

Long-Run Sectoral Reallocation, Job to Job Transitions, and Earnings Inequality: An Empirical Investigation*

Florian Hoffmann
University of British Columbia
(florian.hoffmann@ubc.ca)

Anton Laptiev
University of British Columbia
(laptiev.anton@gmail.com)

Shouyong Shi
Pennsylvania State University
(sus67@psu.edu)

August 2016

Abstract

Using data from the Census, ACS and the CPS, we document that a large part of long-run reallocation of labor from the goods to the service sector took place within narrowly defined groups of workers. In particular, sectoral reallocation reflects a labor market trend that is distinct from the automatization of the goods sector and the increase in female labor force participation. A relative increase in direct monthly job-to-job transitions from the goods to the service sector can explain a sizeable share of the rise in service-sector employment. This finding is robust to the methodology of adjusting for trend breaks in monthly transition rates as measured from the CPS. We use this result to test a central empirical prediction of job-search theory with on-the-job search: that there is an intrinsic link between the earnings structure and the rates at which employed workers change jobs. Using two different empirical strategies and data disaggregated to various levels we find robust evidence that the earnings distribution in the service sector became more unequal relative to the goods sector when the rate at which workers flew into the service sector rose. This relationship is particularly strong among age groups with stable labor force participation rates. We document some evidence based on the distribution of earnings changes computed from the SIPP against the hypothesis that our results are driven by worker sorting on unobserved heterogeneity.

*Hoffmann and Laptiev: Vancouver School of Economics, University of British Columbia, 6000 Iona Drive, Vancouver, B.C., Canada, V6T 1L4. Shi: Department of Economics, Pennsylvania State University, 502 Kern Building, University Park, PA 16802, USA. This research was supported by the Social Sciences and Humanities Research Council of Canada. Shi gratefully acknowledges the Pennsylvania State University for financial support.

1. Introduction

The reallocation of labor from the goods producing sector to the service sector has been one of the most pronounced aggregate labor market trends of the last half a century. While the aggregate employment share of the service sector was 56 percent in 1960, it has risen to almost 77 percent in 2007, as shown in Figure 1.¹ This sectoral shift took place at a time when several other aggregate changes have shaped the labor market as well, such as the sharp increase of educational attainment of males until the 1970s, the steady rise in female labor force participation and educational attainment over much of the second half of the twentieth century, and the deroutinization of the occupational structure since the 1980's. While many of these changes have been studied extensively, often within the paradigm of frictionless markets, evidence on the impact of sectoral reallocation on individual labor market outcomes is scarce.² In particular, there is no systematic empirical study of the association between sectoral reallocation and worker mobility within and across sectors on the one hand and potential linkages to the evolution of the sectoral earnings structures on the other hand.

This gap is noteworthy for at least two reasons. First, the trend towards a large labor share of the service sector has been remarkably steady and gradual, with a nearly linear-in-levels trend between 1960 and 2010. It is reasonable to conjecture that such a pronounced aggregate economic force – a situation of perpetual economic turbulence, associated with new chances for many, but also with new risks for others – had a major impact on labor market outcomes on the micro-level, such as the sector-specific earnings structure and the rate at which individuals change jobs. Indeed, relative average log-earnings between the two sectors have been far from constant, and the same is true for relative earnings inequality. As shown in Figure 2, differences in average log-earnings between the service sector and the non-service sector for the fully employed, computed from Census- and ACS-data, declined slightly over the first two decades of the Census sample and increased steadily thereafter. The increase since 1980 has been large: Earnings in the service sector increased by almost 10 percentage points relative to the non-service sector. The corresponding differences in the variance of log earnings followed a similar U-shaped pattern.³ Furthermore, the post-1980 trends were accompanied by a dramatic rise of the rate at which workers moved from the non-service to the service sector relative to the corresponding worker transition rate in the opposite direction. As documented in figure 3, the probability that a

¹The figure shows the time series constructed from two sources of data, the Census/ACS and the March CPS. We include all individuals between 18 and 64 years of age. We will exclude the post-2007 period from our analysis due to the impact of the Great Recession that is unlikely to be related to sectoral reallocation. It is worthwhile noting that the service sector share has continued to increase in the aftermath of the Great Recession. We use the same industry-based classification of the service sector as Lee and Wolpin (2006).

²We offer a systematic review of the literature in the next section.

³A detailed description of our data is provided in a separate section. At this point we note that for most of our analysis we use annual earnings for the fully employed rather than hourly wages since data on hours of work are usually quite noisy and missing for many observations. Our results hold when we use hourly wages instead.

worker made a direct inter-sectoral job transition into the service sector in two consecutive months, computed from the monthly CPS and adjusted for measurement error, has increased more than threefold since 1980.⁴ In contrast, the corresponding probability for transitions into the non-service sector has remained constant.

Second, these empirical regularities make sectoral reallocation a particularly interesting context for testing models with labor market frictions against a frictionless neo-classical multi-sector economy. The distinctive feature of frictional models with on-the-job search, like Burdett and Mortensen’s (1998) framework, is their ability to endogenize job transition rates and to generate an intrinsic link between worker flows and the *entire* earnings structure.⁵ It is then precisely at times of large and persistent reallocation at which the impact of labor market frictions on individual careers can be discerned. This seems to be consistent with the empirical facts documented above, which suggest a link between a marked change in the sectoral earnings structures on the one hand and the patterns of job-to-job flows on the other hand.

This paper is an attempt to fill this gap in the literature. We provide a detailed empirical analysis of the relationship between sectoral reallocation, sectoral earnings structures and employment stocks and flows, and we explore systematically if our findings are consistent with the predictions of a two-sector model with labor market frictions and on-the-job search. To accomplish this task we proceed in several steps, thereby relying on publicly available data sets including the Census, ACS, CPS and SIPP.

In a first step we use Census and ACS to decompose the increasing share of employment in the service sector – a stock variable – into a part that can be explained by changes in worker and job characteristics and a residual “unexplained” part that we call *pure sectoral reallocation*. We find that the last component contributes to a quarter of the rising employment share in the service sector for the entire sample period and to almost the entire rise in the employment share in the post-1990 period. Thus compositional changes of the labor force alone, including the change in the occupational structure as discussed for example in Acemoglu and Autor (2011), cannot explain the observed extent of sectoral reallocation. Importantly, older male workers were among the groups that experienced particularly large increases, especially in the period starting in 1980, contrary to the common belief that the rise of service sector employment came mainly from labor market entrants.

In a second step we quantify the importance of direct job-to-job transitions for the rise of the service sector, using counterfactual experiments on the parameters of a dynamic system of labor flow

⁴We describe the details of our adjustment procedure in section 4.3, which combines various approaches existing in the literature with a regression-discontinuity like regression adjustment for the 1994 trend break. For this reason, it is unlikely that the peak of transition rates into the service sector in the mid 1990s that is clearly visible in figure 3 is generated by this trend-break. Rather, it coincides with the tech-boom and thus may reflect a true temporary increase in transition rates.

⁵We present a two-sector version of this model in section 6, building on Hoffmann and Shi (2016).

variables. This dynamic system features a monthly transition matrix of workers within and between the two sectors as well as between different states of labor force participation, which we compute from the monthly CPS starting in 1976. Given the notorious difficulties in measuring flows from these data we use various methods that correct for measurement error, including a regression discontinuity-like approach for cleaning the data from a major trend break in 1994.⁶ The central finding is that increases in intersectoral labor flows had a substantially larger contribution to the increased share of employment in the service sector than inflows from unemployment or non-employment. This contrasts sharply with findings in Cortes, Jaimovich, Nekarda and Siu (2015) who document that *occupational* reallocation took place predominantly via transitions from unemployment and non-employment. We interpret this as further evidence that sectoral reallocation and occupational reallocation are, at least to some extent, two distinct economic phenomena.

In a third step, these findings lead us to testing a central prediction of equilibrium job search models with on the job search, namely that the rate at which workers switch jobs and the earnings structure are intimately connected. We first analyse if the relative increase in service sector earnings inequality can be explained by composition effects due to sorting of workers on observed worker or job characteristics. This is important because job search models can endogenize residual earnings inequality, which is the leftover part of inequality after adjusting for composition effects. We thereby rely on modern statistical decomposition methods that can be applied to any functionals of distributions, not only first and second moments. We find that compositional effects had a negligible effect on the evolution of sectoral earnings structures. Most of the increase in service-sector earnings inequality relative to the non-service sector cannot be explained by variables such as gender, age, education, or the routine-content of an occupation.

We then bring together our analysis of employment stocks and flows on the one hand and of the sectoral earnings structure on the other hand, thereby explicitly relating changes in worker flows and changes in the earnings structure. We use two different empirical approaches. The first approach is based on data from repeated cross-sections, capturing sectoral reallocation by changes in the employment share of the service sector. This approach has the advantage that we can rely on the data from the Census and the ACS which have very large sample sizes. As a consequence, we can carry out our analysis with high precision by running regressions in first differences or differences-in-differences on data that are disaggregated to a very fine level. The most important finding is a strong and robust link between within-group sectoral reallocation and changes in within-group earnings inequality. In particular, groups that reallocated the fastest were, on average, those who saw the most pronounced changes in intersectoral differences of the earnings structure. This result is particularly strong among age groups with a stable labor force participation rate in the post-1980 period.

⁶None of our main findings are sensitive to the choice of method.

The primary disadvantage of our first strategy is that it does not relate measures of direct job-to-job flows between sectors to sectoral earnings structures. It thus does not map well into job search models. Our second strategy therefore matches our transition rates constructed from the monthly CPS to measures of intersectoral differences in earnings structures computed from the Census and ACS. The central problem of this approach is that cell sizes in the CPS become small so that parameter estimation may suffer from imprecision. We therefore proceed in the spirit of a recent literature on industrial composition and *local* labor market outcomes, such as Beaudry, Green and Sand (2012) and Chodorow-Reich and Wieland (2015), and disaggregate the data to the state-level instead. We find a strong and robust link between intersectoral labor flows and the change in the relative earnings structure. Conditional on flows into and out of unemployment and NILF and conditional on worker observables, states in which service sector worker recruitment from the goods sector has accelerated particularly fast relative to the recruitment in the other direction have seen a stronger divergence in the sector specific earnings structures.

It is important to note that by the very nature of our question it is hard, if not impossible, to come by quasi-experimental variation in the explanatory variable of interest. For example, a recent literature relating short- to medium run fluctuations in wages and the employment structure uses initial conditions of the employment structure as instrumental variables. Because of the focus of our study on long-run reallocation this strategy cannot be used since we want to use the entire sample period for which we have time-consistent variables. Interestingly, it is close to the "initial condition" between 1960 and 1980 where sectoral reallocation was dominantly driven by compositional factors. It was also exactly during this time period when there was no substantial relative rise in residual inequality in the service sector. Only when the service sector started to recruit increasingly from the non-service sector did the earnings structure diverge. This finding documents that using data from the entire sample period is highly informative and insightful.

To strengthen the evidence for a deep relationship between intersectoral transition rates and the sectoral earnings structure we therefore need to complement the main part of our empirical analysis with a more descriptive approach to rule out alternative explanations. The central cause for concern we need to address is that our findings are driven by worker sorting on unobserved heterogeneity. We thus complete our empirical investigation with an exploration of the sources of earnings growth, using data from the SIPP for the time period 1996 to 2003. Most importantly we find that intersectoral job-to-job transitions are associated with significant earnings gains, while the distribution of earnings growth within job is centered at the origin. If worker skills affect both earnings levels and earnings growth, then the latter finding is inconsistent with the service sector increasingly attracting the highest skilled workers. At the same time, workers switch jobs at a significantly higher rate in the service sector

than in the goods sector. This is exactly what one would expect in standard models of the job ladder, such as Burdett and Mortensen (1998). Neoclassical models of the labor market with worker sorting cannot explain these findings. We therefore argue that our estimates are not likely to be driven entirely by omitted variable bias from worker sorting. More generally, we conclude that our findings are difficult to reconcile with frictionless models of the labor market, even in the presence of worker heterogeneity, or with job-search models that do not feature direct job-to-job transitions, such as Pissarides and Mortensen (1994) or Shimer (2005a). Rather, a unified theory of earnings dynamics, job flows and the rise of the service sector needs to incorporate sectoral reallocation via EE flows.

Our study uses a very different approach from the current empirical literature that tests for a link between aggregate economic changes and individual-level labor market outcomes. This literature has predominantly focused on the effect of business cycle fluctuations on career outcomes on the one hand and on the role of occupational mobility in generating trends in aggregate earnings inequality on the other hand. Compared to this approach, studying the relationship between sectoral reallocation and the rising earnings inequality in the service sector relative to the goods sector instead has a number of advantages. First, job-to-job transitions that are likely to be associated with a change in the wage structure via the job ladder take place across two large sectors rather than many small firms. This allows for a fine-grained analysis of net reallocation within narrowly defined subgroups of the population, which we exploit extensively in our empirical approach and which does not have a counterpart in the literature relating worker flows, the earnings structure and business cycle fluctuations. Furthermore, intersectoral worker transitions are easier to measure than firm-to-firm transitions because they do not require firm identifiers. This is crucial in empirical analyses covering a long time period on the monthly frequency for which U.S. data with firm identifiers do not exist.

Second, sectoral reallocation is likely to be less related to unobserved skills than other types of labor reallocations, such as occupational mobility. Rather, intersectoral transitions within occupational groups, possibly defined by their task content, are more likely to be related to changes in aggregate demand for goods rather than a change in the aggregate demand for skills. For example, clerical occupations in the goods and service sectors attract a more homogenous group of workers than clerical versus non-clerical occupations. Consistent with this argument, there seems to be a consensus that sector-specific human capital accumulation is unimportant once one controls for occupation-specific human capital.⁷ As a consequence, changes in the sector-specific earnings structures are more likely due to deviations from the frictionless markets assumption. But even if sector-specific earnings distributions are, to some extent, reflecting worker sorting on unobserved heterogeneity, sectoral reallocation provides

⁷See e.g. Kambourov and Manovskii (2009).

a context in which a neoclassical two-sector model can be tested directly from the data. In particular, any model of sorting on unobserved traits in a frictionless environment will predict that the declining sector becomes less unequal because only the workers that are best suited for that sector will stay. This is a powerful argument that explicitly relies on splitting the data into worker groups defined by their sectoral affiliation.

Third, sectoral reallocation is more straightforward to measure than business cycle fluctuations, and it generates substantial time-series variation in earnings structures and worker flows that can be used to test various theories of worker flows and earnings dynamics. Specifically, *cross-sectional* interindustry wage differentials, as documented for example in Krueger and Summers (1988) and Gibbons and Katz (1992), are consistent with various economic theories, such as compensating wage differentials, worker sorting or rent sharing. In contrast, *changes* in interindustry wage differentials and their relationship to *changes* in worker flows impose additional structure on competing theories that can be used to test them against each other. For example, we document that service sector wages have risen relative to the non-service sector since the 1980's within narrowly defined groups. At the same time, the intersectoral worker transition rates into the service sector have increased dramatically relative to the intersectoral transition rates in the opposite direction, ruling out compensating wage differentials as a primary mechanism.⁸ To the best of our knowledge, our study is the first to offer a systematic investigation of the relationship between changes in intersectoral worker flows and changes in sectoral earnings structures. We thereby test explicitly the empirical predictions of a two-sector Burdett-Mortensen model of equilibrium search, as formulated in Hoffmann and Shi (2016). These predictions are generated by a combination of strategic interactions between firms in different sectors competing for the same workers, search frictions and monopsony power, and cannot be generated by *any* frictionless neoclassical two-sector model of the economy because intersectoral transition rates in such models are undefined.

The rest of the paper is organized as follows. In section 2 we relate our paper to the literature. We then provide an overview of data construction. Further issues with data and variable construction are discussed as they come up in subsequent sections. Section 4 presents results from a systematic study of employment stocks and intersectoral labor flows, while section 5 focuses on the decomposition of sectoral earnings structures. Section 6 brings together our analysis of employment stocks and flows on the one hand and differentials in sector specific earnings distributions. Section 7 presents descriptive evidence from the SIPP against the hypothesis that our results are driven by selection on unobservables, while section 8 offers some concluding remarks.

⁸The importance of intersectoral worker flows for testing various theories of interindustry wage differentials has been recognized at least since Krueger and Summers (1988) and Gibbons and Katz (1992).

2. Relation to the Literature

A small, but steadily growing literature on the estimation of structural equilibrium search models, such as van den Berg and Ridder (1998), Cahuc, Postel-Vinay and Robin (2006), Bagger et al. (2014), Menzio, Telyukova and Visschers (2016) and Lise and Robin (2014), quantifies frictional earnings dispersion, that is earnings inequality that is due to labor market search frictions rather than worker heterogeneity or informational frictions. Central to the existence of frictional earnings inequality is that workers have the opportunity to search when employed, as shown formally in Burdett and Mortensen (1998) and discussed in book-length in Mortensen (2005). While on-the-job search provides the strategic link between firms' wage posting strategies that gives rise to equilibrium residual earnings inequality, it also generates a theoretical and computational intractability due to the endogenous earnings distribution acting as an infinite dimensional state variable. While various solutions to this intractability problem exist, such as stationarity in Burdett-Mortensen (1998), Bertrand competition over poached workers in Cahuc, Postel-Vinay and Robin (2006), directed search in Shi (2009) and Menzio and Shi (2010, 2011) or rank-preserving equilibria in Moscarini and Postel-Vinay (2013), they limit the types of specifications that can be considered for unobserved worker- and firm heterogeneity. As a consequence, the estimated quantitative importance of frictional earnings dispersion is likely to be model-dependent. We therefore view a model-free approach to estimation, as used in our paper, as an important complement to these structural approaches.

Intersectoral differences in average wages are a pervasive and persistent empirical regularity in many developed economies, as documented for example for the U.S. in Helwege (1992), for Sweden in Edin and Zetterberg (1992), for France in Goux and Maurin (1999), and for various European countries in du Caju et al (2010). Since these differences survive controlling for worker and job characteristics, they can be interpreted as a particular type of residual wage inequality. Various theories for interindustry wage differentials have been proposed in the literature, including compensating wage differentials, unionization rates, worker sorting on unobserved heterogeneity or rent sharing in frictional environments. Krueger and Summers (1988) and Gibbons and Katz (1992) present strong evidence against explanations based on compensating wage differentials and unionization rates. Gibbons et al. (2005) formulate and estimate a worker-level model of sectoral choices and wages and find that unobserved heterogeneity cannot generate the magnitudes of intersectoral wage differences as observed in the data. They also argue that rent sharing is an important part of interindustry wage differentials, a conclusion that is shared by du Caju et al (2010). While these studies do not explicitly spell out the source of frictions giving rise to differential rent sharing across industries, equilibrium search models with on-the-job-search are a natural conceptual framework for reconciling this evidence with theory. Indeed, as

discussed in Manning (2005) labor market search frictions generate sufficient monopsony power to support a mixed-strategy equilibrium, as long as some employed workers receive outside offers periodically. Yet, to our knowledge there have not been any attempts to test the empirical predictions of multi-sector Burdett-Mortensen models. In Hoffmann and Shi (2016) we study the theoretical properties of a two-sector Burdett-Mortensen model and show that it generates a relationship between the relative rates at which workers move from one sector to another and sector-specific wage structures. In this paper we estimate the reduced form of this relationship, thereby relating changes in within-group differentials of sectoral wage structures to changes in within-group labor market flow rates. This is a fundamentally different empirical strategy for testing theories of rent sharing than those in Krueger and Summers (1988) and Gibbons and Katz (1992). A primary advantage is that it utilizes data on the very source of frictional wage dispersion across sectors, namely job-to-job transitions without an intermittent spell of non-employment. The only work we are aware of that tests multisector equilibrium search using a model-free approach and data on employment and factor prices is Beaudry, Green and Sand (2012). They derive a regression equation that links industry wages to industry-composition on the local level. This type of sectoral linkage in wages is implied by Pissarides-Mortensen (1994) style search models without search on the job and can be tested using repeated cross-sectional data on stocks and factor prices. It exists in our two-sector model as well. However, we explicitly test a model prediction that is central to models with on-the-job search, requiring longitudinal data on worker transitions across sectors.

With job-to-job transitions representing the key mechanism that relates sectoral wage structures, one may wonder how prevalent they are empirically. In a summary article, Davis, Faberman and Haltiwanger (2006) show, using various longitudinal data, that the U.S. labor market is characterized by a remarkable amount of worker and job flows, no matter the business cycle conditions. Fallick and Fleischman (2004) report that 2.6% of employed workers switch their employer from one month to the next, and Nagypal (2008) shows that direct monthly job-to-job transitions represent nearly half of all separations from employers. Furthermore, direct job-to-job transitions display considerable fluctuations over the business cycle, as shown for example in Shimer (2005b) and Nagypal (2008). A number of recent quantitative studies of business cycle models with search frictions, such as Menzio and Shi (2011), Robin (2011) and Moscarini and Postel-Vinay (2016), establish that incorporating endogenous job-to-job transitions into labor market search models help resolving the famous Shimer puzzle (2005a). An important issue however is that measurement of monthly worker flows over a long time horizon is problematic. The longest longitudinal data set in the U.S. that can be used to construct such flow measures is the CPS, but these data have undergone a number of major revisions to the survey design. Nagypal (2008) and Moscarini and Thomsson (2008) provide alternative approaches to

correcting for measurement error in worker flows measured from the monthly CPS data. Since our study crucially relies on measuring intersectoral transitions without an intermittent spell of unemployment, it is affected by the same issues. We therefore heavily borrow from these studies when constructing our data on intersectoral worker transitions.

A continuous process of sectoral reallocation and structural transformation of growing economies has been reported at least since the work by Clark (1940). In particular, as economies grow they tend to reallocate labor from the agricultural sector to the manufacturing sector in a first step of transformation and then to the service sector in a second step. There is now a sizeable exogenous growth literature that attempts to understand how to modify frictionless multisector neoclassical growth models to account for these facts. Common to work in this literature, such as Echevarria (1997), Kongsamut, Rebelo and Xie (2001), Ngai and Pissarides (2007), Herrendorf, Rogerson and Valentinyi (2013), Herrendorf, Herrington and Valentinyi (2015) and Swiecki (2014), is the focus on the evolution of stock variables and aggregate factor prices over time. Equations for the dynamics of factor prices explicitly rely on factor-price equalization across sectors. As a consequence, flows, factor price distributions and the relationship between the two are not studied, which distinguishes our work.

There is very little existing work relating structural transformation and sectoral reallocation to the distribution of labor earnings. An important exception is Buera and Kaboski (2012) and Buera, Kaboski and Rogerson (2015) who show, using time-series and cross-country evidence, that the evolution of returns to education and the sectoral employment share are intrinsically related. They reconcile these facts with a neoclassical model of employment stocks and factor prices in which agents are heterogeneous. Effective wages are equalized across sectors and all of labor reallocation comes through labor market participation rather than cross-sector flows. Another recent and very active strand of the literature on labor reallocation and the wage structure, innovated by Acemoglu and Autor (2011) and Autor and Dorn (2013), focusses on long-run changes in the *occupational* employment structure rather than the rise of the service sector.⁹ This literature highlights the importance of the routine content of occupations for the evolution of the occupational wage structure. In our empirical analysis we always condition on the routine-content of occupations so that we study sectoral reallocation within broadly defined occupational groups. Cortes et al (2015) use longitudinal CPS data to show that a significant share of this reallocation took place either through unemployment or non-employment, in stark contrast to the patterns of sectoral reallocation we document here. An important omission in this part of the literature is an investigation of the link between the evolution of residual inequality and labor market flow rates.

⁹Autor and Dorn (2013) study the role of the low-skill service sector for the well-documented phenomenon of wage- and job polarization. In their study, low-skill services are defined by occupations, not by an industry. Our definition of the service sector is based on industry classifications and is much broader.

Closest to our work is the descriptive analysis provided in Lee and Wolpin (2006). Although their focus is on estimating a dynamic general equilibrium model of sectoral and occupational choices with worker heterogeneity, their exploratory descriptive analysis hints at many of the empirical facts we document here in more detail. Most importantly, they show that relative wages have increased in favor of the service sector since the mid-1980s, even within broad occupational groups, that positive net flows from the goods to the service sector have contributed substantially to service sector employment growth, and that workers are more likely to switch from the goods to the service sector rather than vice versa. We expand substantially this analysis, thereby disaggregating the data much further and using nationally representative longitudinal worker-level data on the monthly frequency rather than the NLSY79. Most importantly, we relate changes in worker flows to changes in the entire wage structure directly in a reduced-form regression framework to test a central prediction of labor market search models.

We also touch on an old literature that attempts to decompose the factors underlying fluctuations in the unemployment rate and labor productivity at the business cycle frequency into a purely sector-specific component and a purely aggregate component. Lilien (1982) finds that sectoral shocks together with frictional reassignment of workers across sectors can explain a large part of cyclical unemployment, while Abraham and Katz (1986) provide evidence against a purely sectoral view of aggregate fluctuations. Chodorow-Reich and Wieland (2015) use more disaggregated data together with an IV-estimation to address this question and find that sectoral reallocation can be an important contributor to the amplification and persistence of business cycles. Building on work by Long and Plosser (1983, 1987), Foerster, Sarte and Watson (2011) show that once one takes into account intersectoral linkages in production, a sizeable share of fluctuations in output is driven by sectoral shocks. Empirical work based on financial markets data, such as Loungani, Rush and Tave (1991) and Brainard and Cutler (1993) reach similar conclusions. Building on theoretical models in Lucas and Prescott (1974) and Rogerson (1987, 2005), recent work by Phelan and Trejos (2000), Pilossoph (2014), Yedid-Levy (2016) and Dvorkin (2014) carry out quantitative structural analyses of multisector business cycle models. A general conclusion from this work is that sectoral shocks together with various forms of labor market frictions can generate sizeable comovement, and, depending on the model, propagation in business cycle fluctuations across sectors.

3. Data - General Description

We investigate the relationship between long-run sectoral transformation and individual labor market outcomes on a disaggregated level. This requires data that cover a long time horizon and that are large in the cross section. The Census data and the American Community Survey (ACS) meet these criteria,

and we rely on them as much as possible. We use variables that are available consistently since 1960. We drop the trial years for the ACS, 2001 to 2004, and the years since the Great Recession in 2008.

Constructing data on flows is more problematic since it inherently needs to rely on panel data. Commonly used panel data such as the PSID or the NLSY are either too small or cover only a limited amount of birth cohorts. Furthermore, they are collected on the annual frequency which generates aggregation biases. We thus use the monthly CPS from 1978 on. This data set is notoriously problematic for constructing flow data because of evidence for large measurement error. Most importantly, a major survey redesign in 1994 that addressed the poor quality of worker-level labor employment transition data had dramatic effects on measures of labor flows between industries and occupations. Similarly, a redesign in 1989 improved the quality of the information on labor force participation. We will describe in detail in a separate subsection how we correct for these measurement issues.

We try to be as consistent as possible when constructing samples from the various data sources. Unless indicated otherwise we include all full-time workers between 18 and 64 years of age who are not in the Armed Forces. To define the service sector we follow Lee and Wolpin (2006) and assume that it consists of the following industries: Transportation and Public Utilities, Trade, Finance, Insurance, Real Estate, Private Household Services, Miscellaneous Services, and Public Administration. For data that are collected retrospectively we define workers as full-time employees if they worked at least 40 weeks in the previous year. All allocated data points are dropped.

Top coding of earnings turns out to be an important issue. In particular, a non-parametric analysis of sector-specific earnings distributions indicated that the mass point at the top code has grown in the service sector relative to the non-service sector, implying that dynamics at the top end of the distribution are an important part of the changing sector-specific wage structures. We therefore use a more systematic way of imputing top-coded earnings than what is the standard in the literature.¹⁰ In essence, we apply a MCMC-type algorithm that imputes top-coded earnings repeatedly from a Pareto distribution until our summary statistics of the earnings distribution converge. We offer a more detailed description of this algorithm in the section on sectoral earnings structures.

¹⁰It is common to multiply top-coded earnings by a factor of 1.5. See for example Beaudry, Green and Sand (2012).

4. Sectoral Employment: Stocks and Flows

4.1. Sectoral Employment and the Role of Compositional Changes

We start our empirical analysis with constructing a measure of *pure sectoral reallocation*, that is, changes in sectoral employment shares that are orthogonal to other observable changes in the labor force composition that took place over a similar time period. We attempt to clean this measure from three labor market trends that are likely correlated with sectoral reallocation: the increase in female labor force participation, the increase in average educational attainment, and the broad changes in the occupational structure commonly referred to as deroutinization or automatization. To this end, we rely on statistical decomposition techniques, which, in our first application, remain fully non-parametric. This exercise is interesting for several reasons. First, it yields a measure of net reallocation within narrowly defined groups that quantifies the extent to which sectoral reallocation can be viewed as a distinct trend from other aggregate changes. Second, it generates a data set of within-group sectoral employment changes that we will match later with a data set on within-group changes in sectoral wage structures. Third, it allows us to identify groups with particularly dramatic labor flows into the service sector.

Let s_t be the aggregate service sector employment share in year t and assume that the population can be characterized accurately by K observables, denoted by G_1, G_2, \dots, G_K , which take on discrete values. We can thus construct $G_1 \times G_2 \times \dots \times G_K$ groups from these observables, which we simply index by g . The service sector employment share in 1960, s_{1960} , can be written as

$$s_{1960} = \sum_g s_{1960}(g) * P_{1960}(g)$$

where $s_{1960}(g)$ and $P_{1960}(g)$ are the service sector share for and the population mass of group g in 1960, respectively. We can then decompose $P_{1960}(g)$ as follows:

$$\begin{aligned} P_{1960}(g) &\equiv P_{1960}(G_1 = g_1, G_2 = g_2, \dots, G_K = g_K) \\ &= P_{1960}(G_2 = g_2, \dots, G_K = g_K | G_1 = g_1) * P_{1960}(G_1 = g_1). \end{aligned}$$

These relationships can be used to construct a counterfactual service sector employment share for 1970 that holds constant everything but the distribution of the first observable characteristic G_1 :

$$s_{1970}^{C(1)} = \sum_g s_{1960}(g) * P_{1960}(G_2 = g_2, \dots, G_K = g_K | G_1 = g_1) * P_{1970}(G_1 = g_1).$$

Here, $C(1)$ stands for the counterfactual that allows only the distribution of the first characteristic to vary over time. The value $(s_{1970}^{C(1)} - s_{1960})$ is the part of sectoral reallocation that cannot be told apart from a purely mechanical effect coming from a changing population distribution of this characteristic.

Using the same logic we can generate the k th counterfactual scenario $C(k)$ for any year $t > 1960$ *cumulatively* using

$$s_t^{C(k)} = \sum_g [s_{1960}(g) * P_{1960}(G_{k+1} = g_{k+1}, \dots, G_K = g_K | G_1 = g_1, \dots, G_k = g_k) * P_t(G_1 = g_1, \dots, G_k = g_k)]$$

This procedure yields the aggregate counterfactual change in the service sector employment share after the K th cumulative counterfactual scenario and is given by

$$s_t^{C(K)} = \sum_g s_{1960}(g) * P_t(g).$$

The part of sectoral reallocation that is systematically unrelated to any changes in the distributions of observable characteristics is therefore

$$s_t - s_t^{C(K)} = \sum_g (s_t(g) - s_{1960}(g)) * P_t(g)$$

and is a weighted average of within-group changes in the service sector employment share, $[s_t(g) - s_{1960}(g)]$. This is the part that we call *pure sectoral reallocation*.

We implement this decomposition using the following four observables as grouping variables: Age, gender, educational attainment, and whether an occupation is primarily associated with routine or non-routine tasks. We consider 5 age groups and 3 educational groups, while gender and task input are dummy variables.¹¹ There are thus 60 groups in total. Our decomposition is entirely non-parametric since we do not impose any parametric structures on either the group-specific service shares or the group-specific population shares. Any counterfactual scenario $C(k)$ with $k < K$ can be used to compute what Firpo, Fortin and Lemieux (2011) call the detailed decomposition, while the K th counterfactual scenario can be used to construct the aggregate decomposition. Generally, any detailed non-parametric decomposition is more informative than the aggregate counterpart, but has the disadvantage that the order of the grouping variables matters. In contrast, the aggregate decomposition is invariant to reordering the counterfactual scenarios even in the non-parametric case.¹²

The main results are shown in Figure 3. It is important to notice that all actual and counterfactual data series condition on employment. The first counterfactual holds group-specific service shares and

¹¹Routine occupations include Sales, Office and Administrative Support, Operators, Fabricators and Laborers. Non-routine occupations include Management, Professional, Technical, Service and Craft. It is standard to divide these further into cognitive and manual occupations. This is problematic in our case since non-routine manual occupations are often defined as Service occupations, which mixes occupational and industry affiliation of the job. However, the analysis in Acemoglu and Autor (2011) indicates that it is the routine vs. non-routine component that relates strongly to employment growth, not whether a job is cognitive or manual.

¹²Order does not matter for the detailed decomposition if one assumes that the relationship between the dependent variable of interest, in our case s_t , and the K grouping variables is linear-in-parameters. In that case, the decomposition is the Oaxaca-Blinder decomposition.

the sample shares of our age-, education-, and occupation- groups at their 1960 values, but adjusts the share of females among the employed to its actual values in each sample year. This counterfactual reflects the increase in female labor force participation since 1960, both in absolute terms and relative to males. Since females are more likely than males to work in the service sector, their increasing representation among the employed generates sectoral reallocation mechanically. In numbers, approximately 3.5 percentage points of the 20 percentage points increase of service sector employment can be explained in this way. The remaining sectoral reallocation of workers has taken place within gender groups. Next, we also adjust the age distribution to its actual outcome. This does not noticeably alter the counterfactual time series of service sector employment, implying that the entire sectoral reallocation of workers must have taken place within at least one of our five age groups. Our third counterfactual adds the mechanical effect of the changing educational composition of the workforce on sectoral employment. This has large explanatory power. Since more educated workers were over-represented in the service sector in 1960, their increasing share of the working population generates a mechanical rise of the service sector employment share. This result is a reflection of what Buera, Kaboski and Rogerson (2015) call *skill-biased structural change*. In total, it can generate an increase in service sector employment of 10 percentage points - about 50 percent of the total sectoral reallocation. Finally, adjusting the employment share of routine occupations can explain another 2 percentage points of the total increase in service sector employment.

As noted above, the particular values coming out of these counterfactuals depend on the order in which they are executed. In figure 1 in the appendix we carry out the same decomposition exercise as in Figure 3, but reverse the order of adjustments of the educational and the occupational composition. This leads to a slightly larger role of the latter, though education remains the observable variable with the strongest explanatory power.

After adjusting the group shares to their actual levels while keeping the service sector share at its 1960 level we are left with an unexplained gap of approximately 5 percentage points, or about one quarter of the entire sectoral reallocation of employment. This is the part of the rise of the service sector that can only be explained by sectoral reallocation within groups, i.e. pure sectoral reallocation. That is, there must be groups for which the service sector employment share has increased substantially since 1960. Pure sectoral reallocation has become stronger since the beginning of the 1990. In fact, nearly all of the sectoral reallocation that has taken place during the 1990's is unexplained. The only observable whose relationship with service sector employment has become stronger since is the non-routine content of occupations. This is to be expected given Autor and Dorn's (2013) finding that the rise of low-skill service occupations was particularly pronounced during that period.

Should one view pure sectoral reallocation as quantitatively important? Since it explains only

approximately one quarter of the total reallocation one may argue that it is not. Indeed, the increase in the educational attainment of the population has twice the explanatory power. However, we think that one quarter is substantial. After all, pure sectoral reallocation as defined above is the part of employment reallocation that is systematically unrelated to *any* observable compositional changes of the workforce. Given that our groups are quite narrowly defined – for example females who are between 27 and 34 years of age, have a high school degree in a non-routine occupation – it is remarkable that we are left with a strong trend in unexplained reallocation of labor to the service sector. This is particularly noteworthy since increases in educational attainment, in female labor force participation, and in the importance of non-routine tasks, may be to some extent reactions to the same force that caused sectoral reallocation. Furthermore, pure sectoral reallocation has become substantially larger since the beginning of the 1990s, which is also the period during which the sectoral earnings structures have diverged particularly strongly, as we will see below.

4.2. Which Groups Have Experienced the Largest Reallocation?

Our analysis in the previous section extracts the part of sectoral reallocation that takes place within narrowly defined groups. We now explore which of these groups reallocated the fastest. This is possible since constructing our measure of pure sectoral reallocation requires calculating within-group changes in the service sector employment share for each of our 60 groups. We therefore treat $\{(s_t(g) - s_{1960}(g))\}_{g=1}^{60}$ as data and regress them on our gender- and routine-task-dummies and a set of fixed effects for our 5 age groups and 3 educational groups.¹³ The omitted group is men between the age of 18 and 26 who do not have a high school degree and work in a non-routine occupation. All regression coefficients need to be interpreted as differences in net reallocation relative to this group.

Results are shown in Table 1. We start with estimating the regressions using within-group changes in service sector employment shares between 1960 and 2007, constructed from our full sample. The estimates are shown in the first column of panel 1. Apart from two age dummies, all coefficients are highly significant and, with the exception of the routine-dummy, positive. That is, between 1970 and 2007, higher educated workers, older workers, and female workers had significantly larger relative employment gains in the service sector than young uneducated men, holding constant their share in the working population. On the other hand, reallocation was slower in routine occupations, a result that likely reflects the shift of low-skill labor into low-skill non-routine service jobs as documented in Autor and Dorn (2013).

The finding that sectoral reallocation has been asymmetric across age groups has not received a

¹³Since we run these regressions to learn about the heterogeneity of the speed of sectoral reallocation across our 60 groups we do not weigh group-level outcomes by their size.

lot of attention. Yet, it has important implications. A popular view is that sectoral reallocation is intrinsically related to the rising labor force participation of women. Indeed, as shown in Table 1, the net flows into the service sector were significantly larger among women than men. At the same time, this view suggests that a lions share of the reallocation has taken place by way of labor market entry rather than through reallocation of labor force participants. Significantly positive age effects are a first piece of evidence against this view because otherwise it should be the youngest workers who reallocate the most.¹⁴ At the same time, if one uses employment changes over a period of almost 40 years, as in column 1 of the table, this argument is problematic because large estimated net flows among the older workers could just as well be driven by a particular strong net flow of labor market entrants into the service sector forty years prior. To address this problem we reestimate our regression models using decadal time changes instead. The results, broken up by decade, are shown in the next four columns.

Interestingly, net flows into the service sector were indeed the largest for the youngest age group during the 70's and for the two youngest age groups during the 80's. However, age groups that have spend a lot of time in the labor market had increasingly larger net flows relative to the youngest in the subsequent two decades. This is also the time period for which our measure of pure sectoral reallocation is particularly large. A conclusion from these estimates is that initially, sectoral reallocation took place among younger workers who were more likely to be labor market entrants. In later decades however, reallocation of labor market participants became more important. Focusing on the role of labor market entry misses this fact.

Given the focus on the relationship between the rise in female labor force participation and the rise of the service sector we find it worthwhile reestimating our regressions of within-group service sector employment shares separately by gender. This is equivalent to allowing all regression parameters on the explanatory variables to vary freely across the two gender groups. Estimates for females and males are shown in panels B and C, respectively. In the extreme case where the rise of the service sector is driven entirely by an increase in female labor market participation, all estimates for the male sample should be insignificant. The service sector employment share of each group of males, defined by age, education and the routine task content of an occupation, would remain constant and there would not be any pure reallocation for this group. Our results show that this is not the case. The estimates on the two education fixed effects and the routine dummy are almost identical for female and male workers. That is, higher educated workers and workers in non-routine jobs were particularly likely to flow into the service sector, no matter their gender.

¹⁴Lee and Wolpin (2006) also document net flows by age groups and find similar results.

4.3. Sectoral Employment: Flows

4.3.1. Constructing Flows Data from the Monthly CPS

How important are direct job-to-job transitions in the reallocation of labor from the goods to the service sector? The answer to this question is crucial for understanding the economic forces that cause reallocation, for formulating appropriate economic models, and for understanding to which extent sector-specific earnings distributions are intrinsically linked. Our evidence from the changes of group-specific service employment shares indicates that sectoral reallocation has not occurred solely by way of directing labor market entrants increasingly to the service sector. Rather, individuals already in the labor force have reallocated as well. This may take place by way of a transition through unemployment or by way of a direct EE transition. Only in the latter case will there be strategic interactions that can, in principle, link the evolution of the entire sector-specific income distributions. Thus, we now need to move beyond repeated cross-sectional data and construct flows from panel data on workers.

The longest-running public-use micro-level panel data set on the monthly frequency is the CPS. Working with these data introduces two major problems. First, it is well known that these data are plagued by measurement error, especially in the years prior to a major redesign of the sampling scheme in 1994. The redesign had the explicit goal to reduce measurement error by introducing dependent interviewing. Second, for each month the data are substantially smaller than a wave of the Census. As a consequence, the time-series on the monthly frequency are noisy.

To address these two problems we proceed as follows. We first use a combination of the data adjustments proposed in Moscarini and Thomsson (2007) and our own correction based on a regression-discontinuity type estimation procedure to create time-series of various transition rates that are cleaned from measurement error. To eliminate noise in the monthly frequency data we then aggregate them up to the annual level. This yields series of average monthly transition rates for each year. We consider all sixteen transitions between service sector employment (S), non-service sector employment (NS), unemployment (U) and non-employment (NILF). We now describe this procedure in more detail.

The CPS interviews a representative sample of households for four consecutive months, and then, after a rest period of eight months, for another four months. In each month, households are added to the sample to keep it nationally representative. The CPS provides a list of variables that allows matching, for each household, the first four with the last four interviews. In practice, the matching rate is significantly below 100 percent for various reasons, such as non-response or a household moving away from a dwelling unit. Nagypal (2008) provides evidence for non-randomness in the sample attrition. We therefore follow the suggestion in Moscarini and Thomsson (2007) and leave the four last interviews unused for *any* household. This does not introduce selection problems as mentioned in Nagypal (2008)

because we drop all of the four last interviews even if they are present in the data.

Significant structural breaks in the transition rate time series are still detectable after this initial sample adjustment. In particular, any transitions involving non-employment change discontinuously in 1989, and intersectoral transitions, which are most likely EE transitions, drop substantially in 1994. We next apply the procedure of dropping “unlikely histories” for correcting occupational transition rates in Moscarini and Thomsson (2007) to our context of intersectoral mobility. Let $\langle e_1, e_2, e_3, e_4 \rangle$ be the tuple of employment states in each of the four interviews, with $e_\tau \in \{S, NS, U, NILF\}$ and $\tau \in \{1, 2, 3, 4\}$. We then compute the empirical frequency in the entire sample of each possible history. Examples of histories are $\langle S, S, U, S \rangle$ or $\langle NS, S, NS, S \rangle$, the former of which is relatively frequent and the latter of which is infrequent. The idea is that histories that are observed particularly infrequently in the post-94 period but relatively frequently in the pre-1994 period are likely to reflect measurement error. We sort the household-level data by the relative frequency of their reported histories and, starting from the least-likely history, drop entire households from the sample, no matter the year, until we detect a significant change in intersectoral transition rates after 1994. This “stopping rule” relies on the assumption that data past 1994 are measured correctly. This is because we interpret changes in post-1994 transition rates from continuing dropping histories as a sign of dropping “too many” histories.¹⁵

The discontinuities in transition rates in 1989 and 1994 become substantially smaller after applying this algorithm, but remain economically and statistically significant. In a final adjustment step we therefore estimate the magnitudes of the structural breaks and use the estimates to generate a set of cleaned time-series. To implement this procedure we follow a Regression-Discontinuity-like identification approach. The idea is that conditional on a flexible time trend, any discontinuous change in measured transition rates that happen exactly at the date of survey redesign, that is between June and July 1988 for transitions to and from NILF and between January and February 1994 for intersectoral transitions, must be solely due to the redesign. We notice that any transition rates to and from NILF remain stable after the break in 1988 and are not affected by the survey redesign in 1994.

We thus start estimating separately for each type of transition to and from NILF an RD-regression, delivering a set of regression estimates for the magnitude in the structural break. Holding constant the departure state, these regression estimates need to add up to zero because the transition rates from that state must add up to one. Consider for example individuals in state S in t . The four transition probabilities from this state are $P(S_{t+1}|S_t)$, $P(NS_{t+1}|S_t)$, $P(U_{t+1}|S_t)$ and $P(NILF_{t+1}|S_t)$. As it turns out, neither $P(NS_{t+1}|S_t)$ nor $P(U_{t+1}|S_t)$ are affected by the structural break. We therefore estimate

¹⁵Since this algorithm also drops some histories that involve transitions between *NILF* and labor force participation, it is likely to address, at least to some extent, classification error in non-employment. As discussed in Nagypal (2008), this classification error generates too many transitions in and out of non-employment. Histories with such transitions that are empirically relatively unlikely will be dropped by our algorithm.

only one structural break from the series $P(NILF_{t+1}|S_t)$, subtract this from the series for any $\tau \leq t$ and add it to $P(S_{t+1}|S_t)$. The procedure is more complicated for transitions from *NILF*, because all four such transitions are affected by the structural break. We thus estimate four coefficients, which do not necessarily add up to zero because of the non-linear time trend we condition on. We thus normalize the adjustment coefficients accordingly. We apply a similar procedure to the structural break in 1994, though this is less involved since none of the transitions into and out of unemployment or *NILF* are affected.

In a final step we aggregate up the transition rates to the annual level to smooth low-frequency fluctuations. This yields a series of average monthly transition rates on the annual level. Intersectoral transition rates to and from the service sector are shown in Figure 4. Monthly transition rates from the service sector to the goods sector have remained nearly constant between 1976 and 2007. In sharp contrast, the rate at which workers move from the goods sector to the service sector has risen dramatically. Apart from a period in the mid-90s, this increase has been quite steady and almost linear. The noticeable deviation from trend in the mid 90s may be viewed as worrisome since it started at around the same time as the survey redesign of the CPS. It is unlikely to be entirely due to measurement error however since then one would expect the transition rate in the opposite direction to display a similar discontinuity. It is more likely that this pattern is partly due to the tech-boom that took place around that time and that came to an abrupt halt at the end of the 90s. This is also consistent with findings in Beaudry, Green and Sand (2016) regarding the importance of the tech boom-bust cycle for explaining occupational skill upgrading. In any case, our results below are not driven by this deviation.

4.3.2. A Simple Descriptive Model of the Evolution of the Service-Sector Employment Share

To quantify the role of job-to-job transitions in reallocating labor from the non-service to the service sector, we simulate forward a system of flow equations using our empirical transition rates and perform counterfactual experiments.¹⁶ Let $\Omega = \{S, NS, U, NILF\}$ be the state-space of the dynamic system and let the 4×1 column vector $\Sigma_{t,\tau}$ be the unconditional distribution over the states at the beginning of year t and month τ . It will become clear in a moment why we need to index time by (t, τ) rather than a unified time index. Also define Λ_t to be the Markov transition matrix acting on $\Sigma_{t,\tau}$. If $e_{t,\tau} \in \Omega$ is the state in period (t, τ) then a generic element of Λ_t is the transition probability from state i to state j given by $\lambda_t^{ij} = P(e_{t,\tau+1} = j | e_{t,\tau} = i)$ for $1 \leq \tau < 12$ and $\lambda_t^{ij} = P(e_{t+1,1} = j | e_{t,12} = i)$ if $\tau = 12$. The estimate of Λ_t , which we write as $\hat{\Lambda}_t$, is composed of the empirical transition rates as computed in

¹⁶See for example Cortes, Jaimovich, Nekarda, Siu (2015).

the previous section. These are average monthly transition rates for year t and thus do not vary across months of that year. This explains why τ does not enter the time index of Λ .

Given the initial state $\Sigma_{1976,1}$, which can be directly computed from the data, the dynamics of $\Sigma_{t,\tau}$ are then described by

$$\begin{aligned}\widehat{\Sigma}_{t,\tau+1} &= \widehat{\Lambda}_t * \widehat{\Sigma}_{t,\tau} \text{ for } 1 \leq \tau < 12 \\ \widehat{\Sigma}_{t+1,1} &= \widehat{\Lambda}_t * \widehat{\Sigma}_{t,12}.\end{aligned}$$

Here we use the hats to highlight that the system generates estimates of the distribution over states that can then be compared with the actual data. These equations demonstrate that we simulate the dynamics of the state vector within year, that is on the monthly frequency, using the average monthly transition rate for that year. Only when the system progresses to the next year do we update the transition matrix. In this way we abstract from high-frequency fluctuations in the state due to seasonality, business cycle shocks, measurement error or statistical noise while utilizing the monthly frequency of the data.

A potential issue with this simple descriptive model is that we may under-estimate flows in and out of the labor force because of the age cutoffs at 18 and 64. If for example 17 years old non-participants enter the market at age 18, and if these inflows are increasingly directed towards the service sector, then we may understate the role of labor market entry in our counterfactual experiments below. We therefore replace the 4th element in $\widehat{\Sigma}_{t+1,1}$, which is given by the beginning-of-the-year non-employment rate $\widehat{P}(e_{t+1,1} = NILF)$, with its actual value in the data, $P(e_{t+1,1} = NILF)$. But since the elements in $\widehat{\Sigma}_{t+1,1}$ need to add up to one, this requires rescaling of the other 3 elements. We scale them so that the probabilities conditional on labor market participation are unaltered when replacing $\widehat{P}(e_{t+1,1} = NILF)$ by $P(e_{t+1,1} = NILF)$. The modified dynamic system is given by

$$\begin{aligned}\widehat{\Sigma}_{t,\tau+1} &= \widehat{\Lambda}_t * \widehat{\Sigma}_{t,\tau} \text{ for } 1 \leq \tau < 12 \\ \widehat{\Sigma}_{t+1,1} &= \widehat{\Lambda}_t * \widehat{\Sigma}_{t,12} \\ \widehat{\Sigma}_{t+1,1} &= \begin{pmatrix} \widehat{P}(e_{t+1,1} = S) \\ \widehat{P}(e_{t+1,1} = NS) \\ \widehat{P}(e_{t+1,1} = U) \\ P(e_{t+1,1} = NILF) \end{pmatrix} \cdot \begin{pmatrix} \frac{1 - P(e_{t+1,1} = NILF)}{1 - \widehat{P}(e_{t+1,1} = NILF)} \end{pmatrix}.\end{aligned}\tag{4.1}$$

4.3.3. Model Fit

The dynamics of the distribution over stocks in (4.1) are simulated using actual data on transition rates. One may therefore expect that the model fits the empirical evolution of the four population shares perfectly. This is not so if, to maximize sample size and to mimic the Census closely, one explores the model fit relative to empirical shares in the March CPS rather than in the January files of

the monthly CPS. Cortes et al. (2015) provide an extensive discussion of this point. In our context one important source of discrepancies is the fact that transition rates are computed from panel data which condition on the same individual being observed in two subsequent months, while the shares in the March CPS are computed from a single cross-section. As a consequence, if individuals who enter (exit) our estimation sample at age 18 (64) are drawn to (withdrawn from) one employment state at a different rate than the rest of the working-age population, then we will not match the data perfectly. Another important source is measurement error in transition rates that is not eliminated by our procedure.

We plot the evolution of the service sector employment share simulated from our measurement corrected transition rates together with their empirical counterpart computed from the March CPS in panel (a) of Figure 5. To keep the figure comparable with Figure 1 we show the service sector employment share relative to the population of employed individuals. As can be seen from the figure, the simulated values of this share are too low before 1994, but fitted almost perfectly afterwards. The largest difference can be observed for the first ten years, while the two time-series track each other quite closely afterwards. It is natural to ask to which extent the mismatch prior to 1994 is generated by our final adjustment step, which we label RD-adjustment.¹⁷ The corresponding figure when based on transitions that are computed without this adjustment is shown in panel (b). In this case we find that the model fits the data much better until 1994 and diverges somewhat afterwards, the latter of which is not surprising since the RD-adjustment does not apply to this period. The growth of the service sector in the post-1994 period is slightly too large. A possible reason is that measurement error in transition rates has not been completely eliminated by the survey redesign in 1994. Another reason may be that those turning 18 (65) are entering (exiting) the service sector at a slower (higher) rate than the prime-age population. Taken together, it is clear that our mismatch is largely due to the RD-adjustment, which generates too little sectoral reallocation in the pre-1985 period. As a test for robustness, we will carry out our counterfactual exercises below using both RD-adjusted and unadjusted transition rates. It is important to highlight that the difference between the two time-series of transition rates has a major impact on both quantitative and qualitative results from our counterfactual analysis only for the transitions from unemployment and non-employment, which reverse their roles. Our qualitative conclusions regarding the role of job-to-job transitions for the rise of the service sector are entirely robust to the transition rates used in the counterfactual analysis or whether we restrict our sample period to the post-1994 period.

¹⁷RD” stands for ”regression discontinuity” since this is, in essence, the method we apply when quantifying the various structural breaks.

4.3.4. The Role of Job-to-Job Transitions in Sectoral Reallocation

We now quantify the importance of direct intersectoral worker transitions for the rise of the service sector. To this end we conduct counterfactual experiments on (4.1). Our main question is by how much slower the reallocation had been if the transition rates to and from the service sector had remained at their 1976 values. We consider three counterfactual scenarios, where we use the subscript “ t, c ” to denote the counterfactual transition rate in period t :

1. **Time-constant intersectoral transition rates:** For all $t > 1976, i \in \{S, NS\}$ and $i \neq j$:

$$\lambda_{t,c}^{ij} = \lambda_{1976}^{ij}; \quad \lambda_{t,c}^{ii} = \lambda_t^{ii} + \left(\lambda_{t,c}^{ij} - \lambda_{1976}^{ij} \right).$$
2. **Time-constant flows from unemployment:** For all $t > 1976, i \in \{S, NS\}$: $\lambda_{t,c}^{Uj} = \lambda_{1976}^{Uj}$;

$$\lambda_{t,c}^{U,U} = \lambda_t^{U,U} + \left(\lambda_{t,c}^{U,S} + \lambda_{t,c}^{U,NS} - \lambda_{1976}^{U,S} - \lambda_{1976}^{U,NS} \right).$$
3. **Time-constant flows from NILF:** For all $t > 1976, i \in \{S, NS\}$: $\lambda_{t,c}^{NILF,j} = \lambda_{1976}^{NILF,j}$;

$$\lambda_{t,c}^{NILF,NILF} = \lambda_t^{NILF,NILF} + \left(\lambda_{t,c}^{NILF,S} + \lambda_{t,c}^{NILF,NS} - \lambda_{1976}^{NILF,S} - \lambda_{1976}^{NILF,NS} \right).$$

The first counterfactual isolates the role of intersectoral transitions for sectoral reallocation. It holds constant the intersectoral transition rates and allocates the difference to the actual transition rate to the probability of remaining in a sector from one to the next month. The latter step guarantees that transition probabilities from each of the two employment states add up to one. The second and third counterfactuals quantify the importance of flows from unemployment and from NILF respectively in a similar way. We continue replacing, after each period, the simulated with the actual non-employment rate to avoid potentially understating the importance of flows from NILF, as described in the previous section.

The results are shown in Figure 6 and Table 2. As before, we perform simulations using RD-adjusted and unadjusted transition rates to check the robustness of our results. Furthermore, the employment shares are relative to employed workers rather than the entire population. As shown at the bottom of the table the service sector employment share grew by 9.42 percentage points between 1976 and 2007, compared with model-generated increases of 10.45 percentage points if we use RD-adjusted rates and of 10.63 percentage points if not. If intersectoral transition rates had stayed constant at their 1976 RD-adjusted values, which is our first counterfactual scenario, the predicted increase of the service sector employment share relative to all employed prime-aged workers would have been 6.10 percentage points, that is approximately 3.3 percentage points lower. This is about a third below the actual reallocation of labor. If we use non-RD-adjusted rates instead, the counterfactual increase in service sector employment share is just 2.6 percentage points, that is a quarter of the model-fitted reallocation

of 10.63 percentage points. In either case, direct transitions between sectors can explain a significant part of the rise of the service sector.

Results from simulating the second counterfactual scenario are shown in the next column. In this scenario we hold the transition rates out of unemployment into the two sectors at their 1976 values. This, too, can generate a significant drop in the size of the service sector. If we use RD-adjusted rates, the population share of service sector employment decreases by 3.42 percentage points, and once we use RD-unadjusted rates this value decreases to approximately 2.5 percentage points.

Our third counterfactual scenario quantifies the importance of transitions from NILF into one of the two sectors. Interestingly, this is found to be somewhat less important, with generating a drop in sectoral reallocation by less than a percentage point with RD-adjusted transition rates and almost three percentage points with RD-unadjusted rates. One should notice that once one uses RD-unadjusted rates, transitions out of NILF and out of unemployment reverse, to some extent, their roles in generating a rise in service sector employment. This is most likely a reflection of the fact that prior to 1989 the CPS was problematic for distinguishing these two types of transitions. Our adjustment procedure is evidently sensitive to this issue. However, our finding of the central role of direct intersectoral labor flows for the rise of the service sector is robust to our adjustment procedure.

The entire evolution over time of service sector employment under each of the counterfactual scenarios is shown in Figure 6. The top panel is based on RD-adjusted rates and the bottom panel is based on RD-unadjusted rates. This uncovers two important results. First, the counterfactual experiments have relatively small effects before the mid-80s. It is only afterwards that the counterfactual series diverge significantly, and in some cases dramatically, from the model-predicted (and actual) evolution. We conjecture that this is partly due to the importance of increasing female labor force participation during the 70's on sectoral reallocation, as documented in an earlier section on the evolution of stocks. In particular, because of our age cutoffs we may underestimate the effect of labor market entry that is directed towards the service sector. Second, the most robust and largest counterfactual effect stems from keeping intersectoral transition rates at the 1976 value. This effect is particularly strong after 1985. This shows that the importance of direct job-to-job transitions for the rate of sectoral reallocation has been unjustifiably neglected in the literature. If one wants to understand sectoral reallocation one needs to understand why workers have become more likely to switch from the non-service to the service sector rather than in the opposite direction. Theories that emphasize the role of labor market entry miss this point.

5. Sectoral Income Structures

5.1. Statistical Decomposition of the Income Structure

We have shown that a significant part of the increase in the service sector employment share cannot be explained by changes in the composition of the labor force. Rather, sectoral reallocation has taken place within narrowly defined groups of the prime-age population, and a substantial part of this within-group reallocation has taken place by way of direct job-to-job transitions. As shown in Figure 2, sectoral reallocation came with pronounced changes in the sector-specific earnings structures. For example, both the differences in the averages and the variances of log-earnings between the two sectors have evolved in a convex shape: A decline in the 70's and a large increase afterwards. We now ask if these patterns can be explained by a change in the composition of worker characteristics across sectors. If we control for sufficiently many worker characteristics, then a neoclassical model with two sectors predicts that the answer to this question is "Yes". For any group of workers that can be viewed as sufficiently similar with respect to their labor market skills and tastes, competition should eliminate any earnings differences. Since we decompose changes rather than levels in earnings structures, this argument applies even in the presence of time-constant compensating wage differentials. As in previous sections, we rely on statistical decomposition methods to address this question. With two continuous random variables to be decomposed - earnings in each sector - the decomposition techniques are substantially more involved. We therefore start with an overview of our methodology.¹⁸

Let $\phi(\cdot)$ be any functional of the earnings distributions $f(w)$, such as moments or quantiles. Because we allow earnings distributions to be different across sectors and time periods, it is convenient to use the notation ϕ_t^j , $j \in \{S, NS\}$ to denote the value of the functional when evaluated at $f_t^j(w)$. We are interested in decomposing the intersectoral differences and its evolution over time, $\Delta_t = \phi_t^S - \phi_t^{NS}$, which can be rewritten as follows:

$$\begin{aligned}\Delta_t &= (\phi_t^S - \phi_t^{Sc}) + \phi_t^{Sc} - (\phi_t^{NS} - \phi_t^{NSc}) - \phi_t^{NSc} \\ &= [(\phi_t^S - \phi_t^{Sc}) - (\phi_t^{NS} - \phi_t^{NSc})] + \{\phi_t^{Sc} - \phi_t^{NSc}\} \\ &= \Delta_t^{Unexp} + \Delta_t^{Exp}.\end{aligned}\tag{5.1}$$

where ϕ^{jc} , $j \in \{S, NS\}$ is the functional of the sector-specific counterfactual wage distribution, that is, the distribution of earnings which would have prevailed if workers had been paid according to the base year schedule but their characteristics had been distributed as in the year of interest. This type of counterfactual isolates the composition effect due to a change in the distribution of observable worker characteristics, denoted by Δ_t^{Exp} in equation (5.1), from the effects of sources that alter within-group inequality, denoted Δ^{Unexp} . As discussed in Firpo, Fortin and Lemieux (2011), there are two broad

¹⁸For an excellent summary article of decomposition methods, see Firpo, Fortin and Lemieux (2011).

types of sources for the latter. The first are structural effects reflecting changes in preferences or technologies. The second are composition effects due to a reallocation of workers based on unobserved characteristics. Without explicit parametric assumptions it is not possible to tell apart these two channels underlying Δ_t^{Uexp} . In this section we quantify the composition effect Δ_t^{Exp} . In the final section of this paper we will argue that the changes in within-group inequality are not entirely caused by the sorting of workers.

The decomposition (5.1) requires estimation of functionals of the counterfactual wage distribution. This can be done using the reweighting method proposed by DiNardo et al. (1996). To this end we start with expressing the distribution of earnings in 1970 as

$$f_{1970}(w) = \int_{x \in \Omega_x} f_{1970}(w, x) dw dx = \int_{x \in \Omega_x} f_{1970}(w|x) g(x|t_x = 1970) dx, \quad (5.2)$$

where w , $f(w)$, x and $g(x)$ denote wages, wage density, individual attributes and their population distribution respectively. The counterfactual wage density is then given by¹⁹

$$f_{1970}^C(w) = \int_{x \in \Omega_x} f_{1970}^C(w, x) dw dx = \int_{x \in \Omega_x} f_{1970}(w|x) g(x|t_x = t) dx, \quad (5.3)$$

where $t > 1970$. DiNardo et al. (1996) observed that the counterfactual distribution (5.3) differs from its actual counterpart only by the term $g(x|t)$, i.e. the distribution of individual attributes conditional on a sample year. Defining $\psi_x = g(x|t_x = t)/g(x|t_x = 1970)$, this equation can be rewritten as

$$\int_{x \in \Omega_x} f_{1970}^C(w, x) dw dx = \int_{x \in \Omega_x} f_{1970}(w|x) \psi_x g(x|t_x = 1970) dx. \quad (5.4)$$

This translates the exercise of estimating the non-parametric counterfactual wage distribution into estimating the weighting function ψ_x . Since $g(x|t_x = t)$ is a multivariate density function, this can be computationally intensive and fraught with imprecision, for example because some values of x in the numerator or denominator can have probability mass zero in finite samples. To overcome this implementation problem, DiNardo et al. (1996) propose to apply the Bayes' rule:

$$Pr(x|t_x = t) = \frac{Pr(t_x = t|x) dG(x)}{\int_x Pr(t_x = t|x) dG(x)} \quad (5.5)$$

The weighting function can thus be computed as

$$\psi_x = \frac{Pr(t_x = t|x) Pr(t_x = 1970)}{Pr(t_x = 1970|x) Pr(t_x = t)} \quad (5.6)$$

¹⁹Notice that to be able to write the counterfactual distribution in this way, we need to assume that the wage structure in 1970 does not depend on the vector of individual attributes. This assumption is rather innocuous for the purpose of our study since we are not interested in the *causal* effect of x on the wage distribution.

Notice that the term $Pr(t_x = t)$ is just a fraction of year t observations in the pooled sample of 1970 and year t .²⁰ Moreover, the term $Pr(t_x = t|x)$ can be estimated with a binary response model on a pooled sample where the treatment variable represents “being observed in year t ”. DiNardo et al. (1996) showed that using the sample estimates of the propensity score weighting function, one can compute various statistics from the counterfactual wage distribution. For instance, the counterfactual mean of wages can be computed as follows:

$$\bar{W}_t^c = \sum_i \hat{\psi}_{Xi} \omega_i W_{it}, \quad (5.7)$$

where ω_i is the sample weight. The counterfactual variance of wages can be computed in a similar fashion:

$$Var(W_t^c) = \sum_i \hat{\psi}_{Xi} \omega_i (W_{it} - \bar{W}_t)^2, \quad (5.8)$$

Similarly, one can estimate any counterfactual statistic of the wage distribution by using the composite weight $\hat{\psi}_{Xi} \omega_i$. Moreover, under the assumption that unobserved characteristics are uncorrelated with x , we can investigate how unexplained and explained mean sectoral wage differentials are related to worker characteristics.

Equations (5.7) and (5.8) permit fully non-parametric decompositions of the first two moments of the earnings distribution. Decompositions of other functionals of the earnings distribution require additional assumptions. Suppose that earnings are linearly related to individual attributes, possibly by way of group fixed effects. Using the law of iterated expectations we can estimate the mean of the counterfactual wage distribution in year t as $\bar{W}_t^c = \bar{X}_t \hat{\beta}^c$, where

$$\hat{\beta}^c = \left(\sum_i \hat{\psi}_{Xi} \omega_i X_{it} X'_{it} \right)^{-1} \sum_i \hat{\psi}_{Xi} \omega_i W_{it} X_{it}, \quad (5.9)$$

This methodology cannot be extended to quantiles since the quantile function is not a linear operator and the law of iterated expectations does not apply. Firpo, Fortin and Lemieux (2009) address this problem using the Recentered Influence Function (hereafter RIF) regression decomposition method. This requires computing the recentered influence function of the unconditional quantile of wages, defined as follows:

$$RIF(W, Q_\tau) = Q_\tau + \frac{\tau - \mathbf{1}(w \leq Q_\tau)}{f_W(Q_\tau)}, \quad (5.10)$$

where Q_τ is the τ th quantile of the wage distribution and $\mathbf{1}(\cdot)$ is an indicator function. Notice that to compute the recentered influence function of quantiles of the counterfactual wage distribution we

²⁰This term does not vary across individuals, so it can be ignored in practice.

just need to replace Q_τ in equation (5.10) with Q_τ^c , where Q_τ^c is computed using propensity score reweighting. The central property of the RIF function is that it aggregates back to the statistics of interest. In our case, $\int RIF(W, Q_\tau) f(w) dw = Q_\tau(F_W)$. In its simplest form, the conditional expectation of the quantile RIF can be expressed as a linear function of individual attributes:²¹

$$\mathbb{E}[RIF(W, Q_\tau | X)] = X \gamma + \epsilon \quad (5.11)$$

Therefore, similarly to the case of mean wages, we can estimate the predicted unconditional quantile of the counterfactual wage distribution as $\bar{R}IF(w_t, Q_\tau) = \bar{X}_t \hat{\gamma}^c$ in a first step. The predicted actual and counterfactual RIFs can be used to decompose the sectoral wage differential into the explained and the unexplained parts in a second step as in equation (5.1).

5.2. Implementation and Results

To implement the decomposition techniques for intersectoral difference in means, variances and various quantiles we focus on the distribution of labor earnings rather than hourly wages in the Census and the ACS.²² The primary reason is that we analyze changes in the earnings structure within narrowly defined groups so that we need large sample sizes. The hours-of-work variable in the Census is of significantly lower quality than the information on labor market earnings and introduces additional noise. We therefore avoid using it in the main body of our analysis.²³ For the earnings variable we impose the same sample restrictions as in the rest of the analysis. Importantly, we exclude any imputed observations. Similarly to the service share decomposition analysis described in Section 3, our explanatory variables are represented by age, education, gender and a dummy indicator for routine occupations. To avoid parametric restrictions we estimate earnings regressions flexibly, using the third degree polynomial in age fully interacted with three education dummies, gender dummy, and the routine occupation dummy.

As discussed above, the issue of top-coding of earnings needs to be addressed since it affected the sector-specific earnings distributions differently. We use a simulation-based imputation procedure. First, for each group defined by sector, age, gender, education and the routine dummy we approximate the right tail of the wage density using a two-parameter Pareto distribution.²⁴ Second, we replace

²¹Similarly to mean wages, we estimate the RIF function using the third order polynomial in age fully interacted with three education dummies, gender dummy, and the routine occupation dummy.

²²We have also carried out the decompositions using the march CPS data. The results, though noisier, are consistent with the findings we document in this section.

²³As it turns out, we reach the same conclusions when using hourly wages rather than earnings, although the results are somewhat noisier.

²⁴Since the distribution has only two parameters, it is possible to estimate it using data from only two points in the observed distribution of wages. Following the literature we choose the 90th percentile of the observed distribution as the first point, and the topcode value as the second point.

top-coded values by random draws from these group-specific distributions and compute our statistics of interest. Third, we iterate on this simulation procedure until the cumulative average of these statistics converge.

Results are displayed graphically in Figures 7a to 7c. These figures show the actual and the counterfactual evolution of the intersectoral difference of means, variances and the 90th percentile of the log-earnings distribution. As explained above, the counterfactual holds the group-specific income structure constant at its 1970 value, but adjusts the distribution of observed group-characteristics to its actual level. The difference between the actual and the counterfactual series can stem from either a change in the sector-specific within-group earnings structure or a change in the intersectoral composition of unobserved worker skills. In the case of the mean wage differential, there is no substantial difference between 1960 and 1980. However, the stark increase in the intersectoral earnings differential in favor of the service sector starting in 1980 cannot be explained by a change of the composition of observed worker characteristics. If anything, the counterfactual scenario generates a slight decrease in this differential. Since we control for gender and the routine-content of occupations these results suggest that neither the increase in female labor force participation nor the deroutinization of occupations can explain the rise of earnings in the service sector relative to the non-service sector.

Similar conclusions can be reached when decomposing intersectoral differentials in variances and percentiles of the earnings distributions. The counterfactual series for the former is almost flat, showing that the increase in earnings inequality in the service sector relative to the non-service sector has taken place within narrowly defined groups. The decomposition of a measure of upper-tail inequality as reflected by intersectoral differentials in the 90th percentile of the earnings distribution is particularly interesting. This differential has risen dramatically in the post-1980 period, in accordance with the popular view that high-paying jobs in finance and the legal and medical professions have seen the most dramatic rises in earnings. As shown by our analysis, this evolution has taken place within groups indeed, raising the question of why the non-service sector has fallen behind in competing for similar workers.

6. Relationship between Flows and the Wage Structure

Is sectoral reallocation associated with a change in the sectoral earnings structure? It is well known that intersectoral earnings differences are pervasive and robust. It is less known how such differences correspond with intersectoral worker flows. Neoclassical multi-sector models predict that they are unrelated as long as observable worker characteristics are good proxies for worker skills. In this section we address this question using two empirical strategies. First, we use repeated cross-sectional data

from the Census and the ACS to investigate if changes in earnings differentials are systematically related to a rise in net flows into the service sector. This brings together our group-level analysis of employment stocks and the earnings structures in previous sections. Second, we use CPS panel data on workers to relate more directly intersectoral employment flows to the earnings structure. To the best of our knowledge, this is the first attempt at relating measures of job-to-job transitions to the earnings structure. Finally, we use some further evidence to discuss whether our findings can be reconciled with a neoclassical model with heterogeneous agents, or if worker flows themselves are likely to be an important factor behind the decoupled dynamics of intersectoral earnings structures as observed in the data.

6.1. Evidence from Repeated Cross Sections

Our first empirical strategy consists of estimating a repeated cross-sectional regression of group-level intersectoral log-earnings differentials on the group-level service sector employment share. As in the decomposition exercises we define groups by age, gender, educational attainment, and whether an occupation is primarily associated with routine or non-routine tasks. The dependent variable is constructed by calculating the difference in average log-earnings between the service and the non-service sector for each group and each year. The explanatory variable is the corresponding service sector employment share. We call these variables w_{gt}^d and s_{gt} respectively, where the superscript d stands for "intersectoral differential". Our preferred regression specification is given by

$$w_{gt}^d = \beta * s_{gt} + \mu_g + \varepsilon_{gt}$$

where μ_g is a group fixed effect. The coefficient of interest is therefore identified from first differences and measures the statistical association between within-group changes in intersectoral earnings differentials and within-group changes in the service sector employment share. The latter is given by $(s_{gt} - s_{gt-1})$ and is a measure of group-level net flows into the service sector between periods $t - 1$ and t that can be computed from repeated cross sectional data. The other four outcomes we use are group-level intersectoral differences in the standard deviation and the 10th, 50th and 90th percentile of the earnings distribution. For each outcome, we also show results from a specification that adds time fixed effects to the regression. In this case, the regression coefficient can be interpreted as a difference-in-difference, measuring the effect of sectoral reallocation relative to an aggregate trend on the earnings distribution. For each specification and each outcomes we weigh observations by group size and cluster standard errors on the group level.

Results are shown in Table 3. Each coefficient in the table is an estimate from a different regression. The first row corresponds to estimates that are obtained when group-level variables are constructed

from the entire sample. With the exception of intersectoral differentials in the 10th percentile, all coefficients are positive. They are significant for both specifications when using intersectoral differentials in standard deviations, medians and the 90th percentiles as outcomes. For example, groups whose service sector employment grew by 10 percentage points experienced, on average, an increase of the standard deviation of log-earnings of .024 in the service sector relative to the non-service sector. This is a large effect given that the sample average of intersectoral differences in standard deviations is .077 (shown in the second to last row of the table). Likewise, they saw the 90th percentile grow by .38 log points, reducing the negative intersectoral difference of $-.079$ by almost fifty percent. These results are robust to the inclusion of time fixed effects.

The descriptive analysis of net flows in Table 1 and of intersectoral transition rates in Figures 5 and 6 indicates that an increase in labor force participation rates of females rather than an increase of direct job-to-job transitions into the service sector was the dominant force behind the rise of the service sector in the 1970s. The opposite is true in subsequent decades. In the next two rows we thus reestimate our regressions splitting the sample into these two episodes. The evidence is quite clear: Most of our estimates are insignificant in the earlier period, whereas they are highly significant in the post-1980 period, with the exception of, again, intersectoral differentials in lower-tail inequality. Interestingly, it is only this part of the distribution that was affected by sectoral reallocation in the earlier period, possibly reflecting the increasing labor force participation rates of females together with gender differentials in earnings.

In the fourth row we exclude the youngest age group, which includes individuals between the ages of 18 and 26. It is this group for which labor force participation is likely to be particularly important for explaining the rise of the service sector. Furthermore, younger individuals are less likely than the other age groups to work less than full-time so that the hours-of-work margin may be dominating the observed changes in the earnings structure. Once excluded from the sample, our results become particularly strong, suggesting that changes in labor force participation rates are weakening rather than driving our estimates.

We also experimented with functional forms, using intersectoral ratios rather than differentials in various measures of earnings inequality as outcome variables. Results are shown in Table 1 in the appendix. Generally, these results are fully consistent with those in Table 3.

We conclude that there is a strong and robust association between within-group sectoral reallocation and within-group intersectoral differentials in the earnings distribution. This result is unlikely to be driven by changes in labor force participation rates, especially in the later years of our sample. The result that the relationship between sectoral reallocation and intersectoral differentials is weakest in the

lower end of the earnings distributions is not surprising if one interprets the evidence in terms of a model with a job ladder. It is at the low end of the earnings distribution where the supply effect, according to which workers are willing to accept earnings cuts to get into the growing sector, is dominant, as we discuss in the next section.

6.2. Evidence from Data on Flows

6.2.1. Theoretical Background

To derive explicit hypotheses about the relationship between changes in intersectoral worker transition rates and the sectoral earnings structures and to interpret our results below, it is helpful at this point to formulate a simple two-sector search model with on-the-job search. This also helps demonstrating by which channel an accelerated rate of worker flows from the goods to the service sector may be related to the rise in relative average earnings and earnings inequality in the service sector.

Let $f^i(w, y)$ be the earnings offer distribution in sector $i \in \{G, S\}$ if per-worker revenue in the service sector (S) relative to the goods sector (G) is y . Assume that time is continuous and that a continuous mass of workers who are employed at earnings w receive offers from sector i at rate λ^i and from the other sector at rate σ^{-i} . Also assume that workers and firms are homogeneous, so that the only source of heterogeneity is the intersectoral differential in revenue per worker, y . Taking into account exogenous job breakups that result in a spell of unemployment at rate δ and a discount rate of ρ , the value function of such a job is

$$\begin{aligned} \rho * V^i(w, y) = & w + \lambda^i * \int_{\bar{w}^{-i}}^{\bar{w}^i} [V^i(w', y) - V^i(w, y)] f^i(w', y) dw' \\ & + \sigma^{-i} * \int_{w_R^i(w)}^{\bar{w}^{-i}} [V^{-i}(w', y) - V^i(w, y)] f^{-i}(w', y) dw' \\ & + \delta * [V^u(y) - V^i(w, y)]. \end{aligned} \quad (6.1)$$

Here, $w_R^i(w)$ is the reservation wage that makes a worker indifferent between staying in sector i and earning w per unit of time and moving to sector $-i$ and earning $w_R^i(w)$.

Now suppose that y increases, consistent with the finding in Buera and Kaboski (2012) that value added per worker has increased in the service sector relative to the goods sector. In a model with labor market frictions it is plausible to assume that such a change in the aggregate environment will cause the dispersion of $f^S(w, y)$ to increase relative to the dispersion of $f^{NS}(w, y)$. After all, holding earnings constant, service sector firms now make more profits per worker. Market forces will then increase the mass of firms that post high earnings, which also tends to increase \bar{w}^S relative to \bar{w}^{NS} . Under the empirically realistic assumption that $\lambda^i > \sigma^{-i}$ this also implies that workers will be willing to take earnings cuts to move to the service sector as quickly as possible because of its improved career

prospects. This will have the effect of lowering reservation earnings of the unemployed for service sector employment on the one hand and of increasing the rate at which workers move from the goods to the service sector on the other hand.

This discussion treats $f^i(w, y)$ as exogenous objects. A central question in search theory is whether and how non-degenerate earnings distributions can exist in equilibrium absent any source of heterogeneity. In Hoffmann and Shi (2016) we endogenize $f^i(w, y)$ in a two-sector Burdett-Mortensen equilibrium search model with wage posting.²⁵ We show that residual earnings inequality in each sector is a feature of equilibrium, but characterization is quite involved. Loosely speaking, this is a model of rent sharing in which firms have partial monopsony power because of search frictions. They then need to decide which wages to post. High-wage-posting strategies attract workers at high rates and lose workers at low rates, but per-worker profit is small. The opposite is the case for low-wage-posting strategies. In equilibrium, firms play mixed-strategies over a continuous and bounded support. This theory can rationalize (i) gross worker flows in excess of net worker flows because each sector has some high-wage and some low-wage firms; (ii) a higher rate of job-to-job transitions from the goods- to the service-sector than in the opposite direction; (iii) permanent and substantial differentials in sectoral earnings structures, with higher average wages and higher dispersion in the service sector; (iv) an increase of intersectoral differentials in worker transitions and earnings structures in the process of sectoral reallocation. Importantly, the economic mechanism behind these implications, namely the combination of strategic interactions, search frictions and monopsony power, is absent from *any* frictionless neoclassical two-sector model of the economy because intersectoral transition rates are not defined. Rather, it is the change in the service-sector employment share, a stock variable, that summarizes the reallocation of workers and firms across sectors, even in the presence of unobserved heterogeneity.

6.2.2. Empirical Results

The central prediction coming out of equilibrium search models with on-the-job search is that the rate of job-to-job transitions is intrinsically linked to the earnings structure. This is also the empirical prediction that separates models with on-the-job-search from models in which job transitions must take place via a spell of unemployment, such as in Pissarides and Mortensen (1994) or Shimer (2005). We thus now move on to an empirical design that relies on direct measures of intersectoral transition rates

²⁵ Alternatively one can think of a model in which an increase in y , i.e. a revenue enhancing effect that is common to all firms in the service sector, causes a more heterogeneous set of firms to enter the service sector than the goods sector, causing an increase in the dispersion of $f^S(w, y)$ relative to the dispersion of $f^{NS}(w, y)$. See e.g. Meghir, Narita and Robin (2015) for an application to the study of formal and informal employment. In either case there will be a reduced-form relationship between changes in net flows into the service sector via direct job-to-job transitions and changes in the relative earnings structure between the two sectors.

rather than gross or net flows. This poses a problem because the CPS does not provide repeated observations on individual earnings on the monthly frequency. Yet, to keep as long and as representative a sample as possible we want to continue using these data.²⁶ We therefore construct transition rates on an appropriate level of disaggregation from the CPS and match them to measures of intersectoral differences in the earnings structures computed on the same level of disaggregation from the Census/ACS. In the following, the level of disaggregation is state-group-time rather than group-time. We introduce across-state variation for two major reasons. First, it provides an additional level of variation that permits more precise estimation of our parameters of interest.²⁷ Second, it allows us to estimate regression models that have stronger internal validity than models estimated on data with one observation per group-year. This is in the spirit of a recent literature on industrial composition and *local* labor market outcomes, such as Beaudry, Green and Sand (2012) and Chodorow-Reich and Wieland (2015).²⁸

More precisely, we calculate for each state, group and month all 16 transition rates between the two sectors, unemployment and non-employment, following the procedure outlined in section 3.3. For the growth of the service sector in a particular state it does not matter where new recruits come from. We thus include workers coming from out-of-state when calculating these rates. To avoid thin-cells we define groups by gender, education and the routine-dummy, but not by age. This is unproblematic given the minor role of age we have found in our decomposition exercises above. Because of the substantial noise and substantial high frequency movements in these rates we average monthly transition rates for all months in a symmetric 5 year interval around year t , where $t \in \{1980, 1990, 2000, 2005 - 2007\}$ are the Census/ACS years. We then match these rates to state-group-time level data on intersectoral earnings structures computed from the Census/ACS.

Let w_{gct}^d be the intersectoral difference in average earnings (or any of the four other statistics of the sector-specific earnings distributions used as outcomes above) for group g , work in state c in period t . Likewise, let λ_{gct}^{ij} be the monthly transition rates from employment state i to j , where $i, j \in \{N, S, U, NILF\}$. These rates are defined exactly as in section 3, but disaggregated to cells uniquely defined by (g, c, t) and averaged over the 5 years closest to t . The most flexible regression model we run is as follows:

$$w_{gt}^d = \sum_{i,j} \gamma^{ij} * \lambda_{gct}^{ij} + \mu_g + \mu_c + \mu_t + \varepsilon_{gct} \quad (6.2)$$

where μ_g, μ_c, μ_t are fixed effects for groups, states and Census/ACS years respectively. The parameters of interest are the $\gamma^{S,NS}$ and $\gamma^{NS,S}$ which measure the statistical association between intersectoral job-

²⁶An analysis of the direct relationship between earnings and job mobility on the worker level will be carried out in the next subsection, using data from the SIPP for the period 1996-2007.

²⁷Disaggregation to the MSA-level is infeasible because this would generate thin cells in the CPS data.

²⁸A central difference to these papers is that we cannot use a Bartik-type instrumental variable strategy because of our focus on long-run rather than medium-to-short run labor reallocation.

to-job transitions and intersectoral differences in the earnings distributions. According to equilibrium search models with on-the-job search these coefficients should be statistically significant and should have the opposite sign. In a two-sector Burdett-Mortensen model as in Hoffmann and Shi (2016), whether $\gamma^{NS,S}$ is positive or negative depends on the model parameters and the statistics one uses to summarize the earnings distributions. However, for calibrated and empirically plausible values of the parameters the signs are as follows (see Hoffmann and Shi, 2016):

- $\gamma^{NS,S} > 0; \gamma^{S,NS} < 0$ for intersectoral differentials of the earnings distribution in averages, medians and the 90th percentile.
- $\gamma^{NS,S} < 0; \gamma^{S,NS} > 0$ for intersectoral differentials in the 10th percentile of the earnings distribution.

These hypotheses reflect forces of labor supply and demand. On the one hand, the growing sector makes more revenue per worker and can thus offer higher wages. Since this force is likely to be strong at the upper end of the earnings distribution in the growing sector, it widens intersectoral earnings differential at that tail. On the other hand, if workers receive outside offers at a higher rate from a sector, then they are willing to take a wage cut to move into that sector. This force is likely to *reduce* the intersectoral earnings differential, or even create a negative intersectoral earnings differential, at the lower end of the earnings distribution.

Three further issues are noteworthy. First, we include all rates in the regression above because intersectoral worker transition rates are likely correlated with rates at which workers move from unemployment or NILF to the growing sector. To isolate the unequalizing force of direct job-to-job transitions, we thus need to control for the other transition rates as well. Second, since the transition rates must add up to one, $\sum_j \lambda_{gct}^{ij} = 1$, there is multicollinearity. We thus need to exclude four of the sixteen rates from the regression. Third, it follows from the dynamic system (4.1) that in the absence of measurement error and aggregation bias, the evolution of the stock of service sector employment is perfectly predicted by the transition rates. We therefore cannot include the service sector employment share as a control variable.

Estimates of the parameters of interest, $\gamma^{S,NS}$ and $\gamma^{NS,S}$, from various regression specifications, including our preferred model (6.2), are shown in Table 4. Each column in this table represents one regression. As in previous sections, we use as outcomes the intersectoral differentials in five summary statistics of earnings distributions, namely the average, the standard deviation, and 10th, 50th and 90th percentiles. Since the data are disaggregated to the group-state-time level, we only include the remaining transition rates as additional control variables. All regression models are weighted by cell

size, and standard errors are clustered on the group-state level. In Table 2 in the appendix we show the corresponding results when using ratios rather than differences of the measures of sectoral earnings inequality. As it turns out, conclusions remain unaltered and we thus focus our discussion on Table 4.

The first three columns show results from three different specifications when using intersectoral differences in average earnings as dependent variable. The first specification includes group fixed effects only and thus identifies the parameters of interest from within-group variation across states and over time. The parameter $\gamma^{NS,S}$ is highly significant and positive, in accordance with our hypothesis. Holding constant all other transition rates included in the regression, a rise in direct job-to-job transition into the service sector by one standard deviation (.027) is associated with an increase in the intersectoral differential in average earnings by $(.555)*(.027) = .015$ log points. This is slightly more than 10 percent of the standard deviation of the outcome variable (which is .125; see bottom of the table). Since on average earnings are lower in the service sector, this implies that sectoral reallocation has closed the intersectoral earnings gap. In the next two columns we first add state fixed effects and then time fixed effects. The regression coefficients remain positive, but are only weakly significant when estimating our preferred specification in the third column. On the other hand, the coefficient on $\gamma^{S,NS}$ is insignificant in all specifications and changes sign as we experiment with different regression models.

While these results provide some evidence in favor of our hypothesis, it is not surprising that they are not stronger. As argued above, structural transformation in form of an increase of revenue per worker in the service sector generates supply and demand effects whose strength differs across the earnings distributions and may be in balance at the center of the distribution. The model however makes sharper predictions about measures of the earnings dispersion. Indeed, Burdett-Mortensen models are mainly used as a framework for rationalizing residual earnings inequality. In the next three columns we thus use the intersectoral differentials in standard deviations as outcome. This yields results that are robust across specifications and fully in line with our hypothesis. It is remarkable that estimates of $\gamma^{NS,S}$ are positive and estimates of $\gamma^{S,NS}$ are negative in all specifications. Apart from one specification they are also highly significant.

Our results so far show that there is a strong statistical association between job-to-job transitions and earnings inequality, as predicted by frictional models with search on the job. The next nine columns explore which part of the distribution was most affected. The results can be summarized as follows: There is no evidence that intersectoral differentials in the 10th percentiles changed in the process of structural transformation. There are at least two explanations for this finding. On the one hand, the labor demand effect according to which service sector employers use higher earnings to attract workers from the other sector may approximately balance the labor supply effect according to which workers who move into the service sector become more willing to start in the sector with lower wages. On the

other hand, it is likely that minimum wage laws and other labor market regulations are binding in both sectors for this segment of the earnings distribution. In contrast, there is strong evidence that an increase in job-to-job transitions into the service sector is significantly related to an increase in median earnings relative to the non-service sector, holding all other rates constant. A similar relationship seems to exist for the differentials in the upper tail of the earnings distributions, but this effect is not estimated precisely when adding state- or time fixed effects. At the same time, in all but one case the regression coefficient on $\gamma^{S,NS}$ has the opposite sign. Again, this coefficient is not precisely estimated in most cases so that it lacks significance in all but one case.

Importantly, in all but one regression, namely our preferred specification estimated on the differentials in the 90th percentile, we can reject the joint null hypothesis that there is no joint effect of $(\gamma^{NS,S}, \gamma^{S,NS})$ on intersectoral differentials in the earnings structure. Even in the case where we cannot reject this null hypothesis are the point estimates consistent with those from the other specifications. Combined with our robust and strong evidence in favor of our hypothesis when using differentials in standard deviations as outcome we conclude that structural transformation via job-to-job transitions affected the earnings structures in the middle and upper rather than the lower part of the earnings distributions, that is the 90/10 and 50/10 percentile ratios. Again, this is consistent with an equilibrium job-search model because the labor demand effect of increasing revenue per worker is particularly dominant at the upper tail of the earnings distribution.

6.3. Discussion: Unobserved Heterogeneity or Frictions?

We have shown a strong statistical association between the evolution of the earnings structure in the service sector relative to the non-service sector on the one hand and the rate of sectoral reallocation on the other hand. In particular, subgroups of the population - whether they are defined by a combination of education, occupational task content, age and gender or by state of residence - that experienced a strong reallocation to the service sector saw a particularly pronounced rise in sectoral differentials of average earnings and various measures of earnings inequality. We have also documented robust evidence that the rate at which individuals reallocated via direct job-to-job transitions plays an important role. This raises the question of how to interpret these results. As pointed out frequently in earlier sections, we tend to view our results as evidence in favor of models that generate an intrinsic link between labor market flows and the earnings structure. Under this hypothesis intersectoral differentials in earnings structures reflect labor market frictions rather than compensating wage differentials or unobserved heterogeneity. Furthermore, changes in the earnings structure in the process of structural transformation are due to changes in the relative revenue per worker, altering the strategic considerations of firms and workers and thus the equilibrium wage posting strategies. It remains an open question whether a

neoclassical frictionless model of sectoral labor reallocation could reconcile the facts documented in this paper as well. Most importantly, it would need to explain why the earnings structure within narrowly defined worker groups has evolved so differently between the two sectors.

In this section we document a number of further empirical regularities that suggest worker flows to be a key driving force behind the rising earnings inequality in the service sector relative to the non-service sector. First, the rate of worker reallocation is undefined in any neoclassical model with two sectors. Rather, reallocation, and with it the types of workers who move to the service sector, is entirely captured by the evolution of the service sector employment share. In principle, a frictionless model with sorting on unobserved heterogeneity may generate a relationship between changes in net flows and changes in the earnings structures. However, it cannot explain the significant effects of job-to-job transition rates on intersectoral differentials in earnings inequality. Of course, empirically the evolution of stocks and flows can be highly correlated. But as shown above, our estimates of these effects are robust to inclusion of all transition rates between sectors, unemployment and non-employment as control variables.

Second, a neoclassical model is consistent with residual earnings inequality only in the presence of unobserved worker or firm heterogeneity. Since parts of our analysis disaggregate data to groups that are defined by age, education, gender and occupational task content, this unobserved heterogeneity would have to exist within these groups. If for example the best workers within each group reallocate the fastest to the service sector, then residual inequality and average earnings will increase in the service sector. At the same time, a neoclassical frictionless model with unobserved worker or firm heterogeneity would predict that residual inequality *decreases* in the non-service sector because an increasingly large part of the worker skill distribution will be pulled away from it. However, in the data residual inequality increases in both sectors. We show this by plotting in Figure 8 the kernel densities of earnings for each sector in 1960 and 2000. This finding is more consistent with a two-sector search model in which some non-service sector firms try to compete with high-earnings posting strategies of service sector firms.

Two final pieces of evidence in favor of a theory that puts intersectoral labor market flows at the center of understanding the change in the sectoral earnings structures rely on data on earnings changes and worker mobility within and across sectors. The only publicly available panel data set that contains this information on the monthly frequency is the SIPP. Due to a major survey redesign of these data between 1995 and 1996 we are forced to restrict our analysis to the post-1996 period. In particular, we utilize the two SIPP waves starting in 1996 and 2001 and apply, as much as possible, the same sample restrictions as in the Census and CPS. Our remaining empirical analysis focuses on EE transitions and earnings changes in two consecutive months. We thus keep only workers between the ages of 18 and 64 who are full-time employed, that is, who are working at least 135 hours and who are not on layoff. After

cleaning the earnings data from outliers following the procedure in Hoffmann and Shi (2016) and after only keeping pairs of observations on workers in two consecutive months we are left with 1,617,564 observations.

The first finding from these data is that monthly growth of real earnings is not statistically different from zero among workers who keep the same job in two consecutive months. In fact, there is a large mass point at zero, and the average monthly growth is .0005. Furthermore, the kurtosis is 14.7, making our findings consistent with those from administrative data on the annual level in Guvenen et al (2015). If the rise of within-group earnings inequality in the service sector relative to the goods sector was entirely due to worker sorting on unobserved heterogeneity in a frictionless environment, and if some of this heterogeneity was with respect to the ability to accumulate human capital, as postulated in recent papers on Ben-Porath type models such as Huggett, Ventura and Yaron (2010) or Guvenen, Kuruscu and Ozkan (2014), then one would expect within-job earnings growth to be significantly different in the two sectors. This however is not the case. In both sectors, earnings growth of workers who do not switch jobs is not significantly different from zero.

The second finding is that earnings growth at job-to-job transitions within sector is significantly different from zero no matter the sector. At the same time, job-to-job transitions are significantly larger in the service sector than the goods sector. More specifically, workers who change jobs within a sector see their earnings grow by 2.6 percent on average, a value that is highly statistically significant. Hence, job-to-job transitions within sector are an important source of earnings growth. These transitions take place at a rate that is about one-third higher in the service sector. Notice that it is important to focus on within-sector job-to-job transitions because search theory implies that on average such transitions should come with positive earnings growth. In contrast, without solving a particular model there are no sharp predictions on earnings growth when workers change sectors because they may or may not be willing to take an earnings cut to move into a growing sector. Our finding documented in this section is fully consistent with a two-sector search model. In particular, job changes within sector come with significant earnings growth, and transition rates are higher in the sector that displays more earnings dispersion.

7. Conclusion

We use long-run sectoral reallocation of labor from the goods to the service sector as a context for testing whether labor market frictions are an important source of aggregate residual earnings inequality. To this end we test directly the central prediction of search models with on-the-job-search, namely that job-to-job flows and the wage structure are intrinsically related. We argue that this has various advantages

relative to the more traditional approach of studying flows between firms within sector on the one hand or occupational reallocation and the wage structure on the other hand. First, in our context worker transitions take place between two large sectors rather than many small firms, allowing for a fine-grained analysis; second, sectoral reallocation has generated substantial time-series variation even within narrowly defined worker groups, permitting a fixed-effects research design; and third, job-to-job transitions in the process of structural transformation are likely driven by changes in labor demand due to changes in goods prices rather than skill-bias, at least once one conditions appropriately on worker and job characteristics.

Using several empirical approaches and data disaggregated to various levels we find that (i) sectoral reallocation took place within narrowly defined groups, including those with stable labor force participation rates; (ii) intersectoral job-to-job transitions played an important role in reallocating workers from the goods to the service sector; (iii) changes in intersectoral differentials in the earnings structure are strongly associated with the rate at which workers reallocate; (iv) this link is not driven by sorting on observables and is unlikely driven by sorting on unobservables. Although it is difficult to establish causality of this link due to the lack of an instrumental variable, we argue that it is difficult to rationalize the findings by any other theory than models with search frictions, rent sharing and search-on-the-job. We conclude that future work on sectoral reallocation and the earnings structure needs to incorporate labor market frictions and on-the-job-search.

References

- [1] Abraham, K. G. and L. F. Katz, 1986, "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *Journal of Political Economy* 94(3), 507-522.
- [2] Acemoglu, D. and D. H. Autor, 2011, "Skills, Tasks and Technologies: Implications for Employment and Earnings" *Handbook of Labor Economics*, Vol. 4, Elsevier.
- [3] Alvarez, F. and R. Shimer, 2011, "Search and Rest Unemployment," *Econometrica*, 79(1), 75-122.
- [4] Autor, D. H. and D. Dorn, 2013. "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market," *American Economic Review*, 103(5), 1553-97.
- [5] Bagger J., F. Fontaine, F. Postel-Vinay and J.-M. Robin, 2014, "Tenure, Experience, Human Capital and Wages: A Tractable Equilibrium Search Model of Wage Dynamics," *American Economic Review*, 104(6), 1551-96.
- [6] Beaudry, P, D. Green and B. Sand, 2012, "Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory," *Econometrica*, 80(3), 1063-1104.
- [7] Beaudry, P, D. Green and B. Sand, 2016, "The Great Reversal in the Demand for Skill and Cognitive Tasks," *Journal of Labor Economics*, 34(S1), S199-S247.
- [8] Brainard, L. and D. M. Cutler, 1993, "Sectorial Shifts and Cyclical Unemployment Reconsidered," *Quarterly Journal of Economics*, 108(1), 219-243.
- [9] Buera, F.J. and J.P. Kaboski, 2012. "The Rise of the Service Economy," *American Economic Review*, 102(6), 2540-69.
- [10] Buera, F.J., J.P. Kaboski and R. Rogerson, 2012. "Skill-Biased Structural Change," mimeo, University of Notre Dame.
- [11] Burdett, K. and D. Mortensen, 1998, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257-273.
- [12] Cahuc, P., F. Postel-Vinay, and J.-M. Robin, 2006, "Wage Bargaining with On-the-job Search: Theory and Evidence," *Econometrica*, 74, 323-364.
- [13] Chodorow-Reich, G. and J. Wieland, 2015, "Secular Labor Reallocation and Business Cycles," mimeo, Harvard University.
- [14] Clark, C., 1940, "The Conditions of Economic Progress," London: Macmillan.
- [15] Cortes, G.M., N. Jaimovich, C. Nekarda and H. Siu, 2015, "The micro and macro of disappearing routine jobs: a flows approach," mimeo, University of British Columbia.
- [16] Davis, S. J., R. J. Faberman, and J. Haltiwanger, 2006, "The Flow Approach to Labor Markets: New Data Sources and Micro-Macro Links," *Journal of Economic Perspectives*, 20(3), 3-26.
- [17] du Caju, Philip, A. Lamo, S. Poelhekke, G. Katay, and D. Nicolitsas, 2011, "Inter-industry Wage Differentials in EU Countries: What do Cross-Country Time Varying Data add to the Picture?" *Journal of the European Economic Association*, 8(2-3), 478-486.
- [18] Echevarria, C., 1997, "Changes in Sectoral Composition associated with Economic Growth," *International Economic Review*, 38(5), 431-452.

- [19] Edin, P.-A., and J. Zetterberg, 1992, "Interindustry Wage Differentials: Evidence from Sweden and a Comparison with the United States," *American Economic Review*, 82(5), 1341-1349.
- [20] Fallick, B. and C. A. Fleischman, 2004, "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," *Finance and Economics Discussion Series 2004-34*, Board of Governors of the Federal Reserve System.
- [21] Firpo, S., N. Fortin, and T. Lemieux, 2011. "Decomposition Methods," Chapter 1, *Handbook of Labor Economics*, vol. 4.
- [22] Foerster, A. T., P.-D. G. Sarte, and M. W. Watson, 2011, "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production," *Journal of Political Economy* 119(1), 1-38.
- [23] Gibbons, R. and L. Katz, 1992, "Does Unmeasured Ability Explain Inter-Industry Wage Differentials," *Econometrica*, 59(3), 515-535.
- [24] Gibbons, R., L. Katz, T. Lemieux, and D. Parent, 2005, "Comparative Advantage, Learnings and Sectoral Wage Determination," *Journal of Labor Economics*, 23(4), 681-723.
- [25] Goux, D. and E. Maurin, 1999, "Persistence of Interindustry Wage Differentials: A Reexamination using Matched Worker-Firm Panel Data," *Journal of Labor Economics*, 17(3), 492-533.
- [26] Guvenen, F., F. Karahan, S. Ozkan and J. Song, 2015, "What do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?," mimeo, University of Minnesota.
- [27] Guvenen, F., B. Kuruscu and S. Ozkan, 2014, "Taxation of Human Capital and Wage Inequality: A Cross-Country Analysis," *Review of Economic Studies*, 81, 818-850.
- [28] Helwege, J., 1992, "Sectoral Shifts and Interindustry Wage Differentials," *Journal of Labor Economics*, 10(1), 55-84.
- [29] Herrendorf, B., R. Rogerson and A. Valentinyi, 2013, "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103(7), 2752-2789.
- [30] Herrendorf, B., C. Herrington and A. Valentinyi, 2015, "Sectoral Technology and Structural Transformation," *American Economic Journal: Macroeconomics*, 7(4), 104-133.
- [31] Hoffmann, F. and S. Shi, 2016. "Burdett-Mortensen Model of On-the-Job-Search with Two Sectors," *Review of Economic Dynamics*, forthcoming.
- [32] Huggett, M., G. Ventura, and A. Yaron, 2007, "Sources of Lifetime Inequality", *American Economic Review*, 101(7), 2923-54.
- [33] Kambourov, G. and I. Manovskii, 2009, "Occupational Specificity of Human Capital", *International Economic Review*, 50(1), 63-115.
- [34] Kongsamut, P., S. Rebelo and D. Xie, 2001, "Beyond Balanced Growth," *Review of Economic Studies*, 68, 869-882.
- [35] Krueger, A. B., and L. H. Summers, 1988, "Efficiency Wages and the Inter-Industry Wage Structure," *Econometrica*, 56(2), 259-293.
- [36] Lee, D. and K. Wolpin, 2006, "Intersectoral Labor Mobility and the Growth of the Service Sector," *Econometrica* 74, 1-46.

- [37] Lilien, D., 1982, "Sectoral Shifts and Cyclical Unemployment," *Journal of Political Economy* 90(4), 777-793.
- [38] Lise, J. and J.-M. Robin (2014), "The Macro-Dynamics of Sorting between Workers and Firms," mimeo, University College London.
- [39] Long, J. B. Jr. and C. I. Plosser, 1983, "Real Business Cycles," *Journal of Political Economy* 91(1), 39-69.
- [40] Long, J. B. Jr. and C. I. Plosser, 1987, "Sectoral vs Aggregate Shocks in the Business Cycle," *American Economic Review*, 77(2), 333-336.
- [41] Loungani, P. and R. Rogerson, 1989, "Cyclical Fluctuations and Sectoral Reallocation: Evidence from the PSID," *Journal of Monetary Economics*, 23, 259-273.
- [42] Lucas, R. and E. Prescott, 1974, "Equilibrium search and unemployment," *Journal of Economic Theory*, 7(2), 188-209.
- [43] Manning, A., 2005, "Monopsony in Motion: Imperfect Competition in Labor Markets," Princeton University Press.
- [44] Meghir, C., Narita, R. and J.M. Robin, 2015, "Wages and Informality in Developing Countries," *American Economic Review* 105, 1509-1546.
- [45] Menzio, G. and S. Shi, 2010, "Block Recursive Equilibria for Stochastic Models of Search on the Job," *Journal of Economic Theory*, 145, 1453-1494.
- [46] Menzio, G. and S. Shi, 2011, "Efficient Search on the Job and the Business Cycle," *Journal of Political Economy*, 119(1), 38-62.
- [47] Menzio, G., I. Telyukova and L. Visschers, 2016, "Directed Search over the Life-Cycle," *Review of Economic Dynamics*, 19, 468-510.
- [48] Mortensen, D., 2005, "Wage Dispersion: Why Are Similar Workers Paid Differently?" MIT Press.
- [49] Mortensen, D. and C. A. Pissarides, 1994, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61(3), 397-415.
- [50] Moscarini, G. and F. Postel-Vinay, 2013, "Stochastic Search Equilibrium," *Review of Economic Studies*, 80(4), 1545-1581.
- [51] Moscarini, G. and F. Postel-Vinay, 2016, "Wage Posting and Business Cycles: A Quantitative Exploration," *Review of Economic Dynamics*, 19, 135-160.
- [52] Moscarini, G. and K. Thomsson, 2007, "Occupational and Job Mobility in the US," *Scandinavian Journal of Economics*, 109(4), 807-836.
- [53] Nagypal, E., 2008, "Worker Reallocation over the Business Cycle: The Importance of Employer-to-Employer Transitions," mimeo, Northwestern University.
- [54] Ngai, L. R. and C. A. Pissarides, 2007, "Structural Change in a Multisector Model of Growth," *American Economic Review* 97(1), 429-443.
- [55] Phelan, C. and A. Trejos, 2000, "The aggregate effects of sectoral reallocations," *Journal of Monetary Economics* 45, 249-268.

- [56] Pilossoph, L., 2014, "Sectoral Shocks and Move Unemployment," mimeo, Federal Reserve Bank of New York.
- [57] Robin, J.-M., 2011, "On the Dynamics of Unemployment and Wage Distributions," *Econometrica*, 79(5), 1327-1355.
- [58] Rogerson, R., 1987, "An Equilibrium Model of Sectoral Reallocation," *Journal of Political Economy* 95(4), 824-834.
- [59] Rogerson, R., 2005, "Sectoral Shocks, Human Capital, and Displaced Workers," *Review of Economic Dynamics* 8, 89-105.
- [60] Shi, S., 2009, "Directed Search for Equilibrium Wage-Tenure Contracts," *Econometrica* 77, 561-584.
- [61] Shimer, R., 2005a, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95(1), 25-49.
- [62] Shimer, R., 2005b, "The Cyclical Behavior of Hires, Separations, and Job-to-Job Transitions," *Federal Reserve Bank of St. Louis Review*, 87(4), 493-507.
- [63] Starr-McCluer, M., 1993, "Cyclical Fluctuations and Sectoral Reallocation: A Reexamination," *Journal of Monetary Economics*, 31, 417-425.
- [64] Swiecki, T., 2014, "Determinants of Structural Change," mimeo, University of British Columbia.
- [65] Williamson, S. D., 1990, "Sectoral Shifts, Labor Market Sorting, and Aggregate Fluctuations," *International Economic Review* 31(4), 935-952.
- [66] Thaler, R. H., 1989, "Interindustry Wage Differences," *Journal of Economic Perspectives* 3(2), 181-193.
- [67] Van den Berg, G. J. and G. Ridder, 1998), "An Empirical Equilibrium Search Model of the Labor Market," *Econometrica*, 66, 1183-1221.
- [68] Yedid-Levy, Y., 2016, "Why Does Employment in All Major Sectors Move Together over the Business Cycle?" mimeo, University of British Columbia.

FIGURE 1




FIGURE 2




FIGURE 3




FIGURE 4




FIGURE 5a

Model Fit - RD Corrected Transition Rates




Figure 5b

Model Fit - RD-Uncorrected Transition Rates




Figure 6a

Counterfactual Service Sector Employment -
RD-Corrected Transition Rates




Figure 6b

Counterfactual Service Sector Employment -
RD-Uncorrected Transition Rates




Figure 7a




Figure 7b




Figure 7c




Figure 8

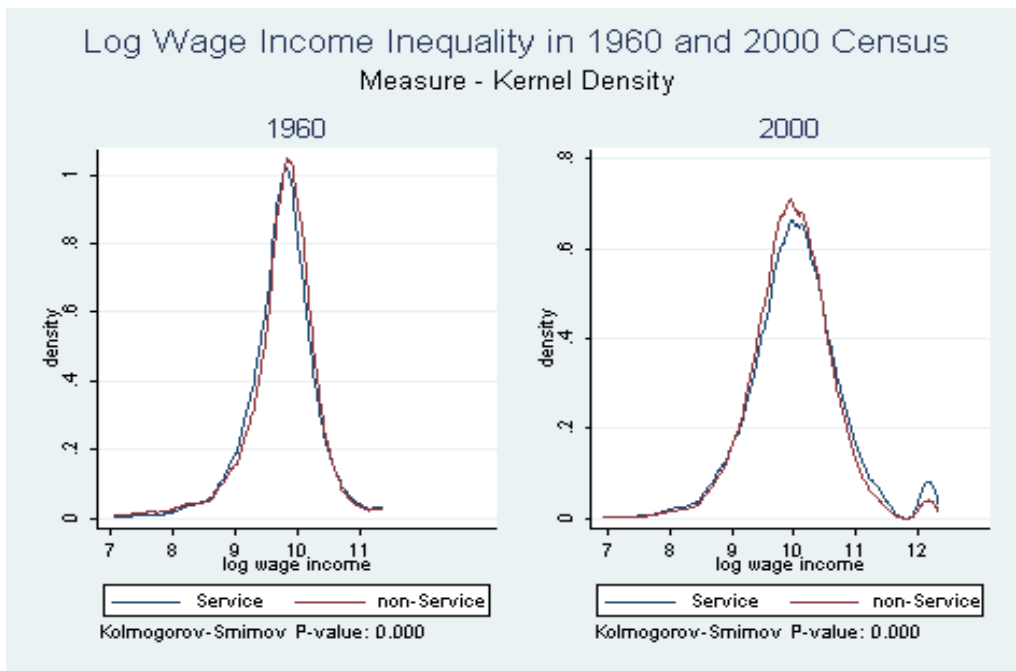


Table 1: Within-Group Changes in Service Sector Employment Shares, Census/ACS 1970 - 2007

PANEL A: MALE AND FEMALE

	Full Sample	1970-80	1980-90	1990-00	2000-07
female	0.0051*** (.0014)	0.0029*** (.0010)	0.0045*** (.0017)	0.005*** (.0016)	0.0053*** (.0019)
age 27-34	0.0035* (.0020)	0.001 (.0018)	0.0065*** (.0022)	0.006** (.0023)	0.0056** (.0028)
age 35-42	0.0035 (.0022)	-0.0048*** (.0007)	0.0067** (.0027)	0.0071*** (.0025)	0.0077*** (.0028)
age 43-52	0.0064** (.0024)	-0.0046*** (.0010)	0.0029 (.0022)	0.0084*** (.0027)	0.0127*** (.0034)
age 53-65	0.0027 (.0022)	-0.004*** (.0012)	-0.0006 (.0028)	0.0019 (.0028)	0.0041 (.0032)
high school degree	0.0118*** (.0017)	0.0093*** (.0008)	0.0113*** (.0018)	0.0127*** (.0021)	0.0134*** (.0026)
some college	0.0233*** (.0018)	0.0137*** (.0012)	0.0238*** (.0019)	0.0251*** (.0022)	0.0255*** (.0027)
routine occupation	-0.007*** (.0013)	-0.0023** (.0009)	-0.0066*** (.0015)	-0.0079*** (.0014)	-0.0091*** (.0017)

PANEL B: FEMALE

	Full Sample	1970-80	1980-90	1990-00	2000-07
age 27-34	0.0069** (.0026)	0.0016 (.0023)	0.011*** (.0026)	0.0107*** (.0026)	0.0103*** (.0030)
age 35-42	0.0065** (.0030)	-0.0049*** (.0008)	0.0099** (.0037)	0.0112*** (.0031)	0.0125*** (.0031)
age 43-52	0.009*** (.0031)	-0.0054*** (.0014)	0.0043* (.0023)	0.0114*** (.0029)	0.0188*** (.0040)
age 53-65	0.0046 (.0033)	-0.0044** (.0017)	0.0013 (.0049)	0.0043 (.0045)	0.0122*** (.0040)
high school degree	0.013*** (.0027)	0.0095*** (.0011)	0.0115*** (.0030)	0.0136*** (.0033)	0.0158*** (.0043)
some college	0.0245*** (.0030)	0.0121*** (.0016)	0.0242*** (.0031)	0.0261*** (.0035)	0.0294*** (.0044)
routine occupation	-0.0076*** (.0018)	-0.0006 (.0012)	-0.0066*** (.0022)	-0.0087*** (.0020)	-0.0116*** (.0024)

PANEL C: MALE

	Full Sample	1970-80	1980-90	1990-00	2000-07
age 27-34	0.0006 (.0026)	0.0005 (.0026)	0.0026 (.0029)	0.0019 (.0030)	0.0005 (.0035)
age 35-42	0.0008 (.0028)	-0.005*** (.0011)	0.0038 (.0041)	0.0035 (.0033)	0.0032 (.0035)
age 43-52	0.0041 (.0034)	-0.0044*** (.0011)	0.0016 (.0034)	0.0058 (.0040)	0.0099** (.0046)
age 53-65	0.0011 (.0027)	-0.0041** (.0016)	-0.0023 (.0035)	-0.0001 (.0032)	0.0059 (.0037)
high school degree	0.0113*** (.0020)	0.0087*** (.0012)	0.0113*** (.0024)	0.0125*** (.0025)	0.0136*** (.0031)
some college	0.0228*** (.0024)	0.0146*** (.0015)	0.0237*** (.0026)	0.0247*** (.0028)	0.0266*** (.0035)
routine occupation	-0.0064*** (.0016)	-0.0036*** (.0012)	-0.0066*** (.0020)	-0.0072*** (.0019)	-0.0085*** (.0022)

NOTES: Standard errors are clustered on the group-level and are in parenthesis. *** (**/*) denotes statistical significance at the 1% (5%/10%) significance level.

Table 2: Percentage Point Change in Service Sector Employment Share from Various Counterfactual Scenarios

	Rates held constant at 1976 Values (all Sector Specific)		
	<i>Intersectoral Transition Rates</i>	<i>UE Transition Rates</i>	<i>NILF-E Transition Rates</i>
RD-adjusted Transition Rates	6.10	6.00	8.82
Non-RD-adjusted Transition Rates	2.6	6.95	6.48
Percentage Point Change in Service Sector Employment Share, 1976-2007:			
<i>Data: March CPS</i>		9.42	
<i>Simulated: RD-adjusted</i>		10.45	
<i>Simulated: non-RD-adjusted</i>		10.63	

NOTES: The service sector employment share is measured relative to the population of working individuals rather than the entire population.

Table 3 - Sectoral Reallocation and Changes in the Relative Earnings Structure - Evidence from Stock Variables

		<i>Outcome: Group-Level Differences of</i>									
		average labor earnings		standard deviations of labor earnings		10% percentiles of earnings distributions		medians of earnings distributions		90% percentiles of earnings distributions	
Sample:											
	Full Sample	0.227 (.138)	0.243 (.207)	0.239*** (.070)	0.164* (.097)	0.169 (.157)	-0.108 (.223)	0.250** (.112)	0.368* (.202)	0.360** (.152)	0.404* (.220)
	1980 to 2005	0.255* (.143)	0.272 (.227)	0.252*** (.077)	0.197* (.101)	0.214 (.173)	-0.025 (.232)	0.308** (.122)	0.436* (.231)	0.376** (.152)	0.425* (.234)
	1970-80	-0.331 (.412)	-0.691* (.399)	0.44 (.392)	-0.159 (.376)	-2.305** (.990)	-2.446** (.926)	0.44 (.392)	-0.159 (.376)	0.489 (.326)	0.113 (.293)
	Full Sample age 27-64	0.457*** (.099)	0.589*** (.192)	0.177** (.082)	0.083 (.107)	0.405*** (.138)	0.176 (.232)	0.332*** (.086)	0.499** (.229)	0.560*** (.169)	0.756*** (.246)
Fixed Effects:											
	Group	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Year	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Sample Averages and Standard Deviations of Dependent Variables											
	average (raw earnings)	-0.163		0.077		-0.298		-0.142		-0.079	
	std (raw earnings)	0.144		0.045		0.212		0.138		0.125	


NOTES: The explanatory variable of interest is the group-time level service sector employment share. Each cell documents the coefficient from a different regression. Standard errors are clustered on the group level and are in parenthesis. *** (**/*) denotes statistical significance at the 1% (5%/10%) significance level.

TABLE 4: Sectoral Reallocation and Changes in the Relative Earnings Structure - Evidence from Employment Flows

	Outcome: Group-Level <i>Differences of</i>														
	average labor earnings			standard deviations of labor earnings			10% percentiles of earnings distributions			medians of earnings distributions			90% percentiles of earnings distributions		
EE-rate (to S)	0.555***	0.139	0.222*	0.048	0.209***	0.180**	0.510***	-0.144	-0.042	0.655***	0.307***	0.375***	0.507***	0.176	0.233
<i>[mean: 0.031, sd: 0.027]</i>	(0.111)	(0.103)	(0.114)	(0.059)	(0.061)	(0.084)	(0.131)	(0.110)	(0.150)	(0.115)	(0.106)	(0.117)	(0.158)	(0.157)	(0.182)
EE-rate (to NS)	0.111	-0.109	0.258	-0.603***	-0.553***	-0.528***	0.114	-0.169	0.223	-0.043	-0.249	0.054	-0.387	-0.583**	-0.245
<i>[mean: 0.012, sd: 0.013]</i>	(0.410)	(0.260)	(0.273)	(0.126)	(0.122)	(0.144)	(0.478)	(0.347)	(0.388)	(0.448)	(0.300)	(0.340)	(0.379)	(0.262)	(0.290)
Fixed Effects:															
Group	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Year	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
Sample Averages and Standard Deviations of Dependent Variables:															
average (raw earnings)	-0.060			0.026			-0.085			-0.065			-0.032		
std (raw earnings)	0.125			0.061			0.156			0.129			0.139		

NOTES: All regressions include as explanatory variables the transition rates between sectoral employment, unemployment, and NILF. Standard errors are clustered on the group-state-level and are in parenthesis. *** (**/*) denotes statistical significance at the 1% (5%/10%) significance level. All regressions include state fixed effects.

Appendix Figure 1



Appendix Table 1: Sectoral Reallocation and Changes in the Relative Earnings Structure, measured in Ratios - Evidence from Stock Variables

		<i>Outcome: Group-Level Ratios of</i>									
		average labor earnings		standard deviations of labor earnings		10% percentiles of earnings distributions		medians of earnings distributions		90% percentiles of earnings distributions	
Sample:											
	Full Sample	0.020 (.014)	0.019 (.021)	0.397** (.152)	0.273 (.192)	0.013 (.017)	-0.018 (.025)	0.023** (.012)	0.034* (.020)	0.032** (.014)	0.034 (.020)
	1980 to 2005	0.023 (.015)	0.022 (.023)	0.431** (.172)	0.349 (.216)	0.018 (.019)	-0.010 (.026)	0.029** (.013)	0.04* (.023)	0.034** (.014)	0.036 (.021)
	1970-80	-0.034 (.042)	-0.07* (.041)	0.831 (.859)	0.489 (.896)	-0.254** (.109)	-0.269** (.101)	0.048 (.040)	-0.013 (.038)	0.049 (.032)	0.012 (.028)
	Full Sample age 27-64	0.044*** (.010)	0.054*** (.018)	0.214 (.168)	0.065 (.217)	0.039** (.015)	0.012 (.025)	0.032*** (.009)	0.047** (.022)	0.051*** (.016)	0.066*** (.022)
Fixed Effects:											
	Group	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Year	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Sample Averages and Standard Deviations of Dependent Variables											
	average (raw earnings)	0.983		1.147		0.967		0.985		0.992	
	std (raw earnings)	0.015		0.084		0.024		0.014		0.012	

NOTES: The explanatory variable of interest is the group-time level service sector employment share. Each cell documents the coefficient from a different regression. Standard errors are clustered on the group level and are in parenthesis. *** (**/*) denotes statistical significance at the 1% (5%/10%) significance level.

APPENDIX TABLE 2: Sectoral Reallocation and Changes in the Relative Earnings Structure, measured in Ratios - Evidence from Employment Flows

	Outcome: Group-Level <i>Ratios</i> of														
	average labor earnings			standard deviations of labor earnings			10% percentiles of earnings distributions			medians of earnings distributions			90% percentiles of earnings distributions		
EE-rate (to S)	0.055***	0.012	0.021*	0.129	0.480***	0.371**	0.054***	-0.017	-0.006	0.065**	0.029***	0.037***	0.046***	0.014	0.020
[mean: 0.031, sd: 0.027]	(0.011)	(0.010)	(0.011)	(0.114)	(0.116)	(0.162)	(0.014)	(0.012)	(0.016)	(0.012)	(0.010)	(0.012)	(0.015)	(0.015)	(0.017)
EE-rate (to NS)	0.017	-0.006	0.032	-1.010***	-0.952***	-0.977***	0.019	-0.010	0.031	0.000	-0.021	0.010	-0.030	-0.050**	-0.017
[mean: 0.012, sd: 0.013]	(0.041)	(0.026)	(0.027)	(0.250)	(0.241)	(0.293)	(0.051)	(0.037)	(0.041)	(0.045)	(0.030)	(0.034)	(0.036)	(0.025)	(0.028)
Fixed Effects:															
Group	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Year	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
Sample Averages and Standard Deviations of Dependent Variables:															
average (raw earnings)	0.994			1.059			0.991			0.994			0.997		
std (raw earnings)	0.013			0.117			0.017			0.013			0.013		

NOTES: All regressions include as explanatory variables the transition rates between sectoral employment, unemployment, and NILF. Standard errors are clustered on the group-state-level and are in parenthesis. *** (**/*) denotes statistical significance at the 1% (5%/10%) significance level. All regressions include state fixed effects.