendorf et al., 2014). We deviate from this existing literature in allowing workers to make career choices subject to training costs and retraining costs to switching between sectors. Since the introduction of these training frictions in the labor market is the main contribution of this paper, we present them in a separate subsection below. Before that, however, it is useful to first consider the firms' decisions that determine the supply side and labor demand schedules of our economy.

3.2 Firms

On the supply side of this economy, firms use labor as the only production factor. The production technologies in each of the sectors $i \in \{a, m, s\}$ are linear. We denote output of each respective sector by $Y_{i,t}$ and the sectoral Total Factor Productivity (TFP) by $A_{i,t}$, such that

$$Y_{i,t} = A_{i,t} L_{i,t}. \tag{8}$$

where $L_{i,t}$ is the amount of labor used in production in sector i. Note that, because workers differ in their productivity levels in this economy, $L_{i,t}$ is measured in terms of efficiency units of labor. What makes this a model of structural transformation is that we assume that the three sectors in the economy are subject to three different rates of TFP growth, g_i, such that

$$A_{i,t} = (1 + g_i) A_{i,t-1}. \tag{9}$$

Each of the three sectors is perfectly competitive in that firms are price and wage takers, and that there is free entry. Free entry of firms occurs until price equals the average (and marginal) cost of production:

$$p_{i,t} = \frac{w_{i,t}}{A_{i,t}}. \tag{10}$$

Here $w_{i,t}$ is the real wage paid per efficiency unit of labor in sector $i \in \{a, m, s\}$ in period t. An important difference between our model and other studies of structural transformation is that we consider sector-specific labor markets. This is the reason that $w_{i,t}$ is not equated across sectors and, hence, it is denoted by a subscript i.

3.3 Career decisions and the labor supply

The reason that wages differ between sector-specific labor markets is that individual workers' labor supply is not perfectly elastic across sectors. That is, individual workers do not simply choose to work for the sector that pays the highest wage. Instead, their sectoral choice is affected by three particular factors.

First, each worker receives an idiosyncratic sector-specific productivity shock $z_{i,t}$ in each period. These shocks are drawn from an exponential distribution with mean 1. For notational purposes, we combine these three shocks in the vector, $\mathbf{z}_t = [z_{a,t}, z_{m,t}, z_{s,t}]'$.

Second, it is costly for individual workers to get *trained* to acquire the skills necessary to work in a particular sector at the beginning of their career at age $h = 0$. Finally, it is also costly for them to get *retrained* in case they decide to work in a different sector mid-career at age $h > 0$. These latter two factors, i.e. the *training* and *retraining* costs, are the labor market frictions that are the main focus and contribution of this paper.

Both the *training* and *retraining* costs reflect that it takes workers time to initially get trained to start their career in a particular sector and then to get retrained in case they switch sectors. We capture these costs in terms of two parameters. The *training*-cost parameter $\phi \in [0, 1]$ is the fraction of the period when the worker is of age $h = 0$ that the worker spends on getting trained to work in a particular sector. The *retraining*-cost parameter $\gamma_h \in [0, 1]$ is age-specific and reflects the fraction of a period that a worker spends on being retrained when he or she decides to switch sectors of employment after the initial training at age $h = 0$.

Household members share their income and are fully insured against these costs. Because of this, each individual worker chooses his or her career path to maximize the expected present discounted value of lifetime earnings. At time t, this choice depends on the worker's age, h, industry of employment, i, and productivity shocks $\mathbf{z}_t$. The expected present discounted value of net future lifetime earnings by the individual equals $V_t(i, h; \mathbf{z}_t)$.

Given that the workers make optimal career decisions to maximize their $V_t(i, h; \mathbf{z}_t)$, we can write the expected present discounted value of lifetime earnings as the following Bellman equations. At age $h = 0$ the worker is not employed in a particular sector yet. She chooses an initial sector to start her career, taking into account the productivity shocks $\mathbf{z}_t$ and the fact that in order to get trained she will only work a fraction $(1 - \phi)$ of the first period of her career. The Bellman equation associated with this choice reads

$$V_t(0; \mathbf{z}_t) = \max_{i \in \{a,m,s\}} \left\{ (1 - \phi) z_{i,t} w_{i,t} + \frac{1 - \delta}{1 + r_t} \mathbb{E}_t V_{t+1}(i, 1; \mathbf{z}_{t+1}) \right\}. \tag{11}$$

Here r_t is the real interest rate in period t and $\mathbb{E}_t$ is the expectation conditional on information available at time t. This expectation is over all possible realizations of the worker-sector-specific productivity shocks, i.e. $\mathbf{z}_{t+1}$. The value function on the left-hand side does not have a sector index here because workers are not yet employed in a specific sector at the beginning of their career.

At age $h > 0$ the worker has started a career and the Bellman equation that determines the value for a worker of age h employed in industry i in period t and faced with productivity shocks $\mathbf{z}_t$ is given by

$$V_t(i, h; \mathbf{z}_t) = \begin{cases} \max\limits_{j \in \{a,m,s\}} \{(1 - \mathbb{I}(j \neq i)\gamma_h) z_{j,t} w_{j,t} + \frac{1-\delta}{1+r_t} \mathbb{E}_t V_{t+1}(j, h+1; \mathbf{z}_{t+1})\} & \text{if } h = 1, \ldots, H-1 \\ \max\limits_{j \in \{a,m,s\}} \{(1 - \mathbb{I}(j \neq i)\gamma_h) z_{j,t} w_{j,t}\} & \text{if } h = H \end{cases}. \tag{12}$$

Here, the indicator function $\mathbb{I}(j \neq i)$ reflects that a worker spends a fraction γ_h of her time on being retrained in the period when she decides to switch sectors and that she does not have to spend any time on retraining if she remains in the same sector. The latter case reflects that workers of age $h = H$ die with certainty and, thus, do not have a continuation value to their careers.

The result is that workers' career decisions involve a dynamic discrete choice problem. This discrete choice problem can be summarized in four variables that are important for the equilibrium dynamics of the labor supply and, thus, the economy.

The first two of these variables have to do with workers' *training* decisions at the beginning of their career, at age $h = 0$. First is the probability that a worker of age $h = 0$ is trained to work in sector i at time t, $\Phi_t(i)$. Second is the average number of efficiency units of labor that these workers supply to sector i at time t, $\tilde{z}_t(i; 0)$. The zero here denotes the age of the workers making the training decision.

The latter two variables are determined by the *retraining* decisions in period t, which depend on the worker's age h, the industry of employment i, and the productivity shocks, $\mathbf{z}_t$. The first is the probability that a worker of age $h > 0$ who works in sector i at time t decides to get retrained and starts working in sector j, $\Gamma_t(i, j; h)$. The other variable is the average productivity level for working in sector j of workers of age h who switch from sector i to j in period t, $\tilde{z}_t(i, j; h)$.

The assumption that the idiosyncratic sector-worker-time specific productivity shocks have exponential distributions allows us to solve $\Phi_t(i)$, $\tilde{z}_t(i; 0)$, $\Gamma_t(i, j; h)$, and $\tilde{z}_t(i, j; h)$ in closed form as a function of the wages in each of the sectors and the continuation values of working in each of the sectors. These closed-form solutions are algebraically intense and we therefore leave them for Section B in the Appendix. What is important in the rest of our analysis is that these four variables are sufficient to describe the dynamic evolution of the labor supply in our model.

3.4 Equilibrium

Product markets

Because output is only used for consumption, product market equilibrium requires

$$Y_{i,t} = C_{i,t}, \text{ for } i \in \{a, m, s\}. \tag{13}$$

Through the inverse demand function implied by (6), this determines the relative prices as a function of relative demands as

$$p_{i,t} = \lambda_i \left(\frac{C_t}{C_{i,t}} \right)^{\frac{1}{\varepsilon}} = \lambda_i \left(\frac{Y_t}{Y_{i,t}} \right)^{\frac{1}{\varepsilon}}. \tag{14}$$

where we define aggregate output as $Y_t = C_t$.

Labor markets

Free entry of producers drives down the price to equal the average cost of production, which, using (10) and (14), determines real wages as a function of relative output levels, $Y_{i,t}$, and TFP levels, $A_{i,t}$:

$$w_{i,t} = A_{i,t} p_{i,t} = A_{i,t} \lambda_i \left(\frac{Y_t}{Y_{i,t}} \right)^{\frac{1}{\varepsilon}}. \tag{15}$$

Using the production functions, (8), we can write this expression for the real wages in terms of labor inputs and relative productivity levels

$$w_{i,t} = \lambda_i A_{i,t}^{\frac{\varepsilon-1}{\varepsilon}} \frac{\left(\sum_j \lambda_j A_{j,t}^{\frac{\varepsilon-1}{\varepsilon}} L_{j,t}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{1}{\varepsilon-1}}}{L_{i,t}^{\frac{1}{\varepsilon}}}. \tag{16}$$

These are the inverse labor demand functions that determine how many efficiency units of labor firms will hire at given real wages and productivity levels.

If labor markets were frictionless and labor was homogenous then the labor supply for each sector would be perfectly elastic. As a result $w_{i,t} = w_t$, as in canonical models of structural transformation. Here the dynamics of the labor market are more complicated because the labor supply depends not only on current wages, but also on past and current career decisions of workers, which in turn depend on past and future wages.

In fact, because it is costly for workers to get trained to work in different sectors, the age-industry structure of the labor supply in this economy is a state variable whose law of motion is determined by demographics and the career decisions of workers. We derive the law of motion of the labor supply in three steps. In the first, we follow how many workers of age h work in sector i at time t. We denote this number by $E_t(i; h)$. In the second step, we consider how many efficiency units of labor these workers supply, based on their endogenous, career choices. In the final step, we aggregate these efficiency units over workers of all ages to get the aggregate labor supply for each sector i at time t.

The number of workers of age $h = 0$ who work in sector i at time t is equal to the number of persons of age $h = 0$ at time t, i.e. $N_t(0)$, times the fraction of them who decide to get trained to work in sector i, i.e. $\Phi_t(i)$. Thus,

$$E_t(i;0) = \Phi_t(i)\, N_t(0)\,. \tag{17}$$

The number of workers of age $h > 0$ who work in sector i at time t, $E_t(i;h)$ is equal to the number of workers of age $h-1$ who worked in the sector a period ago and didn't die, $(1-\delta)E_{t-1}(i;h-1)$, times the share of them who do not switch sectors, $\Gamma_t(i,i;h)$, plus the sum of the workers of age $h-1$ who worked in other sectors in period $t-1$ that survived and decided to switch to sector i in period t. Mathematically, the boils down to

$$E_t(i;h) = \sum_{j\in\{a,m,s\}} (1-\delta)\Gamma_t(j,i;h)\, E_{t-1}(j;h-1)\,, \text{ for } h = 1,\ldots,H. \tag{18}$$

The $3\times(H+1)$-dimensional tuple, $\{E_t(i;h)\}_{i,h}$ is the state variable in this economy that determines the labor supply.

This state variable is measured in terms of numbers of workers. The labor inputs for each sector, $L_{i,t}$, are measured in terms of efficiency units of labor instead. The labor supply can be transformed from number of workers into efficiency units of labor by multiplying the number of workers by their average productivity level that depends on their career choice and by the net (of training and retraining time) number of hours that these workers supply. To do this, we denote the number of efficiency units of labor supplied to sector i by workers of age h in period t by $L^s_t(i;h)$. This allows us to write

$$L^s_t(i;0) = (1-\phi)\,\tilde{z}_t(i;0)\,\Phi_t(i)\,N_t(0)\,, \tag{19}$$

and

$$L^s_t(i;h) = \sum_{j\in\{a,m,s\}} (1-\mathbb{I}(j\neq i)\,\gamma)\,\tilde{z}_t(1-\delta)\,(j,i;h)\,\Gamma_t(j,i;h)\,E_{t-1}(j;h-1)\,, \text{ for } h = 1,\ldots,H. \tag{20}$$

These equations define the industry-age-specific labor supply curves, in terms of efficiency units of labor.

Equilibrium in the labor market is when the sector-specific real wages, $w_{i,t}$, adjust such that the total number of efficiency units of labor demanded in a sector, i.e. $L_{i,t}$, equals the aggregate supply of efficiency units of labor to this sector. That is,

$$L_{i,t} = \sum_{h=0,\ldots,H} L_t^s(i;h). \tag{21}$$

Here, the left-hand side variable depends on the real wages through the inverse labor demand function, (16), while the right-hand side variables depend on the real wages through the workers' career choices.

4 The Impact of (Re-)Training Costs

In this section we consider the impact of (re-)training costs on the equilibrium dynamics of our model. We do so by comparing the dynamic equilibrium path of our economy with such costs with a baseline case in which such costs are not present. We call this baseline case the *Flexible Benchmark*. We illustrate the impact of (re-)training costs in four steps.

First, we describe the main properties of the equilibrium of the Flexible Benchmark case. This case is very similar to the model introduced in Bárány and Siegel (2018) and we discuss the similarities as well as emphasize the properties that are important to understand when we add (re-)training costs for workers. The most important property of the flexible benchmark is relative wages in the service sector are increasing compared to those in manufacturing and agriculture.

Next, we show that retraining costs accelerate the process of structural transformation in the economy rather than slow it down. We do so with an example in which retraining costs are the same for workers of all ages (*flat retraining costs*). The counterintuitive result that retraining costs accelerate structural transformation is because there is an option value to working in the service sector in anticipation of future wage gains. This option value is higher for young than for old workers. Because of this, under flat retraining costs the model has the counterfactual implication that older workers disproportionately switch back from the service sector to agriculture and manufacturing.

In the third step of this section we show that the absence of such career switchbacks in the data implies that, in the context of our model, retraining costs need to be increasing in age.

Finally, we explain why retraining costs in this model mainly affect the trends in relative wages across sectors rather than the trends in employment shares. This is a consequence of the near-Leontief preferences in the parameterization of our model.

The fact that sectoral employment shares are not affected much by retraining costs does not mean that these costs have not effect on output. We show that retraining costs reduce output in for two reasons. The first is that they siphon off labor from production into training. The second is that they distort the workers labor supply decisions reducing the efficiency units of labor employed in each sector.

4.1 Flexible Benchmark and solution method

Flexible Benchmark

Throughout the rest of this section we use the case in which there are no (re-)training costs, i.e. $\phi = \gamma_h = 0$ for $h = 1, \ldots, H$, as our main baseline for comparison. This *flexible benchmark* is a useful baseline because it is similar to the transitional dynamics studied in other analyses of structural transformation. Most notably, our flexible benchmark is very similar to the equilibrium dynamics in Bárány and Siegel (2018).

When workers do not face any training or retraining costs, their period-by-period labor supply decision in this model simply involves choosing to work in the sector i that pays the maximum compensation given their idiosyncratic productivity draws, $z_{i,t}w_{i,t}$. Thus, in this case, workers' career choices neither depend on their future career opportunities, nor on their initial industry of employment, nor on their age. Moreover, because we have abstracted from capital as a factor of production, the level of output per worker is also not affected by the population growth rate and workers' life expectancy. This means that we use the this baseline to consider how workers' career decisions change relative to it when workers face (re-)training costs and how the aggregate level of output per worker is affected by these changes in career decisions.

Because we rely on numerical methods for our analysis, we have to choose a set of baseline parameters to evaluate the dynamics of the flexible benchmark. Following Ngai and Pissarides (2008), we choose $\varepsilon = 0.1$, which is in the range of estimates implied by postwar U.S. national income data. We discipline our choice of the other parameters by having the flexible benchmark match the historical U.S. employment shares in agriculture, manufacturing, and services at the beginning and end of our sample, i.e. 1870 and 2010. We also match the average annualized historical growth rate of real GDP per capita over the sample. This calibration is described in more detail in Appendix C. We transform the annualized parameters in our model to reflect the length of a period which we set to 10 years.

The demographic parameters, i.e. the population growth rate, n, and the mortality rate, δ, do not matter for equilibrium in the flexible benchmark. Thus, we cannot use the flexible benchmark path to quantitatively discipline them. Instead, we choose n to match the average annual population growth rate in the U.S. between 1870 and 2010 and δ to match the average annual mortality rate for persons aged 10-70 born between 1904 and 1942 from Carter et al. (2006). In addition, we set the discount factor to $\beta = 0.95$ (annualized).

We assume that persons live for 6 periods in our model, i.e. $H = 5$. Given the period length of 10 years, one can interpret this as covering individuals from age 10 through 70 (similar to the data we analyzed).

The equilibrium path of our economy in the flexible benchmark closely resembles that of the model introduced in Bárány and Siegel (2018).[8] The main difference is that the labor supply in our model is made up of cohorts of workers who make lifetime career decisions. These career decisions, and how they compare to the evidence we presented in Section 2 is what we focus on here.

In the flexible benchmark, the choice of the sector that pays the highest compensation does not depend on a worker's age. As a result, all cohorts make the same career decisions. Panels (i) - (a) through (c) from Figure 3 show this. The are the model-equivalent of Figure 1. Contrary to the data, in the flexible benchmark the fraction of workers that are employed in each sector is the same across cohorts. This is why the lines in the panels in row (i) of Figure 3 overlap.

This also shows that the *changes* in employment shares across sectors are the same across cohorts in the benchmark. Table 3 reports the between-cohort share from the ANOVA of the changes of aggregate employment shares for three model specifications. These are the model-equivalent of Table 1, and the first row shows it for the flexible benchmark. These between-cohort shares are much lower than in the data. Even though all cohorts make the same career decisions, the between share is not zero. This is because average employment patterns over their life cycle differ across cohorts because they are alive during different periods.

Panel (a) of row (i) of Figure 4 shows the trends in the relative (log) wages across sectors that drive workers' career decisions. In the flexible benchmark wages in agriculture initially exceed those in manufacturing and services. Most importantly, relative wages in the service sector increase over time. This is driven by the increase in the relative price of services along the transitional path of this economy. As Ngai and Pissarides (2007) point out, the

[8] We illustrate this using a detailed set of results in Appendix D

complementarity of the goods and services produced by the three sectors in this economy and the relatively low productivity growth results in an increase in the relative price of services (and an increase in the share of value added) along the transitional path.

Two other effects reduce the trend in relative wages in services. The first is the low productivity growth in services, which puts downward pressure on real wage growth in the service sector. The second is the selection of workers into the service sector. Consistent with the evidence provided in Young (2014), workers of increasingly low average productivity are drawn into services. That is, $\tilde{z}_s, t$ declines over time. These workers are drawn into the service sector by the increase in relative wages driven by the different productivity growth rates and resulting relative price trends across sectors. This downward trend in labor quality in the service sector also puts downward pressure on the growth rate of average labor productivity (per worker).

On net, however, the productivity and worker-selection effects are smaller than the relative price effect. As a result, in this economy real wages in the sector with the lowest productivity growth grow the fastest.

Panel (b) of row (i) of Figure 4, for $h = 1, \ldots, H$, shows the covariance between the within-cohort and aggregate changes in employment shares, normalized by the variance of the aggregate changes. This ratio, from (1), is the regression coefficient of a regression of the age-specific changes in the employment shares on the aggregate changes in the employment shares. Because the career profiles of each cohort change in lockstep with the aggregate distribution of employment, these regression coefficients are 1 for all cohorts in the flexible benchmark. This stands in stark contrast to the variation in the data we documented in Figure 2

Thus, compared to the data, the flexible benchmark generates much less between-cohort variation in career profiles. Moreover, in the absence of (re-)training costs, changes in each cohort's career profile follows that of the overall economy, which is not the case in the data.

Implementation of Extended Path method

Because workers' career choices in the flexible benchmark do not depend on their previous decisions or age, the flexible benchmark case does not have a state variable and can be solved relatively straighforwardly on a period-by-period basis. In the presence of retraining costs, however, the equilibrium path of this economy depends on the initial age-industry distribution of the labor supply, $\{E_0(i; h)\}_{i,h}$.

This equilibrium path can be reduced to a path of (real) wages in each of the sectors, $\{w_{i,t}\}_{i,t}$ that, at each point in time, equates the demand and supply in each of the three labor markets. When the wages result in equilibrium in the labor market, Walras' Law implies that the product market will be in equilibrium as well. Because the equilibrium depends on the complicated evolution of the age-industry distribution of the labor supply, which in turn is determined by the workers' dynamic discrete career choices, it is not possible to find a closed-form solution for the equilibrium path. Instead, we have to resort to numerical methods.

The solution method that we use, described in Section E of the Appendix, is an application of the "Extended-Path "method, which was first discussed in Fair and Taylor (1983) and applied in, for example, Greenwood and Yorukoglu (1997) and Hobijn et al. (2006). Normally, the extended path method solves the transtional dynamics of a model between one steady state in period $t = 0$ and another at $t = T$. Because our model does not have a steady state or balanced growth path, we use a slightly different approach.

We solve the transitional dynamics of our model economy over the period from $t = -\tilde{t}_l$ through $t = T + \tilde{t}_r$. We assume that the initial state of the economy, at $t = -\tilde{t}_l$, is the one from our flexible benchmark. Moreover, we assume that the economy is on a balanced growth path, in which all sectors grow at the same rate, with no (re-)training costs after $t = T + \tilde{t}_r$. The reason that we add the left- and right-padding, i.e. $\tilde{t}_l > 0$ and $\tilde{t}_r > 0$, to the solution path is to reduce the impact of the assumed initial and final conditions on the part of the solution path, namely $t = 0, \ldots, T$, that we focus on.

4.2 Retraining costs accelerate structural transformation

The first thing we illustrate is that the addition of retraining costs to the model *accelerates* rather than slows down the process of structural transformation, as captured by the shift in employment from agriculture through manufacturing to services.

We illustrate this property of the model for a case with flat retraining costs. In particular, we look at the case where $\phi = 0.65$ and $\gamma_h = 0.5$ for all $h = 1, \ldots, H$. Under the restriction of flat retraining costs, i.e. $\gamma_h = \gamma$ for $h = 1, \ldots, H$, this combination of parameters gets the closest to the between share for the United States reported in Table 1 and the cohort-career regression coefficients shown in Table 2.

Figure 5 illustrates the difference between the path of the sectoral employment shares in the flexible benchmark (hashed bars) and in the case with flat retraining costs. It shows that the employment shares of manufacturing and services in the early stages of the structural transformation are higher under the retraining costs than in the flexible benchmark. At first glance, this might seem like a very counterintuitive result, because we tend to think of adjustment costs slowing down adjustments rather than accelerating them.

The reason for this acceleration is that, when workers face retraining costs, career choices do not only depend on current (real) wages, $w_{i,t}$, but also on the career continuation values, $\frac{1}{1+r_t}\mathbb{E}_t V_{t+1}(i, 1; \mathbf{z}_{t+1})$. These continuation values reflect the option of being employed in a particular sector. This option value is particularly high in our model for the service sector, because it largely captures the present discounted value of the future increases in relative wages in the service sector *over the rest of a worker's career*. As a result, workers facing retraining costs choose to be employed in the service sector more than those that do not face retraining costs. They do so in anticipation of future relative wage increases in services.[9]

In equilibrium, this increase in the labor supply in services results in higher employment in the service sector. It also subdues the trend in relative wages in the service sector compared to the flexible benchmark. This can be seen by comparing Panels (a) of rows (*i*) and (*ii*) of Figure 4. What happens is that retraining costs reduce the gross flows of workers between sectors that are largely driven by their idiosyncratic productivity levels. They, however, result inlarger *net* flows of workers across sectors that are coordinated by the common career continuation values that the workers face.

The younger the worker, the higher this option value, and the more the worker's decision is driven by it. The result is that older workers put more weight on current wages and productivity shocks when they make their labor supply decisions. For example, an old worker in services that draws a high productivity shock, i.e. gets a good opportunity in manufacturing or agriculture, will switch back to one of the shrinking sectors with declining relative wages in the economy. This can be seen from the three panels in the second row of Figure 3. The left and middle panels show the cohort-specific employment shares in agriculture and manufacturing. The career switchbacks of older workers are reflected by the increases in these shares in the last two periods of each cohort's career. These increases are offset by a decline in the share of workers in services. Higher wage gaps between manufacturing (as well as agriculture) and services imply higher career option values. Higher wage gaps also make switchbacks more common. This is why their size increases in Figure 3 over the transition path.[10]

[9]Note that this result does not depend on our assumption of per-period independent idiosyncratic shocks the career continuation values also matter in case of persistent shocks.

[10]The career switchbacks are the result of our assumption that workers are subject to sector-specific

4.3 Retraining costs increase with age

Comparing the three panels in the second row of Figure 3 with those in Figure 1, it is clear that the career switchbacks that the model generates are counterfactual. In order for them not to occur in our model, we need to assume that retraining costs are higher for older workers. These costs need to increase so that they offset the decline in the career option value of workers in shrinking sectors.

To illustrate the effect of increasing retraining costs over the life cycle, we increase the retraining costs for workers of age $h = 4$ and $h = 5$ to $\gamma_4 = 0.9$ and $\gamma_5 = 0.999$. This makes it very costly for older workers to change their sectors of employment, and greatly reduces career switchbacks. We call this the *increasing retraining costs* case and its equilibrium path is plotted in Row (*iii*) of Figures 3 and 4. This is an extreme example. However, it shows that with late-career high retraining costs, career option values play a larger role. Workers move more rapidly from manufacturing to services in anticipation of future relative wage increases. This is, at least qualitatively, more consistent with the cohort career employment patterns depicted in Figure 1. Our data of consecutive cross-sections does not allow us to consider individual employment transitions by workers, but evidence in Menzio et al. (2016) suggests that these transitions are, indeed, declining over the life cycle.

Note, however, that with more workers lining up in services the equilibrium trend of relative wages is lower, compared to the flexible benchmark and flat retraining costs cases. The left panel of Row (*i*) of Figure 4 shows this. It is most evident from the real wages in agriculture: in the flexible benchmark these were declining after 40 years, and with increasing retraining costs, these continue to increase. In fact, the difference in real wages in agriculture at the end of the quarter millennium that we consider is 200 log points. That is about a 700 percent difference. In fact, most of the impact of retraining costs is not on the employment shares but on the relative wages across sectors.

4.4 Bulk of impact retraining costs through relative wages

The reason that most of the adjustment to retraining costs in this economy goes through relative wages is that preferences are almost Leontief. As is commonly done in studies of structural transformation (e.g. Ngai and Pissarides, 2008), we have chosen a very low elasticity of substitution between the goods and services produced by the sectors in the

productivity shocks each period. In our model these shocks are independent over time. If they weren't, highly correlated shocks over time or across sectors would imply smaller career switchbacks.

economy, i.e. $\varepsilon = 0.1$. This elasticity is consistent with the trends in relative prices and value added shares of the three sectors in U.S. time series.

It implies, however, that the marginal rates of substitution between $C_{i,t}$ for $i \in \{a, m, s\}$ vary a lot in response to small changes in the quantities. These marginal rates of substitution affect the relative wages in this economy because they determine the elasticities of the sector-specific labor demand curves. The smaller ε the less elastic the labor demand curves. The inelastic labor demand curves mean that changes in the labor supply mainly result in changes in relative wages rather than changes in sectoral employment shares.[11]

Figure 6 illustrates this for period $t = 50$. It shows the sector-specific labor supply and labor demand curves under the *flexible benchmark* (dashed) and frictional *increasing retrainings costs* (solid). The downward sloping labor demand curves are very inelastic (close to vertical) and do not move much due to the changes in the wages in the other sectors from the flexible to the frictional case. The result is that the equilibrium employment shares do not change a lot between the flexible and frictional case. The shifts in the labor supply curves induced by the retraining costs therefore translate mostly into changes in the equilibrium real wages.

4.5 Output losses due to retraining and allocation of labor

The fact that retraining costs do not affect sectoral employment shares much does not mean, however, that these costs do not affect the level of output in the economy. Output in the case of increasing retraining costs is much lower under increasing retraining costs than in the flexible benchmark for two reasons. The first is the loss in effective labor input due to the (re-)training time of workers. The dashed purple line in Figure 7 shows the size of the gap in output between the increasing costs and flexible benchmark cases due to (re-)training in log points over the transition path.

At the beginning of the transition path, the (re-)training costs result in a loss of more than a third of output. In large part, this is due to the training costs, ϕ, that all workers incur when they start their career. In addition, this reflects the retraining time of workers that switch sectors.[12] Note that this loss declines as more workers are employed in services

[11] In the limiting case of Leontief preferences in which $\varepsilon \downarrow 0$ labor demand (in efficiency units) is pinned down completely by technology parameters and all the effect of retraining costs on the labor supply is reflected in changes in relative wages rather than employment shares.

[12] The level of these costs in this model is relatively high because the variance of the Exponential distribution of workers' idiosyncratic productivity shocks is one. That is, the size of this loss is partly driven by our distributional assumption about workers' productivity shocks that keeps the discrete choice problem

and fewer gross flows of workers occur between sectors. At the end of the transition path, training time absorbs about a quarter of the time available from all workers, which is mainly due to the training cost of workers in the first period of their lives, i.e. because $\phi = 0.65$.

The second reason for the output loss has to do with the allocation of workers across sectors. (Re-)training costs can have a worker choosing to go into services when her current productivity (and wage) is higher in manufacturing or agriculture. This is due to the option value of going into services. The solid line, labeled "Total" in Figure 7 adds this effect on output on top of that for the training time to show the total output loss in the increasing retraining costs case compared to the flexible benchmark.

The part due to the allocation of labor, i.e. the difference between the two lines, peaks after about 30 years at around 15 percent and then declines as the employment distribution shifts towards services. Thus, because (re-)training costs affect the allocation of workers across sectors they reduce average labor productivity in this economy. This effect eventually goes to zero when all workers are employed in services.

In sum, in the long run retraining costs result in about a 20 percent reduction in output, solely due to the training of young workers to prepare them for careers in the service sector. However, during the process of structural transformation, the retraining time involved with the reallocation of workers across as well as workers' forward-looking career decisions double this impact. That is, the total output loss peaks after 30 years at around 40 percent of the level of output in the flexible benchmark.

4.6 Unanticipated acceleration of structural transformation

The insight that most of the impact of retraining costs in a canonical model of structural transformation is on wages instead of employment shares is important also for understanding the effect of an unanticipated acceleration in structural transformation. We illustrate this using a scenario in which the economy is on the same path as in the previous subsections for *flexible benchmark* and *the increasing retraining costs* cases. At time $t = 0$ structural transformation unexpectedly accelerates in that both g_a and g_m permanently increase by 50 percent compared to their calibrated values.[13]

Figure 8 plots the changes in the paths of the employment shares and the log of the average real wage paid per employee by sector. The two panels in Row (i) of the figure show the

of workers tractable and solvable.

[13]That is, g_a increases from 0.045 (annualized) to 0.0675 and g_m from 0.020 to 0.030.

results for the *flexible benchmark* and the panels in Row (ii) for the *increasing retraining costs*.

The left panels in each of these rows show the change in the employment shares by sector over the post-shock equilibrium path of the economy. Both panels show how an increase in g_a and g_m results in a faster reallocation of labor from the agriculture and manufacturing sectors to services. For the reasons we already explained in subsection 4.4 above, the impact of the acceleration in technological change on the employment transition patterns is very similar for both cases.

The impact of the shock differs mostly in terms of the paths of the sectoral wages in the two cases. This can be seen by comparing the right-hand side panels in Rows (i) and (ii) of Figure 8. To understand the impact of the shock on wages under retraining costs it is important to distinguish between the initial impact and the subsequent dynamics.

In the absence of the retraining costs there is no initial impact of the shock on wages because it does not affect the sectoral productivity levels, $A_{i,0}$ for $i \in \{a, m, s\}$. The presence of retraining costs makes more younger workers line up in services in anticipation of future wage gains in that sector in response to the shock. This increases the labor supply in services at the time of the shock and lowers real wages in the service sector.[14] For young workers, this decline in the wage in the service sector at $t = 0$ is offset by subsequent increases in expected earnings later on in their careers. For older workers employed in the service sector at the time of the shock, this decline in the real wage is not offset by future wage increases and, in hindsight, some of these workers would have preferred to be employed in agriculture or manufacturing but are stuck in the service sector due to their high retraining costs when the shock occurs.

In the subsequent periods the shock results in a higher increase in the growth rate of real wages in the service sector than in the other two sectors. In this sense, it results in a stagnation of relative wage growth in agriculture and manufacturing. In the first 5 decades after the shock, this stagnation is more profound for the case of increasing retraining costs when the allocation of workers across sectors is still affected by the initial allocation at the time of the shock. In the longer-run retraining costs dampen the effect of the shock on relative real wage growth across sectors for the same reason that they dampen the trends in relative wages in the left-hand side panels in Figure 4. This can be seen from the very different scales on the vertical axes of the two right-hand side panels in Figure 8.

[14] Most of these workers turn out to be drawn out of the agriculture sector, which is where wages increase in response to the shock.

5 Conclusion

Using data on sectoral employment patterns by birth cohort from more than 50 countries around the world, we show that the bulk of the shift in the allocation of employment across major sectors of economic activity that is typical for structural transformation is because younger cohorts are disproportionately employed in growing sectors. We argue that the importance of between-cohort differences in sectoral employment shares is indicative of workers incurring substantial retraining costs when they switch sectors.

To illustrate the aggregate implications of such retraining costs, we introduce a model of structural transformation with overlapping generations that face (re-)training costs when they make their labor supply decisions. On the aggregate level, we follow Ngai and Pissarides (2007) and we model structural transformation as being driven by different levels of TFP growth across sectors. These sectors produce different types of output that are gross complements in aggregate CES preferences and drive demand patterns.

Our main theoretical contribution is modeling the career decisions of different cohorts as a discrete choice problem. In each period workers decide what sector they would like to work in based on their current productivity levels, the future wage prospects in each of the sectors, and the potential retraining costs they incur.

To match the between-cohort contribution to structural transformation in our model, we need to include substantial training costs for workers at the beginning of their careers and retraining costs when they switch sectors. We obtain four important insights from our model once we introduce (re-)training costs.

First of all, (re-)training costs *accelerate* the reallocation of labor to the growing service sector rather than slow it down. This is because workers decide to take jobs in the service sector in anticipation of future relative wage increases.

Secondly, because there is no option value to being in services for workers at the end of their careers, some of these move back to agriculture and manufacturing. These counterfactual career switchbacks in the model suggest that retraining costs are increasing in age and very high for old workers.

Thirdly, the strong gross complementarity of the goods and services produced by the three sectors in our economy results in the bulk of the impact of the retraining costs being on trends in relative wages across sectors rather than trends in employment shares.

Finally, when structural transformation accelerates unexpectedly more young workers will choose to supply their labor in the service sector in anticipation of future relative wage growth in services. This reduces the service sector wage on impact of the shock which most negatively affects the career earnings outlook of older workers in the service sector. In the longer-run the shock reduces relative wages in agriculture and manufacturing.

Our main contribution here is the analysis of the evolution of the labor supply across cohorts in a general equilibrium framework. To isolate the effects of the assumptions we made about the labor supply, we deliberately embed our cohort-specific labor supply model in a simple model of structural transformation. This, deliberately stylized, framework reveals potentially large effects of labor market frictions, generally ignored in models of economic growth, on long-run economic outcomes.

The reallocation of labor due to structural transformation is only one of the potential applications of our theoretical contribution. For example, it can also be used to consider the cohort-specific effects of the labor market impacts of trade (Autor et al., 2013, e.g.). As is common in models of economic growth, we have abstracted from workers' participation and retirement decisions. Adding those margins is a useful extension and the subject of future research.

References

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

Bárány, Zsófia L. and Christian Siegel, "Job Polarization and Structural Change," *American Economic Journal: Macroeconomics*, 2018.

Baumol, William J. and William G. Bowen, *Performing Arts - The Economic Dilemma*, MIT Press: New York, NY, 1968.

Bolt, Jutta and Jan Luiten van Zanden, "The Maddison Project: Collaborative Research on Historical National Accounts," *Economic History Review*, 2014, *67* (3), 627–651.

Buera, Francisco J, Joseph P Kaboski, and Kent Zhao, "Quantifying the Growth in Services: the Role of Skills, Scale, and Female Labor Supply," *Journal of Human Capital*, forthcoming.

Carrillo-Tudela, Carlos, Bart Hobijn, Powen She, and Ludo Visschers, "The Extent and Cyclicality of Career Changes: Evidence for the U.K.," *European Economic Review*, 2016, *84*, 18 – 41.

Carter, Susan B., Scott S. Gartner, Michael R. Haines, Alan L. Olmstead, Richard Sutch, and Gavin Wright, *The Historical Statistics of the United States Online Edition: Millennial Edition*, New York: Cambridge University Press, 2006.

Caselli, Francesco and Wilbur John Coleman, "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation," *Journal of Political Economy*, 2001, *109* (3), 584–616.

Cociuba, Simona and James C. MacGee, "Demographics and Sectoral Reallocations: A Search Theory with Immobile Workers," February 2018. mimeo, University of Western Ontario.

Duarte, Margarida and Diego Restuccia, "Relative Prices and Sectoral Productivity," 2017. NBER Working Paper 23979.

Duernecker, Georg and Berthold Herrendorf, "Structural Transformation of Occupation Employment," 2017. mimeo, Arizona State University.

Echevarria, Cristina, "Changes in Sectoral Composition Associated with Economic Growth," *International Economic Review*, 1997, *38* (2), 431–452.

Fair, Ray C and John B Taylor, "Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models," *Econometrica*, 1983, *51* (4), 1169–1185.

Gaffield, Chad, Peter Baskerville, Sean Cadigan, Marc St-Hilaire, Claude Bellavance, France Normand, Gordon Darroch, Carl Amhrein, Lorne Tepperman, Charles Jones, and Eric Sager, "National Sample of the 1911 Census of Canada [dataset]," Edmonton, Canada: University of Alberta [distributor] 2009.

Greenwood, Jeremy and Mehmet Yorukoglu, "1974," *Carnegie-Rochester Conference Series on Public Policy*, 1997, *46* (1), 49–95.

Hendricks, Lutz, "Cross-country Variation in Educational Attainment: Structural Change or Within Industry Skill Upgrading?," *Journal of Economic Growth*, 2010, *15* (3), 205–233.

Herrendorf, Berthold and Todd Schoellman, "Wages, Human Capital, and Barriers to Structural Transformation," *American Economic Journal: Macroeconomics*, 2018, *10* (2), 1–23.

_ **, Richard Rogerson, and Akos Valentinyi**, "Growth and Structural Transformation," in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 2, Elsevier, 2014.

Hobijn, Bart, Federico Ravenna, and Andrea Tambalotti, "Menu Costs at Work: Restaurant Prices and the Introduction of the Euro," *Quarterly Journal of Economics*, 2006, *121* (3), 1103–1131.

Inwood, Kri and Jack Chelsea, "National Sample of the 1891 Census of Canada," Guelph, Canada: University of Guelph 2011.

Kambourov, Gueorgui and Iourii Manovskii, "Rising Occupational and Industry Mobility in the United States: 1968-97," *International Economic Review*, 2008, *49* (1), 41–79.

Kendrick, John W., *Productivity Trends in the United States*, Princeton: Princeton University Press, 1961.

Kim, Dae-Il and Robert H. Topel, "Labor Markets and Economic Growth: Lessons from Korea's Industrialization, 1970-1990," in "Differences and Changes in Wage Structures," University of Chicago Press, 1995, pp. 227–264.

Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie, "Beyond Balanced Growth," *Review of Economic Studies*, 2001, *68* (4), 869–882.

Lagakos, David and Mike Waugh, "Specialization, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, 2013, *103* (2), 948–980.

Menzio, Guido, Irina Telyukova, and Ludo Visschers, "Directed Search over the Life Cycle," *Review of Economic Dynamics*, January 2016, *19*, 38–62.

Minnesota Population Center, "Integrated Public Use Microdata Series — International: Version 6.3 [Machine-readable database].," Minneapolis: University of Minnesota. 2014. http://www.ipums.org.

_ , "North Atlantic Population Project: Complete Count Microdata.," Minneapolis: Minnesota Population Center 2017. Available online at http://doi.org/10.18128/D040.V2.3.

Ngai, L. Rachel and Barbara Petrongolo, "Gender Gaps and the Rise of the Service Economy," *American Economic Journal: Macroeconomics*, 2017, *9* (4), 1–44.

_ **and Christopher A. Pissarides**, "Structural Change in a Multisector Model of Growth," *American Economic Review*, 2007, *97* (1), 429–443.

_ **and** _ , "Trends in Hours and Economic Growth," *Review of Economic Dynamics*, April 2008, *11* (2), 239–256.

Perez, Santiago, "The Impact of Railroads on Education and Occupational Attainment: Evidence from 19th Century Argentina," 2016. mimeo, Stanford University.

Porzio, Tommaso and Gabriella Santangelo, "Structural Change and the Supply of Agricultural Workers," 2017. mimeo, University of California - San Diego.

Rendall, Michelle, "Brain Versus Brawn: the Realization of Women's Comparative Advantage," 2017. mimeo, Monash University.

Ruggles, Steven, Trent J. Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schoreder, and Matthew Sobek, "Integrated Public Use Microdata

Series: Version 5.0 [Machine-readable database]," Minneapolis, MN: Minnesota Population Center [producer and distributor] 2010. Available online at http://usa.ipums.org/usa/.

Schultz, Theodore William, *The Economic Organization of Agriculture*, New York: McGraw-Hill New York, 1953.

The Digital Archive (The National Archive), Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center, "National Sample of the 1865 Census of Norway, Version 2.0," Tromsø, Norway: University of Tromsø 2008.

_ , "National Sample of the 1900 Census of Norway, Version 2.0," Tromsø, Norway: University of Tromsø 2008.

_ , "National Sample of the 1910 Census of Norway, Version 1.0," Tromsø, Norway: University of Tromsø 2011.

The Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center, "National Sample of the 1875 Census of Norway, Version 2.0," Tromsø, Norway: University of Tromsø 2008.

Young, Alwyn, "Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services," *American Economic Review*, November 2014, *104* (11), 3635–67.

Tables and figures

Table 1: Between Cohort Share of Structural Transformation

Sample	Total	By Industry		
		Agriculture	Manufacturing	Services
United States	53%	71%	36%	41%
International	53%	66%	23%	48%

Table 2: Between Demographic Group Share of Structural Transformation

	Total	By Industry		
		Agriculture	Manufacturing	Services
United States	16%	10%	19%	18%

Table 3: ANOVA in Different Model Specifications

	Total	By Industry		
		Agriculture	Manufacturing	Services
		(i) Flexible benchmark $\phi = 0$, $\gamma_h = \{0, 0, 0, 0, 0\}$		
Share Between Cohorts	18%	18%	18%	18%
		(ii) Flat retraining costs $\phi = 0.65$, $\gamma_h = \{.5, .5, .5, .5, .5\}$		
Share Between Cohorts	53%	24%	66%	-5%
		(iii) Increasing retraining costs $\phi = 0.65$, $\gamma_h = \{.5, .5, .5, .9, .999\}$		
Share Between Cohorts	20%	45%	23%	52%

Figure 1: Structural Transformation in the United States

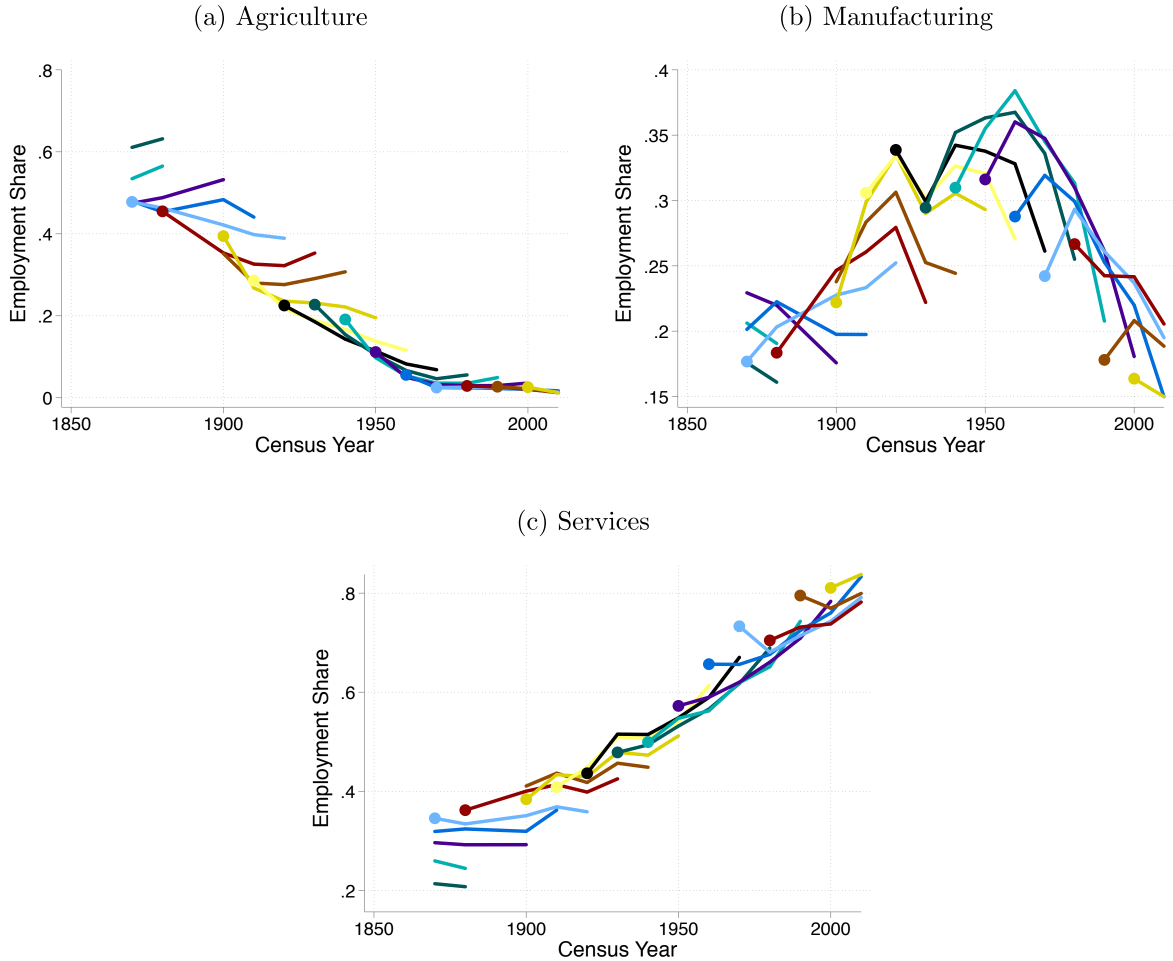

Figure 2: Within cohort effects by age

(a) United States (b) International

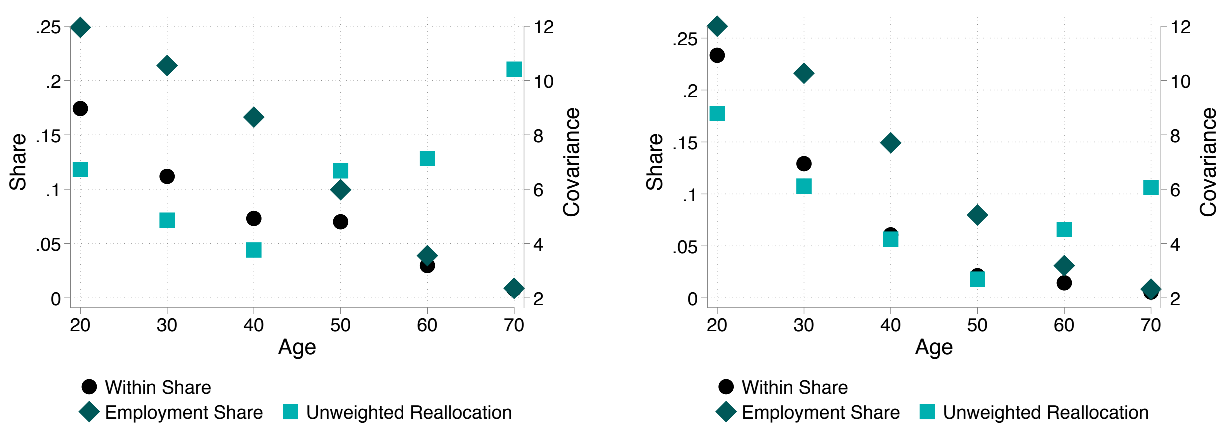

Figure 3: Cohort-specific employment shares in three model specifications

(i) Flexible benchmark, $\phi = 0$, $\gamma_h = 0$ for $h = 1, \ldots, 5$

(a) Agriculture (b) Manufacturing (c) Services

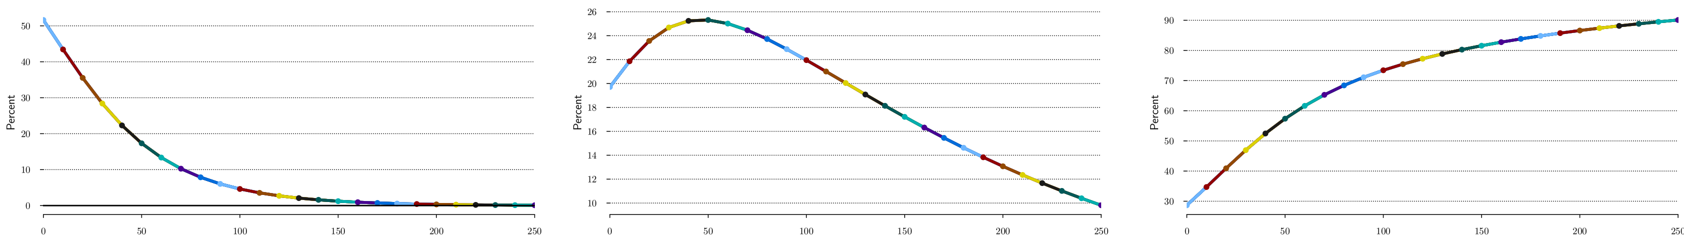

(ii) Flat retraining costs, $\phi = 0.65$, $\gamma_h = 0.5$ for $h = 1, \ldots, 5$

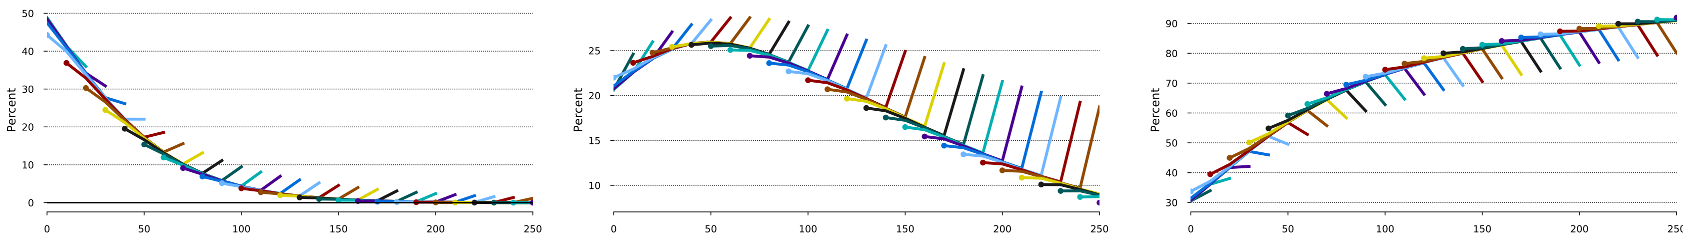

(iii) Increasing retraining costs, $\phi = 0.65$, $\gamma_h = \{.5, .5, .5, .9, .999\}$

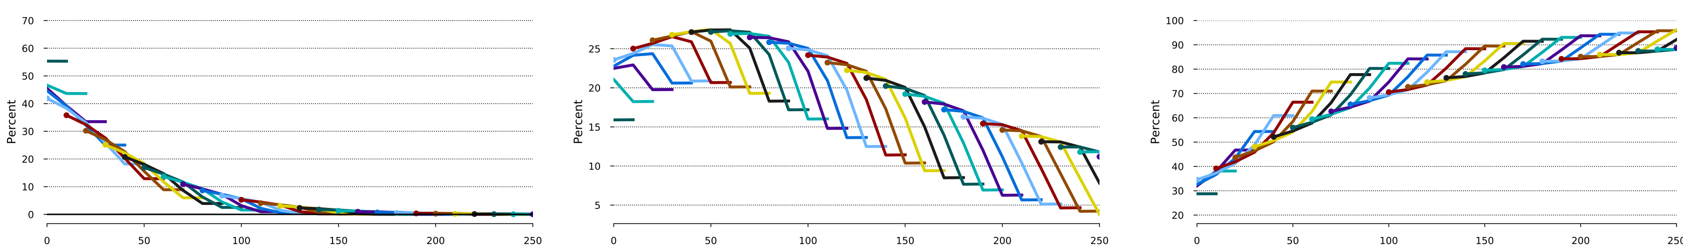

Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010.

Figure 4: Wages and Cohort career profiles in three model specifications

(i) Flexible benchmark, $\phi = 0$, $\gamma_h = 0$ for $h = 1, \ldots, 5$

(a) Real compensation per worker (b) Age-specific within regression coefficient

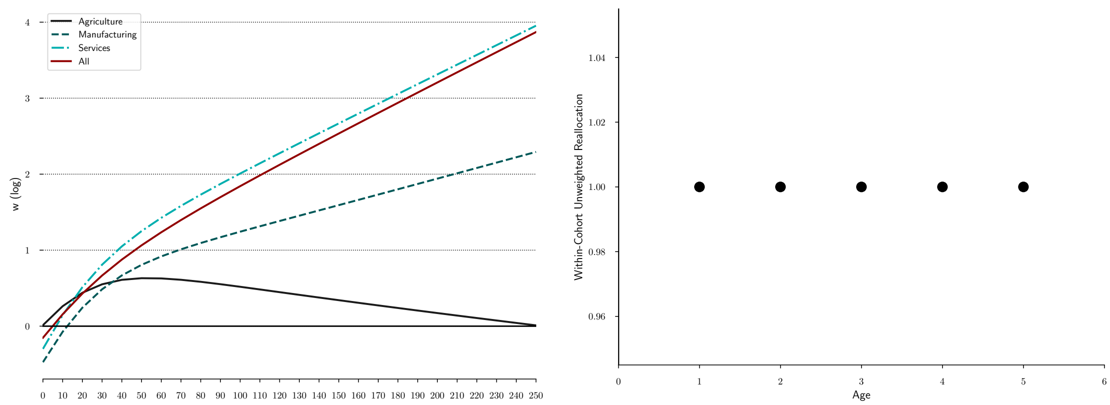

(ii) Flat retraining costs, $\phi = 0.65$, $\gamma_h = 0.5$ for $h = 1, \ldots, 5$

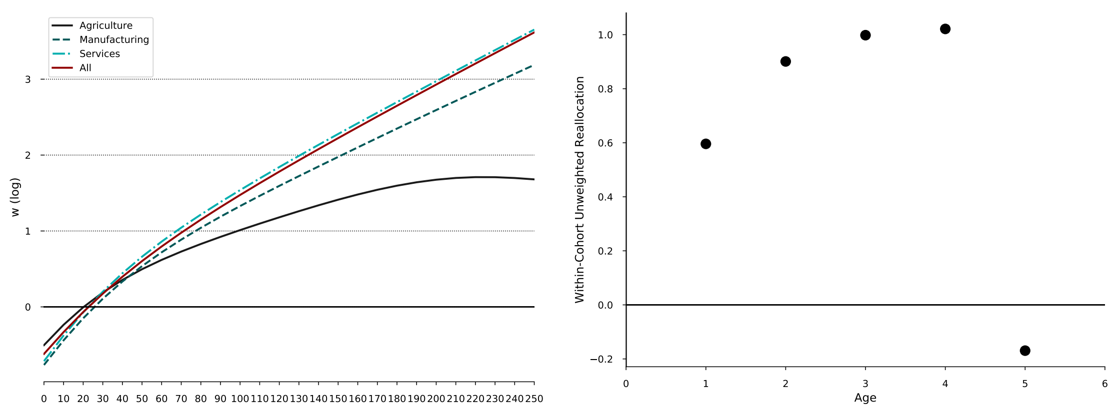

(iii) Increasing retraining costs, $\phi = 0.65$, $\gamma_h = \{.5, .5, .5, .9, .999\}$

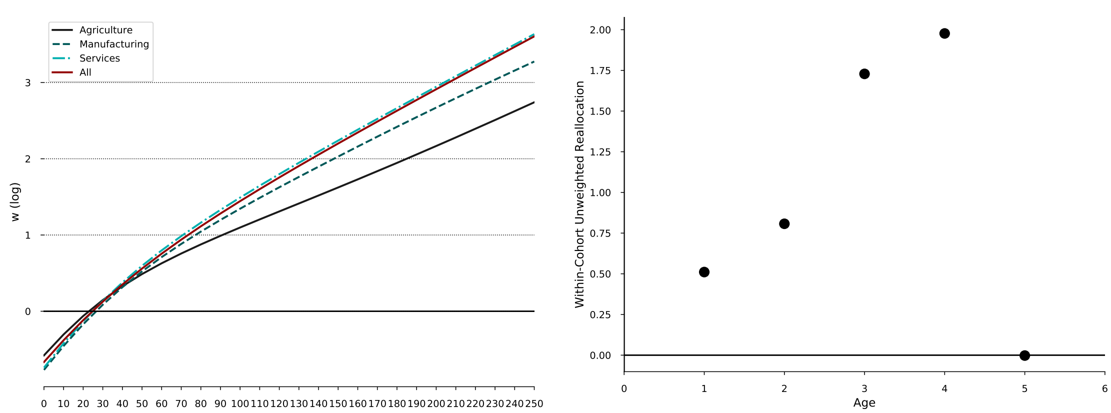

Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010. Age is measured in decades.

Figure 5: Sectoral employment shares in flexible benchmark and under flat retraining costs

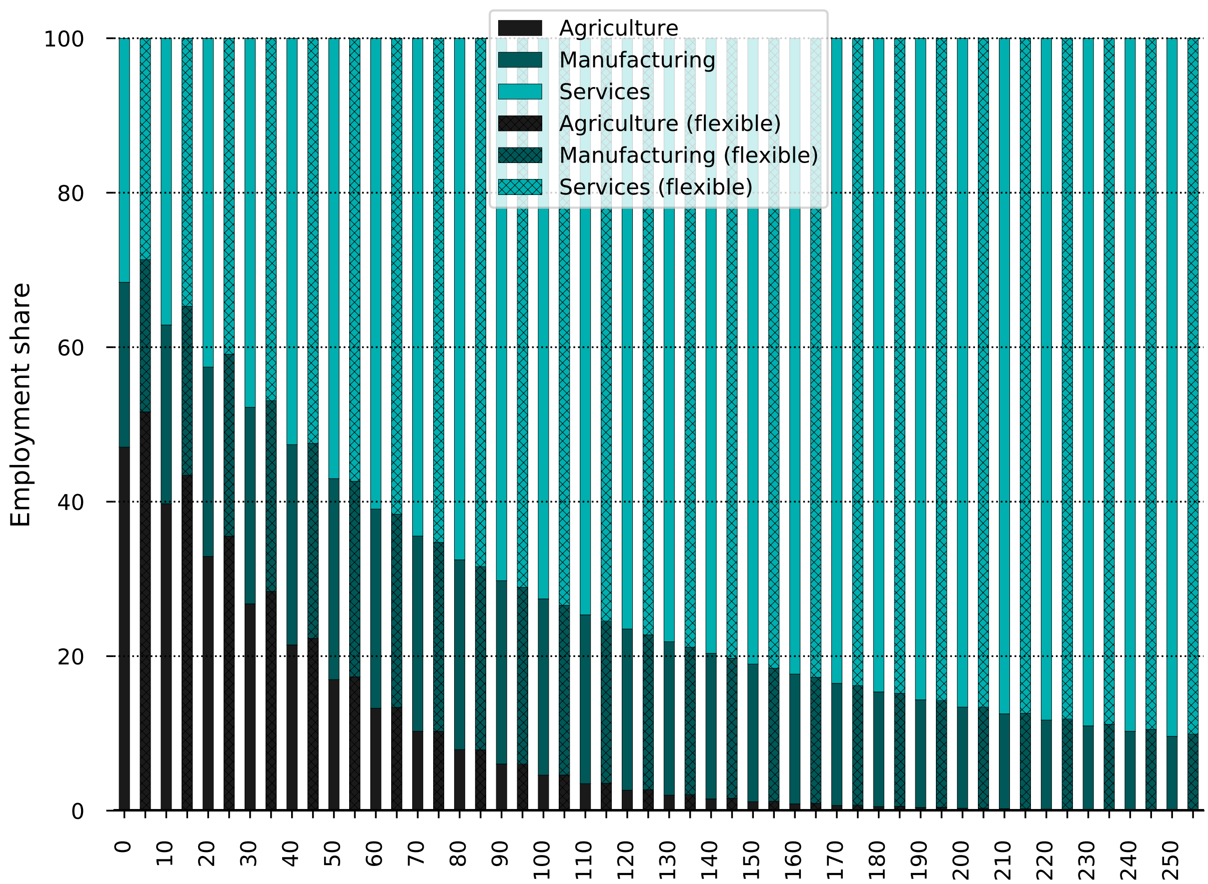

Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010.

Figure 6: Sectoral labor supply and labor demand curves at $t = 50$

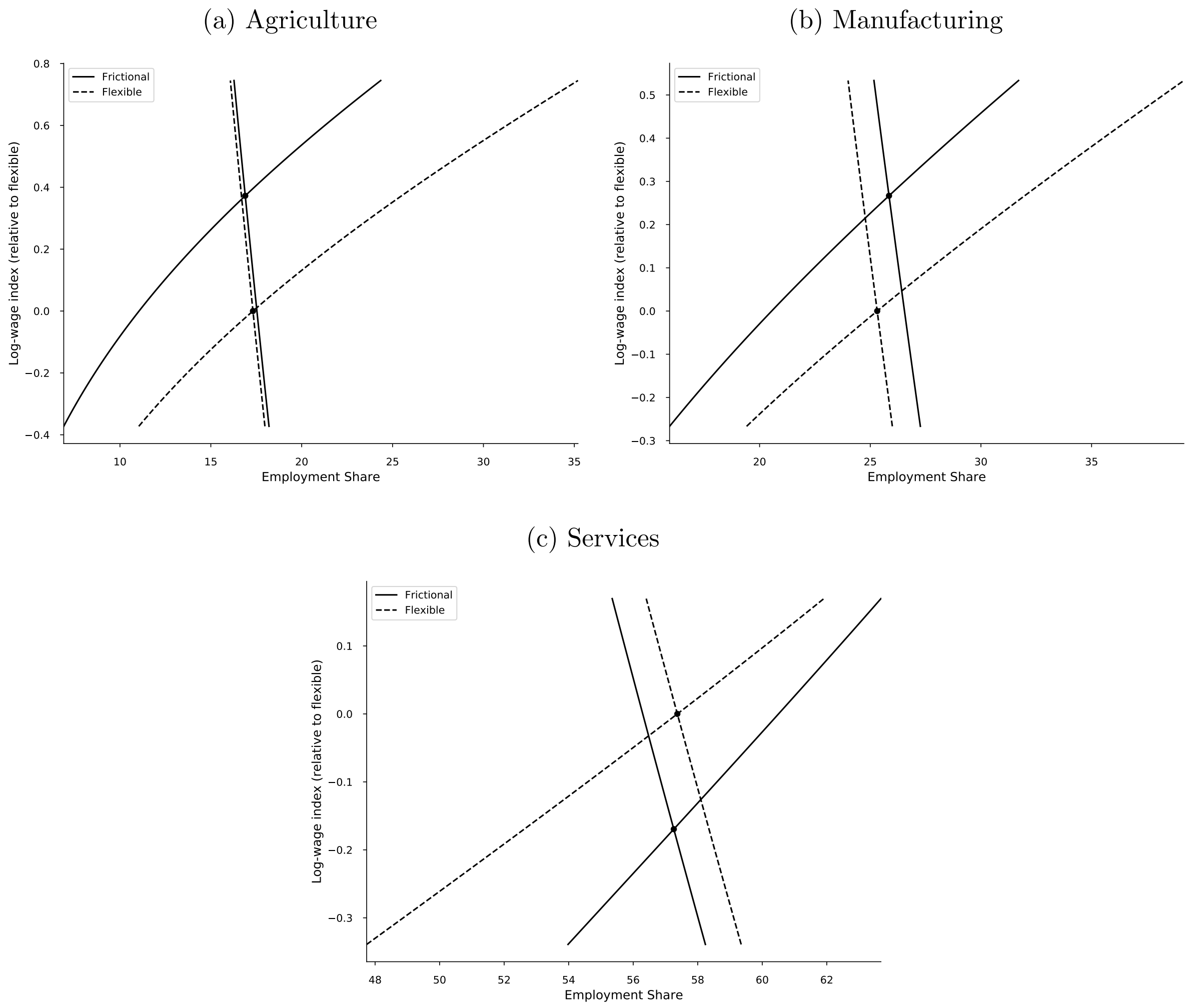

Note: Dashed lines are the labor demand (downward sloping) and labor supply (upward sloping) curves in *flexible benchmark*. Solid line are those in *increasing retraining costs* case. All curves are conditional on the equilibrium real wages in the other two sectors. The vertical axes show the log real wage in deviation from that in the flexible benchmark and the horizontal axes the employment shares $E_{i,t}/E_t$. The dots depict the equilibrium combinations of the log real wages and employment shares.

Figure 7: Difference in output levels between increasing retraining costs and flexible benchmark

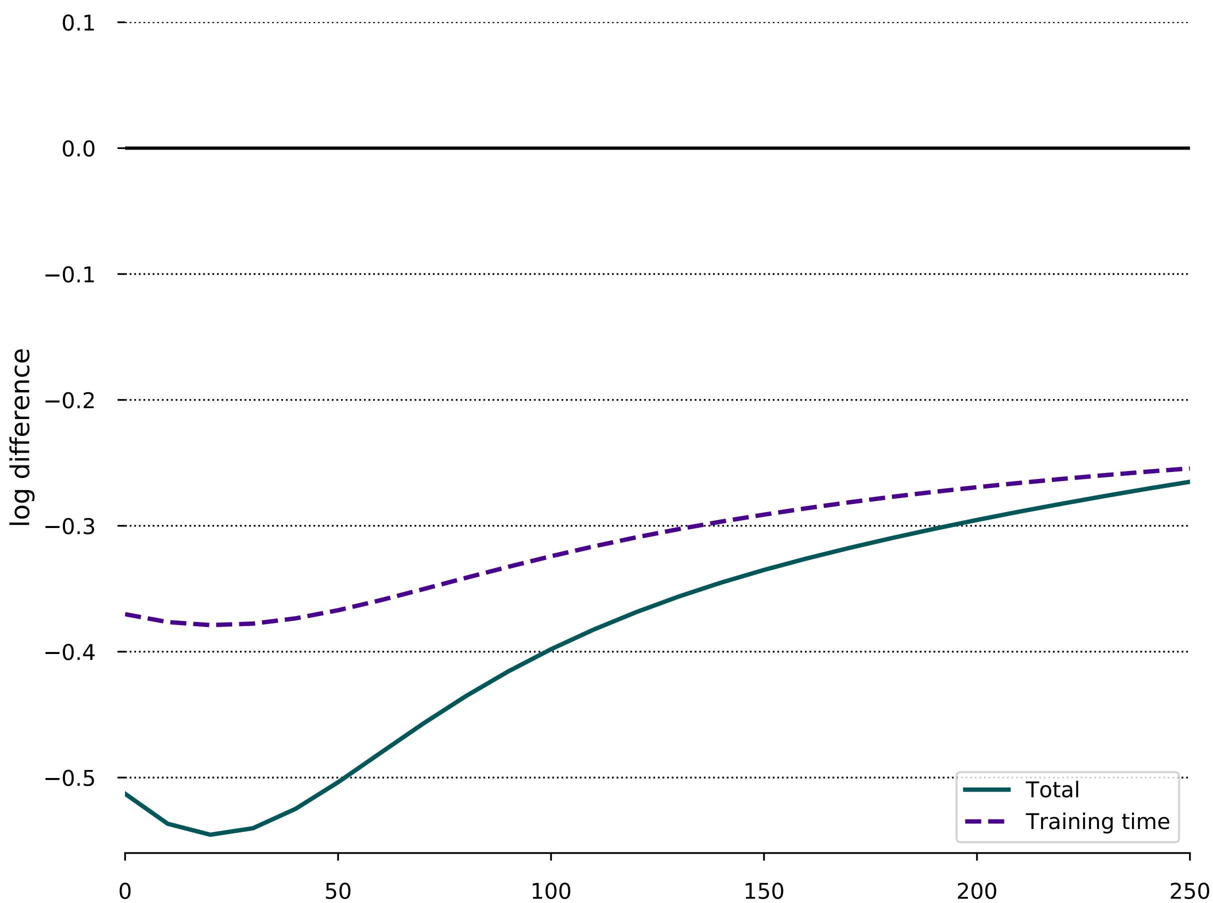

Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010. Gap is measured in terms of log points.

Figure 8: Impact of acceleration of structural transformation on employment and wages

(i) Flexible benchmark, $\phi = 0$, $\gamma_h = 0$ for $h = 1, \ldots, 5$

(a) Employment shares

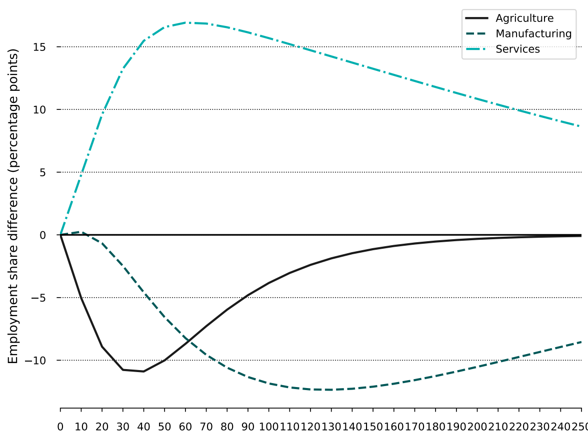

(b) Relative wages

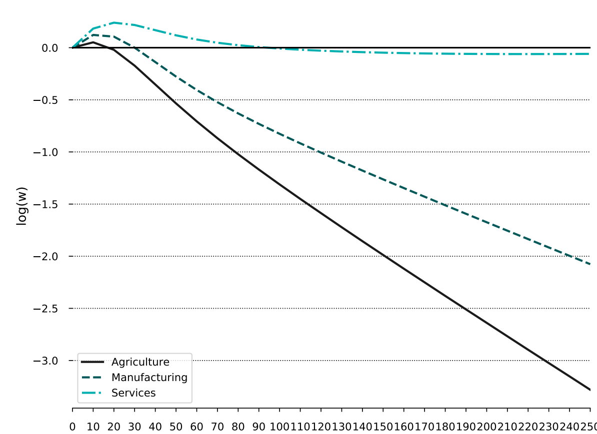

(ii) Increasing retraining costs, $\phi = 0.65$, $\gamma_h = \{.5, .5, .5, .9, .999\}$

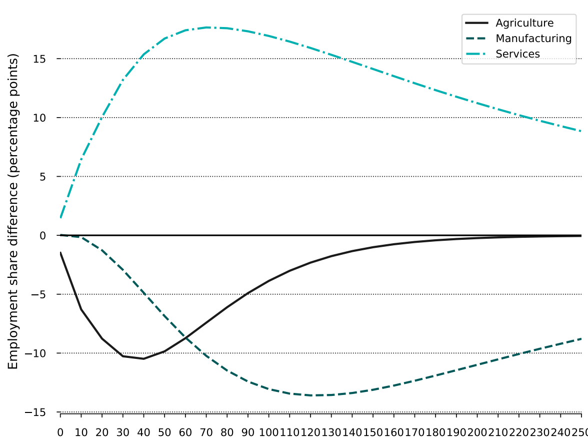

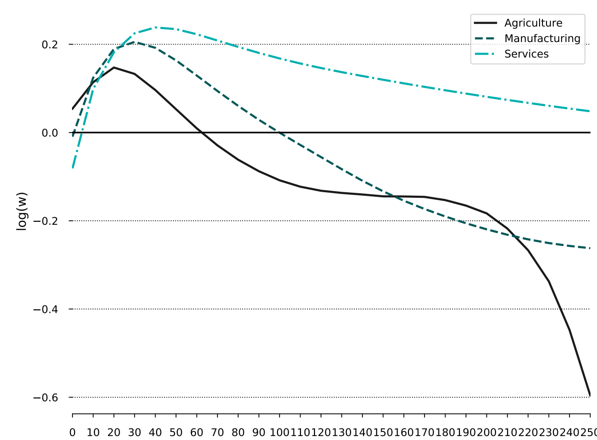

Note: The panels on the left show the percentage point difference between the sectoral employment shares with an acceleration in the rates of technological change in agriculture and manufacturing at time $t = 0$ and those where the rates of technological change are constant. The panels on the right show the differences in the log of the sectoral average wages paid per employee between the path with the shock and without it.

A Data Details and Additional Results

A.1 United States Data

For the United States we draw on census microdata, taken from Ruggles et al. (2010). We exclude the 1850 and 1860 censuses because they did not enumerate slaves. The 1890 census microdata were lost in a fire. We pool the 2008–2012 American Community Surveys to take the place of a 2010 census, in line with usual IPUMS practice. The pooled surveys have very similar questions, responses, and coverage as the 2000 census. Together, we have 14 surveys spanning 140 years at ten year intervals, excluding 1890.

We focus throughout on workers who are employed and have a valid response to age and industry of employment. We aggregate industry codes to three broad industry groups. Agriculture includes farming, forestry, and fishing. Manufacturing includes most remaining goods production, including manufacturing, mining, and construction. Services includes the remaining industries: utilities; retail and wholesale trade; hotels and restaurants; finance, insurance, and real estate; public administration and defense; education; health and social services; private household services; and other/miscellaneous services. We discard reported industry of employment for those younger than 20, under the view that this likely represents part-time or seasonal work and not a serious career choice.

Our demographic analysis focuses on the role of education, gender, and marital status in accounting for structural transformation. We focus here on the years 1940–2010, since education is available in the United States only from 1940 onward. We focus on workers with valid responses to all three questions. We aggregate the detailed variables so that we have broader categories that are easily comparable over time. We have two gender categories (male and female); two marriage categories (married and unmarried, which includes separated, divorced, and widowed); and four education categories (less than primary complete, primary complete, secondary complete, and tertiary complete).

A.2 International Data

Our international sample includes all countries for which we have been able to acquire repeated (at least two) nationally representative cross-sections of microdata that include employment status and industry of employment. We further limit our attention data that include sufficient detail that we can harmonize these key objects in a reasonably consistent

Figure A.1: Countries in Dataset

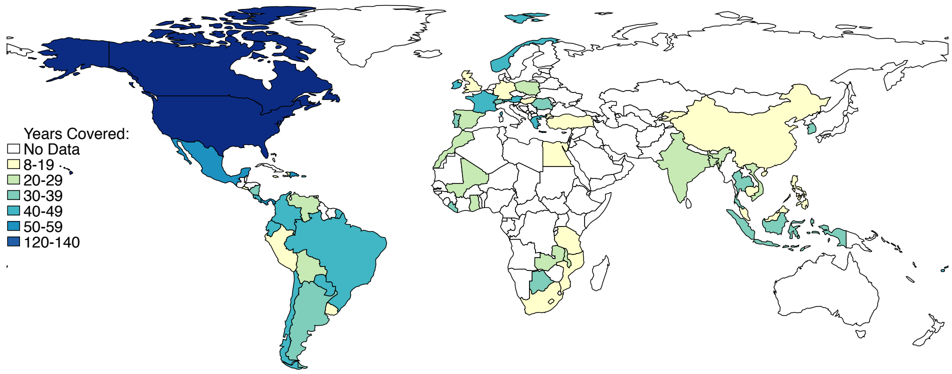

way; we also eliminate a few countries or samples that cover shorter periods (generally, less than eight years) to avoid confusing temporary changes with trends.

Most of our data come from IPUMS. Minnesota Population Center (2014) includes repeated cross-sections for 54 countries with the necessary information. Ruggles et al. (2010) provides the censuses for Puerto Rico. Minnesota Population Center (2017) includes additional cross-sections for Canada as well as new data for Norway from the late 19th and early 20th centuries.[15] Finally, we have identified three countries with independent data of sufficient information that we were able to harmonize and add to this dataset. All told, our dataset for studying worker reallocation includes 201 samples from 59 countries. Figure A.1 shows the countries covered and the total length spanned by country. Our coverage is broad both in terms of geography and PPP GDP per capita. Table A.1 includes a full list of countries and years.

The IPUMS team has devoted a great deal of energy to harmonizing variables and responses across countries and years. The most important for our purposes is that they have re-coded each country's original responses for the industry or sector question (e.g., the one describing the activity or product produced at the respondent's workplace) into a variable they call indgen, which is a slightly modified version of the ISIC 1-digit industry coding scheme. As they note, this coding process is non-trivial, in three main ways. First, for some countries the underlying codes are too coarse to be mapped into indgen at all; these countries are absent from our data. Second, in some countries not all of the original industry codes can be mapped into the indgen classification scheme. Finally, there are inevitably some

[15] Specific data provided by Inwood and Chelsea (2011), Gaffield et al. (2009), The Digital Archive (The National Archive), The Norwegian Historical Data Centre (University of Tromsø), The Digital Archive (The National Archive), and The Digital Archive (The National Archive).

judgment calls when constructing such crosswalks. The main examples described by the IPUMS team involve categories which are small (repair work) or judgment calls that are not relevant for our work (distinguishing among the service industries when mapping an industry). We aggregate these categories into the three broad industry groups as explained in the last subsection.

A.3 Alternative Industry Decompositions

In the text we focus on the classic three-industry description of structural transformation. This industry decomposition is useful for middle income countries, but less so for poor or rich countries, where agriculture or services respectively dominate. Here we show that our results are robust to studying alternative decompositions.

When studying poor countries, it is common to group all of non-agriculture together and focus simply on the transition from agriculture to non-agriculture. We showed in Table 1 that the between share is highest for agriculture. Thus, not surprisingly, this leads to a higher role for the between share in structural transformation, reported in the first row of Table A.2.

For rich countries it might be useful to decompose services, given that it now accounts for a large majority of employment (Duarte and Restuccia, 2017). We follow Herrendorf and Schoellman (2018) by breaking services into unskilled and skilled categories based on average education of the workforce. The former includes primarily personal services, wholesale and retail trade, and hotels and restaurants; the latter includes professional services. Figure A.2 shows graphically how structural transformation looks. Unskilled services display a mixed, possibly inverse-U shape similar to manufacturing, while skilled services grow uniformly. The second row of Table A.2 shows that this extra detail matters little for our basic metric; the between share in this case is 49 percent. To push this point even further, we can use a fifteen sector decomposition (the full detail of indgen as coded in ipums). With this decomposition the between share is still 53 percent. This suggests that the finding of an important role for new cohorts in accounting for structural transformation is robust.

Table A.1: Structural Transformation Sample

Country	N	Years	Country	N	Years
Argentina	4	1970–2001	Austria	5	1971–2011
Bangladesh	3	1991–2011	Bolivia	3	1976–2001
Botswana	4	1981–2011	Brazil	5	1970–2010
Cambodia	2	1998–2008	Canada	7	1891–2011
Chile	5	1960–2002	China	3	1982–2000
Colombia	4	1964–2005	Costa Rica	5	1963–2011
Dom. Republic	5	1960–2010	Ecuador	5	1962–2010
Egypt	2	1996–2006	El Salvador	2	1992–2007
Fiji	5	1966–2007	France	5	1962–2011
Germany (West)	2	1971–1981	Ghana	3	1984–2010
Greece	5	1971–2011	Haiti	2	1982–2003
Hungary	2	2001–2011	India	3	1983–2009
Indonesia	4	1971–2010	Ireland	4	1971–2011
Jamaica	3	1982–2001	Kyrgyzstan	2	1999–2009
Liberia	2	1974–2008	Malawi	3	1987–2008
Malaysia	2	1991–2000	Mali	3	1987–2009
Mexico	5	1960–2015	Morocco	3	1982–2004
Mozambique	2	1997–2007	Nicaragua	3	1971–2005
Norway	4	1865–1910	Palestine*	2	2000–2015
Panama	6	1960–2010	Paraguay	5	1962–2002
Peru	2	1993–2007	Philippines	2	1990–2000
Poland	2	1978–2002	Portugal	4	1981–2011
Puerto Rico	5	1970–2015	Romania	4	1977–2011
South Africa	2	1996–2007	South Korea*	4	1986–2016
Spain	3	1981–2001	Switzerland	4	1970–2000
Tanzania	2	2002–2012	Thailand	4	1970–2000
Trinidad & Tobago	3	1980–2000	Turkey	2	1985–2000
United Kingdom*	3	1997–2014	United States	14	1870–2015
Uruguay	2	1996–2006	Venezuela	3	1981–2001
Vietnam	3	1989–2009	Zambia	3	1990–2010

* Samples derived from independently collected labor force surveys. The remaining samples are from Ruggles et al. (2010), Minnesota Population Center (2014), and Minnesota Population Center (2017).

Table A.2: Between Cohort Share for Alternative Industry Groupings

Number of Broad Industries	Between Share
Two	71%
Four	49%
Fifteen	53%

Figure A.2: Structural Transformation in the United States: 4 Industry View

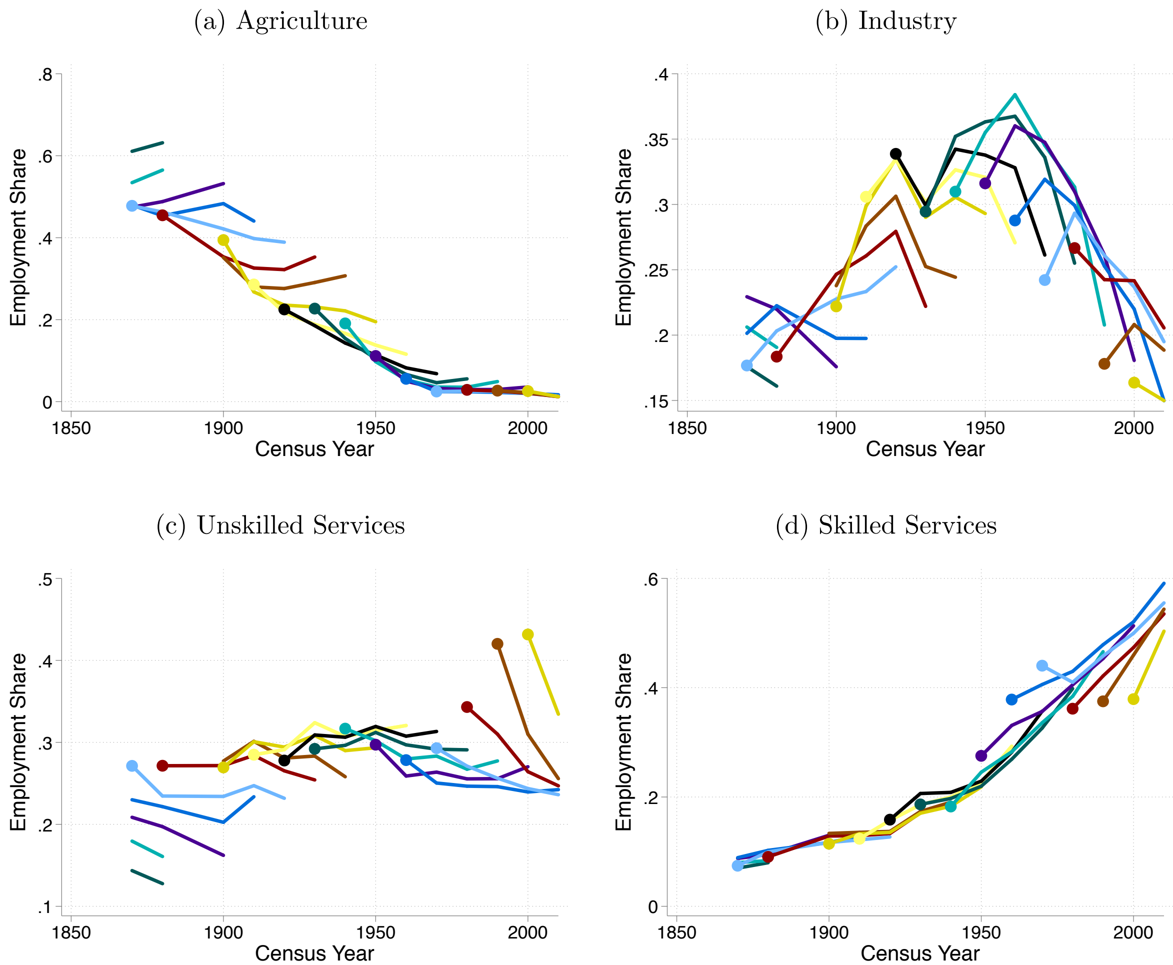

B Solution to career choices

The solution to the dynamic discrete choices that make up the career decisions of workers in this model involves several conditional probabilities and conditional expectations of transformations of independently exponentially distributed random variables. In this section we derive these conditional probabilities and expectations in closed form. Because of the length of these derivations and the resulting expressions, they are omitted from the main text.

The variables we derive in closed form in this section are

- $\Phi_t(i)$, the probability of a worker of age $h = 0$ being trained to work in sector i at time t.
- $\tilde{z}_t(i; 0)$, the average productivity level of a worker trained to work in sector i at time t.
- $\Gamma_t(i, j; h)$, the probability that a worker of age $h > 0$ who works in sector i at time $t - 1$ decides to get retrained and starts working in sector j at time t.
- $\tilde{z}_t(i, j; h)$, the average productivity level of a worker of age $h > 0$ who decided to work in sector j at time t while having worked in sector i in period $t - 1$.

All four of these variables can be derived using two results on the distribution and expectation of the maximum of three linear transformations of exponentially distributed random variables. We derive these results in general first below and then show how they apply to the four variables at hand.

Two main results about exponentially distributed random variables

The two results we consider are about linear transformations of three independently Exponentially distributed random variables, $Z_i \sim \text{Exp}(1)$ where $i = 1, \ldots, 3$. We write these transformations as $X_i = a_i Z_i + b_i$, where $a_i, b_i > 0$.

The first result we use is that for the probability that X_i is the maximum of the sample of three X's, which we denote by π_i. In terms of order statistics, this is

$$\pi_i = P\left[X_i = X_{(3)}\right] = P\left[a_i Z_i + b_i = \max_{j=1,\ldots,3}\{a_j Z_j + b_j\}\right]. \tag{22}$$

Using the notation that $j \neq i$ and $k \neq j$ and $k \neq i$, we can write this probability as

$$\pi_i = \int_{\underline{z}_i}^{\infty}\left(1 - \exp\left(-\frac{a_i Z + b_i - b_j}{a_j}\right)\right)\left(1 - \exp\left(-\frac{a_i Z + b_i - b_k}{a_k}\right)\right)\exp(-Z)\, dZ, \tag{23}$$

where

$$\underline{z}_i = \max_{j=1,\ldots,3}\{\frac{b_j - b_i}{a_i}\}. \tag{24}$$

The above integral can be written in terms of four subintegrals. This yields

$$\begin{aligned}
\pi_i &= \int_{\underline{z}_i}^{\infty} \exp(-Z)\, dZ && (25)\\
&\quad - \exp\left(\frac{b_j - b_i}{a_j}\right) \int_{\underline{z}_i}^{\infty} \exp\left(-\left\{1 + \frac{a_i}{a_j}\right\} Z\right) dZ && (26)\\
&\quad - \exp\left(\frac{b_k - b_i}{a_k}\right) \int_{\underline{z}_i}^{\infty} \exp\left(-\left\{1 + \frac{a_i}{a_k}\right\} Z\right) dZ && (27)\\
&\quad + \exp\left(\frac{b_j - b_i}{a_j} + \frac{b_k - b_i}{a_k}\right) \int_{\underline{z}_i}^{\infty} \exp\left(-\left\{1 + \frac{a_i}{a_k} + \frac{a_i}{a_j}\right\} Z\right) dZ && (28)\\
&= \exp(-\underline{z}_i) && (29)\\
&\quad - \exp\left(\frac{b_j - b_i}{a_j}\right) \left\{\frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_j}}\right\} \exp\left(-\left\{1 + \frac{a_i}{a_j}\right\} \underline{z}_i\right) && (30)\\
&\quad - \exp\left(\frac{b_k - b_i}{a_k}\right) \left\{\frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_k}}\right\} \exp\left(-\left\{1 + \frac{a_i}{a_k}\right\} \underline{z}_i\right) && (31)\\
&\quad + \exp\left(\frac{b_j - b_i}{a_j} + \frac{b_k - b_i}{a_k}\right) \left\{\frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_j} + \frac{1}{a_k}}\right\} \exp\left(-\left\{1 + \frac{a_i}{a_k} + \frac{a_i}{a_j}\right\} \underline{z}_i\right). && (32)
\end{aligned}$$

For the numerical implementation, it is important to also take into account the limiting behavior of this integral when the a's go to zero. If $a_j \downarrow 0$ then (26) and (28) are zero, and the integral only consists of (25) and (27). Similarly, if $a_j \downarrow 0$ then (27) and (28) are zero, and the integral only consists of (25) and (26).

Finally, if $a_i \downarrow 0$ then π_i is not an integral, but instead equals

$$\pi_i = \begin{cases} \left(1 - \exp\left(-\frac{b_i - b_j}{a_j}\right)\right)\left(1 - \exp\left(-\frac{b_i - b_k}{a_k}\right)\right) & \text{if } b_i > b_j \text{ and } b_i > b_k \\ 0 & \text{otherwise} \end{cases}. \tag{33}$$

In this expression, the first term goes to one when $a_j \downarrow 0$ and the second term goes to one when $a_k \downarrow 0$.

We summarize this definition by definining the function

$$\pi_i = \Pi(a_i, b_i, a_j, b_j, a_k, b_k), \tag{34}$$

which is what we will use later on to define two of the four variables of interest.

The second main result that we consider is about the expected value of Z_i conditional on X_i being the third order statistic. We denote this expectation by $\tilde{\zeta}_i$ and, formally, it is defined as

$$\zeta_i = \mathbb{E}\left[Z_i \middle| X_i = X_{(3)}\right] = \mathbb{E}\left[Z_i \middle| a_i Z_i + b_i = \max_{j=1,\ldots,3}\{a_j Z_j + b_j\}\right]. \tag{35}$$

Given this definition, we can write

$$\zeta_i = \frac{1}{\pi_i}\int_{\underline{z}_i}^{\infty} Z\left(1-\exp\left(-\frac{a_i Z + b_i - b_j}{a_j}\right)\right)\left(1-\exp\left(-\frac{a_i Z + b_i - b_k}{a_k}\right)\right)\exp\left(-Z\right)dZ. \tag{36}$$

This, again, can be written in terms of four subintegrals as

$$\begin{aligned}
\zeta_i &= \frac{1}{\pi_i}\int_{\underline{z}_i}^{\infty} Z\exp\left(-Z\right)dZ && (37)\\
&\quad -\frac{1}{\pi_i}\exp\left(\frac{b_j-b_i}{a_j}\right)\int_{\underline{z}_i}^{\infty} Z\exp\left(-\left\{1+\frac{a_i}{a_j}\right\}Z\right)dZ && (38)\\
&\quad -\frac{1}{\pi_i}\exp\left(\frac{b_k-b_i}{a_k}\right)\int_{\underline{z}_i}^{\infty} Z\exp\left(-\left\{1+\frac{a_i}{a_k}\right\}Z\right)dZ && (39)\\
&\quad +\frac{1}{\pi_i}\exp\left(\frac{b_j-b_i}{a_j}+\frac{b_k-b_i}{a_k}\right)\int_{\underline{z}_i}^{\infty} Z\exp\left(-\left\{1+\frac{a_i}{a_k}+\frac{a_i}{a_j}\right\}Z\right)dZ. && (40)
\end{aligned}$$

Each of these subintegrals can be solved by using the result that

$$\int x\exp\left(-bx\right)dx = -\frac{1}{b^2}\left(bx+1\right)\exp\left(-bx\right). \tag{41}$$

Doing so, we obtain that

$$\begin{aligned}
\zeta_i &= \frac{1}{\pi_i}\left(1+\underline{z}_i\right)\exp\left(-\underline{z}_i\right) && (42)\\
&\quad -\frac{1}{\pi_i}\exp\left(\frac{b_j-b_i}{a_j}\right)\left\{\frac{\frac{1}{a_i}}{\frac{1}{a_i}+\frac{1}{a_j}}\right\}^2\left(\left\{1+\frac{a_i}{a_j}\right\}\underline{z}_i+1\right)\exp\left(-\left\{1+\frac{a_i}{a_j}\right\}\underline{z}_i\right) && (43)\\
&\quad -\frac{1}{\pi_i}\exp\left(\frac{b_k-b_i}{a_k}\right)\left\{\frac{\frac{1}{a_i}}{\frac{1}{a_i}+\frac{1}{a_k}}\right\}^2\left(\left\{1+\frac{a_i}{a_k}\right\}\underline{z}_i+1\right)\exp\left(-\left\{1+\frac{a_i}{a_k}\right\}\underline{z}_i\right) && (44)\\
&\quad +\frac{1}{\pi_i}\exp\left(\frac{b_j-b_i}{a_j}+\frac{b_k-b_i}{a_k}\right)\left\{\frac{\frac{1}{a_i}}{\frac{1}{a_i}+\frac{1}{a_j}+\frac{1}{a_k}}\right\}^2\times && (45)\\
&\quad \left(\left\{1+\frac{a_i}{a_j}+\frac{a_i}{a_k}\right\}\underline{z}_i+1\right)\exp\left(-\left\{1+\frac{a_i}{a_k}+\frac{a_i}{a_j}\right\}\underline{z}_i\right).
\end{aligned}$$

Again, it is useful to consider the limiting cases. When $a_j \downarrow 0$ then (38) and (40) go to zero and the integral just consists of parts (37) and (39). When $a_k \downarrow 0$ then (39) and (40) go to zero and the integral consists of (37) and (38). When $a_i \downarrow 0$ then the value of Z_i does not matter for the career choice. As a result, in that case $\zeta_i = \mathbb{E}\left[Z_i\right] = 1$.

We summarize this definition by definining the function

$$\zeta_i = \mathcal{Z}\left(a_i, b_i, a_j, b_j, a_k, b_k\right), \tag{46}$$

which is what we will use later on to define two of the four variables of interest.

Career-choice probabilities and expected productivity levels

To shorten the notation, it turns out to be useful to define

$$\tilde{V}_{t+1}\left(i,h\right) = \frac{1-\delta}{1+r}\mathbb{E}_t V_{t+1}\left(i,h+1;\mathbf{z}_{t+1}\right). \tag{47}$$

Using the above definition and the results derived in the subsection above, we can now write

$$\Phi_t\left(i\right) = \Pi\Big(\left(1-\phi\right)w_{i,t}, \tilde{V}_{t+1}\left(i,h\right), \left(1-\phi\right)w_{j,t}, \tilde{V}_{t+1}\left(j,h\right), \left(1-\phi\right)w_{k,t}, \tilde{V}_{t+1}\left(k,h\right)\Big), \tag{48}$$

as well as

$$\tilde{z}_t\left(i;0\right) = \mathcal{Z}\Big(\left(1-\phi\right)w_{i,t}, \tilde{V}_{t+1}\left(i,h\right), \left(1-\phi\right)w_{j,t}, \tilde{V}_{t+1}\left(j,h\right), \left(1-\phi\right)w_{k,t}, \tilde{V}_{t+1}\left(k,h\right)\Big). \tag{49}$$

Moreover, we can write

$$\Gamma_t\left(i,i;h\right) = \Pi\Big(w_{i,t}, \tilde{V}_{t+1}\left(i,h\right), \left(1-\gamma_h\right)w_{j,t}, \tilde{V}_{t+1}\left(j,h\right), \left(1-\gamma_h\right)w_{k,t}, \tilde{V}_{t+1}\left(k,h\right)\Big), \tag{50}$$

and

$$\Gamma_t\left(i,j;h\right) = \Pi\Big(\left(1-\gamma_h\right)w_{j,t}, \tilde{V}_{t+1}\left(j,h\right), w_{i,t}, \tilde{V}_{t+1}\left(i,h\right), \left(1-\gamma_h\right)w_{k,t}, \tilde{V}_{t+1}\left(k,h\right)\Big), \tag{51}$$

when $j \neq i$.

The associated expected productivity levels are

$$\tilde{z}_t(i,i;h) = \mathcal{Z}\left(w_{i,t}, \tilde{V}_{t+1}(i,h), (1-\gamma_h)\,w_{j,t}, \tilde{V}_{t+1}(j,h), (1-\gamma_h)\,w_{k,t}, \tilde{V}_{t+1}(k,h)\right), \tag{52}$$

and

$$\tilde{z}_t(i,j;h) = \mathcal{Z}\left((1-\gamma_h)\,w_{j,t}, \tilde{V}_{t+1}(j,h), w_{i,t}, \tilde{V}_{t+1}(i,h), (1-\gamma_h)\,w_{k,t}, \tilde{V}_{t+1}(k,h)\right), \tag{53}$$

when $j \neq i$.

C Calibration of flexible benchmark parameters

For the dynamics of employment shares in the model, what matters is the product of the preference parameter, λ_i, and the technology parameter, A_i. For this reason, we normalize the preference parameters[16] to be equal across industries

$$\lambda_a = \lambda_m = \lambda_s = \frac{1}{3}. \tag{54}$$

Moreover, because preferences are homothetic and the production technology exhibits constant returns to scale, the absolute levels of the technology parameters (A_i) do not matter for the equilibrium dynamics, but their relative levels. With that in mind, we also normalize

$$A_a = 1. \tag{55}$$

These two normalizations leave six parameters to be pinned down based on historical U.S. data, namely A_m, A_s, g_a, g_m, g_s, and ε. We choose them such that the flexible benchmark model matches the following historical facts for the U.S. economy:

- *Employment shares in 1870 and 2010:* The percent of workers employed in agriculture, manufacturing, and services, at both the beginning, i.e. 1870, and end, i.e. 2010, of the historical sample that we consider. These shares are taken from the Decennial U.S. Censuses digitized and made available by Minnesota Population Center (2017). Because employment shares add up to one, this provides four moment conditions, two in 1870 and two in 2010, to match.

[16] This normalization affects the value added shares in GDP, which we do not focus on. It does not affect any of the dynamics that we analyze in this paper.

- *Average real GDP per capita growth:* The average annualized growth rate of real GDP per capita from 1870 through 2010, from Bolt and van Zanden (2014). This is the fifth moment we match.

- *Estimates of elasticity of substitution:* Estimates of the elasticity of substitution, ε, based on postwar U.S. national income data reported in Ngai and Pissarides (2008). This is the final moment condition.

In practice, this means that we follow a procedure similar to the one used by Ngai and Pissarides (2008). First, just like Ngai and Pissarides (2008), we choose $\varepsilon = 0.1$ which is consistent with postwar U.S. NIPA data. We then find the values for A_m and A_s under which the employment shares in the flexible benchmark match the 1870 employment shares in the data. Next, we find the relative growth rates of manufacturing and services, i.e. $g_m - g_a$ and $g_s - g_a$, to match the 2010 employment shares. Finally, we choose g_a such that the flexible benchmark matches the average growth rate of real GDP per capita from 1870 through 2010. The values of the non-normalized parameters that we obtain using this method, are $g_a = 0.045$, $g_m = 0.020$, and $g_s = 0.013$ for the annualized TFP growth rates, and $A_m = 2.970$, and $A_s = 1.638$ for the initial relative productivity levels of manufacturing and services.

In terms of unmatched moments in the data, the implied growth rate of agriculture that results from this calibration is about a percentage point higher than the actual growth rate from merged data from Kendrick (1961) and postwar Industry Accounts of the Bureau of Economic Analysis. When $\varepsilon > 0.1$ the implied growth rate of agriculture is even higher.

Figure C.1 plots the implied time-series paths of the employment shares in the model and the data. These shares are equal in 1870 and 2010 by construction.

D Details about Flexible Benchmark

The flexible benchmark in which (re-)training costs are zero, i.e. $\phi = \gamma_h = 0$ for $h = 1, \ldots, H$ is our point of comparison for many of our results. Therefore, we describe the equilibrium path under the flexible benchmark here in more detail. In particular, we focus on the results plotted in Figure D.1. The flexible benchmark equilibrium in this economy is similar to that in Bárány and Siegel (2018) and we touch on many of the same qualitative properties here that Bárány and Siegel (2018) discuss in much more detail.

Figure C.1: Path of employment shares in flexible benchmark and data

Percent of workers employed in each sector in model and data, 1870-2010

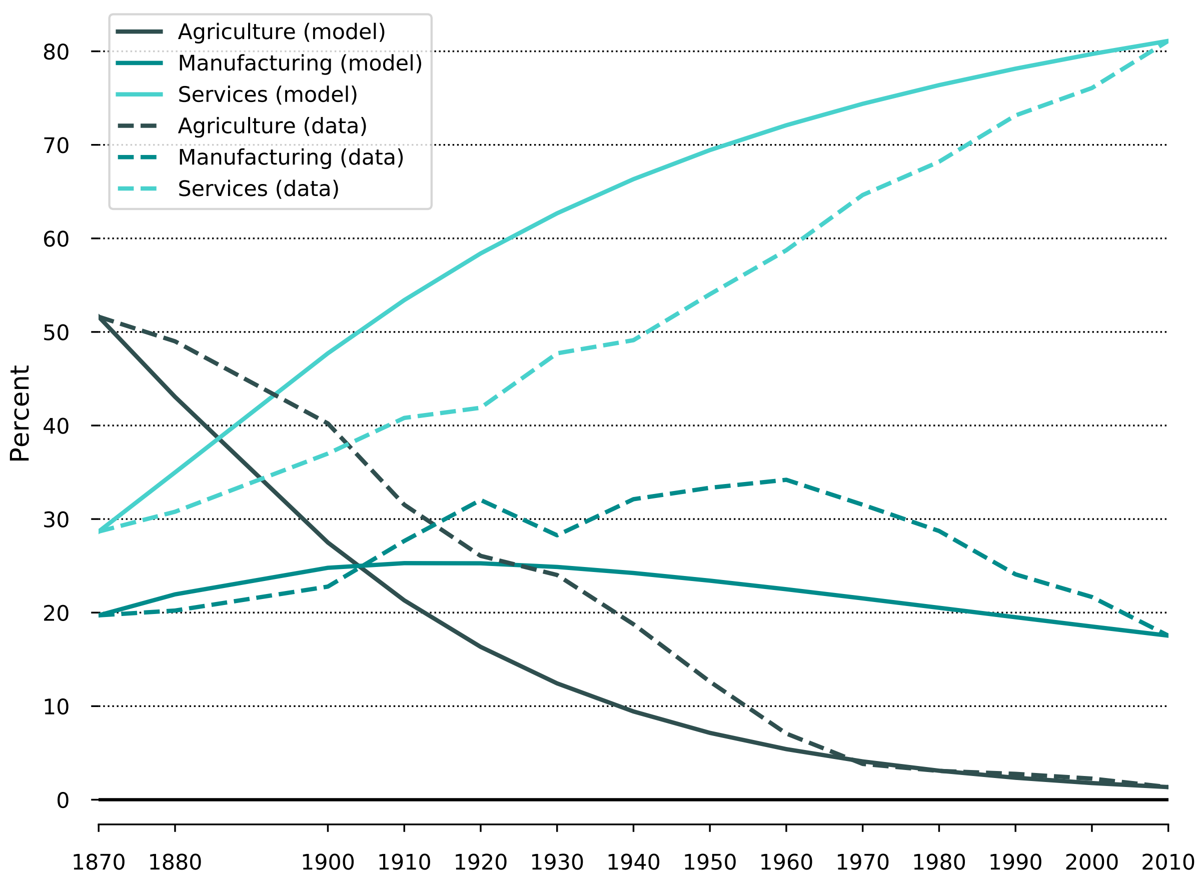

Source: Minnesota Population Center (2017) and authors' calculations

Because the outputs of the three sectors are gross complements, the relative wage of workers in the services sector is increasing over time. This is because these workers are getting hired to provide the services that complement that manufacturing and agricultural output for which relatively little labor is need to produce. The trends in wages is depicted in Panel (a) of Row (i) of Figure 4 in the main text.

The increasing wage gap draws more and more workers into the service sector over time, as can be seen from Figure D.1a. In the intermediate stages of the transition the employment share of manufacturing peaks while that of agriculture monotonically declines.

Even though a larger fraction of workers is drawn into services over time, this does not translate into higher output growth in services because of the lower TFP growth rate in services, g_s, than in manufacturing, g_m and agriculture, g_a. Moreover, as Ngai and Pissarides (2007) emphasize, the CES preferences with $\varepsilon < 1$ imply that the relative price, as well as its value-added share in GDP, of the good that grows most slowly will rise over time. As a result, the slowest growing sector makes up an increasing part of GDP over time and drags down overall GDP growth. This can be seen from Figure D.1b where aggregate GDP increasingly aligns with services output over the transition path. In this sense, our model economy suffers from the Cost Disease described in Baumol and Bowen (1968).

But it is not only the exogenously lower rate of TFP growth, g_s, that drags down output growth in the service sector. Average labor productivity growth in this model is determined by both the exogenous rate of TFP growth as well as the selection of workers into different sectors. Just like in Bárány and Siegel (2018) the average quality of workers in the service sector declines over time. This is because the increasing wedge in real wages between services and the other sectors draws workers of lower quality, $\tilde{z}_s$, into services over time. This can be seen from Figure D.1c, which plots the index of $\tilde{z}_i$ for $i \in \{a, s, m\}$, as well as the average for the whole economy (aggregate). At the end of the quarter millenium transition that we consider this selection of workers across sectors increases labor productivity in agriculture by a factor of two, while reducing that in services by about 40 percent. Because the service sector becomes the predominant sector over time, this decline in labor quality in services drags down aggregate labor quality and, with it, growth in average labor productivity.

This selection effect of workers across sectors in small, however, when compared to the productivity growth rates experienced by each of the sectors. This can be seen by comparing the scale of the vertical axis in Figure D.1c with that of Figure D.1d. The latter plots the log-indices of average labor productivity for the three sectors in the economy as well as aggregate labor productivity. Notice how the Cost Disease of Baumol and Bowen (1968) results in a perpetual productivity slowdown.

Figure D.1: Dynamics of flexible benchmark

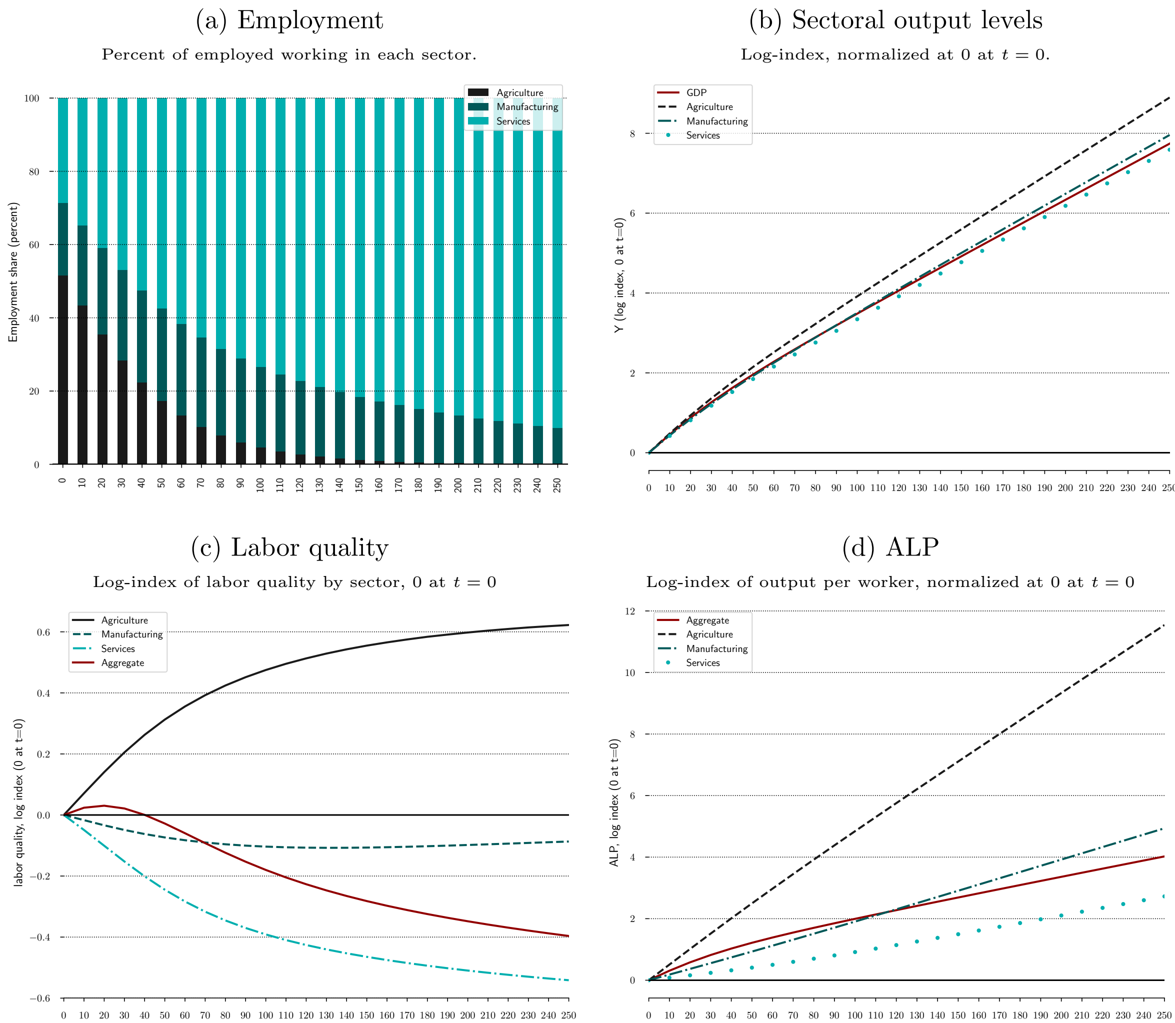

Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010.

E Implementation of extended-path method

For our solution method, we consider the transitional path from $t = 0 - \tilde{t}_l$ until $t = T + \tilde{t}_r$. Here $\tilde{t}_l$ and $\tilde{t}_r$ are padding of the extended path that allow for startup and wind-down periods on the path. We present results for $t = 0, \ldots, T$. For $t > T + \tilde{t}_r$ we assume the economy is on a balanced growth path in which $g_i = g > 0$ for $i \in \{a, m, s\}$. Moreover, we assume that, after $t > T + \tilde{t}_r$ workers do not need to spend time on training and retraining anymore. Hence, our solution method isolates the importance of these costs along the transitional path where structural transformation occurs.[17]

At each point in time, t, the relevant state of the economy consists of three parts. The first are the sector-specific TFP levels, $\{A_{i,t}\}_{i \in \{a,m,s\}}$. The second is the size of the new cohort, $N_t(0)$. The final part is the labor supply, that consists of different age-industry-specific levels $\{\{E_{t-1}(i; h)\}_{h=1}^{H}\}_{i \in \{a,m,s\}}$.

The TFP levels and initial cohorts evolve exogenously over time, according to

$$A_{i,t} = (1 + g_i) A_{i,t-1}, \text{ where } A_{i,0} \text{ is given and } i \in \{a, m, s\}, \tag{56}$$

and

$$N_t(0) = (1 + n) N_{t-1}(0). \tag{57}$$

Our solution method is used to solve for the endogenous evolution of $\{\{E_t(i; h)\}_{h=0}^{H}\}_{i \in \{a,m,s\}}$ in equilibrium along the transitional path for $t = -\tilde{t}_l + 1 \ldots T + \tilde{t}_r$. Note that the initial age-industry-specific levels of the labor supply, at time $t = 0$, are given. Thus, we solve the transitional dynamics of the model conditional on the initial state $\{\{E_0(i; h)\}_{h=1}^{H}\}_{i \in \{a,m,s\}}$. Our method loops over a backward and a forward recursion until reaching convergence.

Backward recursion: Update career choices conditional on labor supply

This recursion starts at time $t = T + \tilde{t}_r$ and runs backwards to $t = -\tilde{t}_l$. It takes the path of the age-industry levels of the labor supply, $\{\{E_t(i; h)\}_{h=0}^{H}\}_{i \in \{a,m,s\}}$ for $t = -\tilde{t}_l + 1 \ldots T + \tilde{t}_r$, as given.

Labor market equilibrium at time t: The main step in this backward recursion is to solve for the equilibrium in the three labor markets at time t, taking as given the age-industry

[17] We have checked the robustness of our results for the choice of T and g as well as the length of padding $\tilde{t}_l$ and $\tilde{t}_r$.

specific levels of the labor supply, $\{\{E_{t-1}(i;h)\}_{h=1}^{H}\}_{i\in\{a,m,s\}}$, as well as the age-industry specific career continuation values, $\{\{\mathbb{E}_t V_{t+1}(i;h)\}_{h=1}^{H}\}_{i\in\{a,m,s\}}$. This requires solving for real wages, $\{w_{i,t}\}_{i\in\{a,m,s\}}$, using a system of 3 equations.

Taking as given the continuation values, $\{\{\mathbb{E}_t V_{t+1}(i;h)\}_{h=1}^{H}\}_{i\in\{a,m,s\}}$, a given set of three wages $\{w_{i,t}\}_{i\in\{a,m,s\}}$ pins down the optimal career decisions that determine $\Phi_t(i)$, $\tilde{\phi}_t(i)$, $\Gamma_t(i,j;h)$, and $\tilde{\gamma}_t(i,j;h)$ through the solution described in Section B. With this, we can find the equilibrium wages $\{w_{i,t}\}_{i\in\{a,m,s\}}$ that clear the labor markets, (21).

We evaluate $\Phi_t(i)$, $\tilde{\phi}_t(i)$, $\Gamma_t(i,j;h)$, and $\tilde{\gamma}_t(i,j;h)$ at these equilibrium wages and the continuation values $\{\{\mathbb{E}_t V_{t+1}(i;h)\}_{h=1}^{H}\}_{i\in\{a,m,s\}}$. Then, we iterate backwards. Because of the idiosyncratic nature of the (re-)training costs, we can write

$$\mathbb{E}_{t-1}V_t(0) = \sum_{i\in\{a,m,s\}} \Phi_t(i)\left((1-\phi)\,\tilde{z}_t(i)\,w_{i,t} + \frac{1-\delta}{1+r}\mathbb{E}_t V_{t+1}(i;1)\right). \tag{58}$$

and for $h = 1,\dots,H-1$

$$\mathbb{E}_{t-1}V_t(i;h) = \sum_{j\in\{a,m,s\}} \Gamma_t(i,j;h)\left[(1-\mathbb{I}(i\neq j)\,\gamma)\,\tilde{z}_t(i,j;h)\,w_{i,t} + \frac{1-\delta}{1+r}\mathbb{E}_t V_{t+1}(j;h+1)\right]. \tag{59}$$

Finally, we have

$$\mathbb{E}_{t-1}V_t(i;H) = \sum_{j\in\{a,m,s\}} \Gamma_t(i,j;H)\,(1-\mathbb{I}(i\neq j)\,\gamma)\,\tilde{z}_t(i,j;H)\,w_{j,t}. \tag{60}$$

These equations now allow us to solve the relevant continuation values using a backward recursion. Of course, this solution is conditional on $\{\{E_t(i;h)\}_{h=0}^{H}\}_{i\in\{a,m,s\}}$. With these continuation values for $t-1$ in hand we can now solve the labor market equilibrium for period $t-1$ conditional on $\{\{E_{t-1}(i;h)\}_{h=0}^{H}\}_{i\in\{a,m,s\}}$ using the same method and role this backward recursion all the way from $t = T + \tilde{t}_r$ to $t = -\tilde{t}_l$.

Forward recursion: Update labor supply conditional on career choices

In the forward recursion we now update the path of the age-industry specific levels of the labor supply, $\{\{E_t(i;h)\}_{h=1}^{H}\}_{i\in\{a,m,s\}}$, using the career choices solved in the backward recursion. This simply involves iterating over the law of motion of the labor supply from (17) and (18). This recursion is initialized using the initial condition that gives the age-industry specific labor supply levels at time $t = 0$, i.e. $\{\{E_0(i;h)\}_{h=1}^{H}\}_{i\in\{a,m,s\}}$. We continue to loop over the backward and forward recursions until the calculated path of the age-industry specific labor supply levels, $\{\{E_t\,(i;h)\}_{h=1}^{H}\}_{i\in\{a,m,s\}}$ for $t = -\tilde{t}_l + 1 \ldots T + \tilde{t}_r$, converges.