

The Cyclical Component of Labor Market Polarization and Jobless Recoveries in the US

Paul Gaggl^a
UNC Charlotte

Sylvia Kaufmann^b
Study Center Gerzensee

October 2019

Abstract: Based on quarterly occupation-level data from the US Current Population Survey for 1976-2013, we exploit common cyclical employment dynamics to identify two clusters of occupations that roughly correspond to the widely discussed notion of “routine” and “non-routine” jobs. After decomposing the cyclical dynamics into a cluster-specific (“structural”) and an occupation-specific (“idiosyncratic”) component, we detect significant structural breaks in the systematic dynamics of both clusters around 1990. We show that, absent these breaks, employment in the three “jobless recoveries” since 1990 would have recovered significantly more strongly than observed in the data, even after controlling for observed idiosyncratic shocks.

JEL: J21, E32, E24

Keywords: employment polarization; jobless recoveries; dynamic factor models

1. Introduction

Over the past several decades, the employment share of the lowest and highest skilled occupations increased, while it declined for middle skilled jobs. Over the same period, wages for middle skilled occupations grew substantially less than wages at the tail ends of the skill distribution. These trends are commonly referred to as “labor market polarization” and are in large part attributed to the widespread adoption of computing technology and the rising importance of offshoring—both of which potentially substitute for tasks performed by middle skilled workers and complement those performed by the highest skilled workers.²

Much less is known about the cyclical aspects of this apparent trend. In pioneering work, [Jaimovich and Siu \(2018\)](#) use an aggregated mapping of skills into jobs and document that 92% of the decline in per-capita employment within “routine” jobs in the US—ones that are considered easily replaceable by technology and require “middle” skills—occurred within a 12 month window of NBER dated recessions since the mid 1980s. Moreover, as is immediately apparent from panel A of Figure 1, “routine” (middle skill) occupations used to strongly rebound after recessions prior to 1990, while these swift rebounds were absent in the last three recessions. To the contrary, “non-routine” occupations—ones considered to directly or indirectly complement technology and comprising both low and high skilled workers—appear to be fairly immune to recessions and do not seem to have experienced a marked change in employment dynamics around 1990.

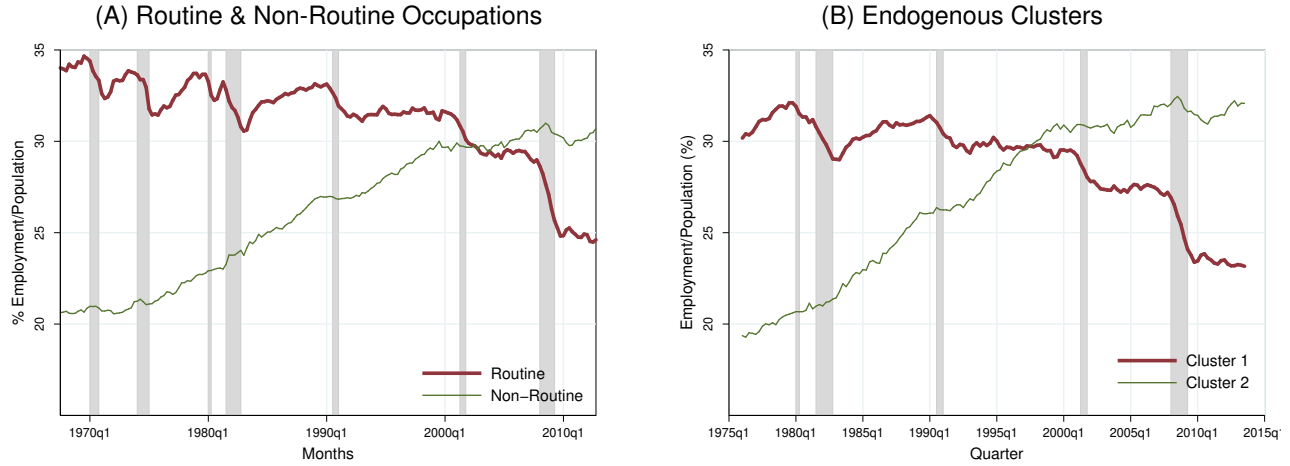
^aCorresponding author. University of North Carolina at Charlotte, Belk College of Business, Department of Economics, 9201 University City Blvd, Charlotte, NC 28223-0001; pgaggl@uncc.edu.

^bStudy Center Gerzensee, Dorfstrasse 2, P.O. Box 21, CH-3115 Gerzensee, Switzerland; sylvia.kaufmann@szgerzensee.ch.

¹We are grateful to Nir Jaimovich, Stan Rabinovich, Ayşegül Şahin, Henry Siu, Jim Nason, Travis Berge, Kyle Hood, David Wiczer as well as participants of the 2013 SEA meetings, the 2019 ASSA meetings and economics seminars at NC State University, Duke University, and UNC Chapel Hill for extremely valuable comments and feedback. We would further like to thank Debapriti Chakraborty and Jonathan Viscount for their help with data collection.

²[Acemoglu \(1999\)](#) was the first to document employment polarization in the US over the period 1983–1993. For more recent periods, [Goos and Manning \(2007\)](#) find similar patterns in the UK, [Goos, Manning, and Salomons \(2009\)](#) for 16 EU countries, and [Autor, Katz, and Kearney \(2008\)](#) as well as [Autor and Dorn \(2013\)](#) for the US. [Autor and Dorn \(2013\)](#) further show compelling evidence that PC adoption was more prevalent in areas with a historical abundance of workers performing “routine tasks”. Note that the rise in employment and wages for the highest skilled workers is likely due to complementarity with information and computing technology ([Akerman, Gaarder, and Mogstad, 2015](#); [Gaggl and Wright, 2017](#)) while the rise at the low end of the skill distribution is less directly related. [Autor and Dorn \(2013\)](#) attribute the latter to individuals’ love for variety. Therefore, a rise in income at the high end of the skill spectrum will cause an increase in the demand for services, mostly provided by low-skill labor.

Figure 1: Group Specific Employment Patterns



Notes: Panel A plots employment trends as grouped in the polarization literature. Data prior to 1983 are taken from the US Department of Labor’s *Employment & Earnings* publications and from FRED thereafter. The occupations are grouped as suggested in [Acemoglu and Autor \(2011\)](#). Panel B illustrates the cumulative growth of employment/population in each occupation assigned to factors 1 and 2 in model (1), which are listed in Table 2. Data for this graph are directly constructed from the monthly basic CPS files for the consistent panel of occupations compiled by [Dorn \(2009\)](#). The levels in both figures are imputed from quarterly growth rates and start with the level of employment/population at the beginning of each sample. We seasonally adjusted all time series from both data sources using the US Census X11 method.

If these stark cyclical patterns are truly due to a distinguishing feature, that is common to occupations within each broad category plotted in Figure 1, then we should be able to identify this group-specific characteristic as well as potential structural breaks from high frequency employment dynamics in the underlying detailed occupations.

Motivated by this observation, we develop a statistical framework that serves two purposes: first, we provide an “agnostic” approach to group detailed occupations into “clusters” that share common business cycle dynamics. Second, and jointly with our classification of occupations, we identify structural breaks in the cluster-specific cyclical dynamics, which allows us to revisit and formally test [Jaimovich and Siu’s \(2018\)](#) hypothesis that labor market polarization may play an important role in explaining the emergence of “jobless recoveries” in the US since 1990.³ Our test refines their original analysis in several ways: we conduct formal statistical inference on the estimated effects, we explicitly model the dynamics of occupation specific per-capita employment, we account for heterogeneity (across occupations) and asymmetry (across business cycle phases) in the effects of aggregate shocks, and we control for idiosyncratic, occupation-specific shocks.

To accomplish this, we estimate a dynamic factor model with latent clusters using detailed occupation level data from the US Current Population Survey (CPS) for the period 1976q1-2013q3. Even though our identification is based entirely on common cyclical employment dynamics across detailed occupations, our model uncovers two occupation clusters that almost perfectly coincide with “routine” and “non-routine” occupations, as identified in the polarization literature.

This finding is remarkable, as the original classification is based on cross-sectional variation in the “task content” of each occupation (see [Acemoglu and Autor, 2011](#), for a survey). While conceptually intuitive, this approach faces numerous practical challenges, including the lack of high quality longitudinal information and difficult to interpret ordinal task-content metrics (see [Autor, 2013](#), for a detailed discussion of these difficulties). Reassuringly, our “agnostic” clustering approach suggests that traditional blue collar jobs as well as sales and administrative support are most strongly associated with the gradually disappearing occupation group (routine/cluster 1), while managerial and

³For a detailed discussion of “jobless recoveries”—a recovery of output starting at the business cycle turning point, without a parallel recovery in employment—see for example [Gordon and Baily \(1993\)](#), [Groschen and Potter \(2003\)](#), [Bernanke \(2003\)](#), and [Bernanke \(2009\)](#). While there is no dominant theory for the emergence of this phenomenon, some recent theoretical contributions include [Koenders and Rogerson \(2005\)](#), [Bachmann \(2012\)](#), and [Berger \(2012\)](#).

service jobs—such as child and health care—are most representative of the strongly growing occupation group (non-routine/cluster 2). Moreover, this result highlights the strong tie between the well documented polarization trend and employment dynamics over the business cycle. A visual comparison of panels A and B in Figure 1 clearly reveals the similarity in the aggregate dynamics between the two classification schemes.

To capture group-specific business cycle dynamics, we adopt a Markov switching structure that accounts for asymmetric growth rates across expansions and recessions and allows for a potential break in these growth rates. Indeed, we find significant breaks in group-specific growth rates around 1990, which indicate that systematic routine employment growth during expansions completely vanished (from 0.16% to -0.03% quarterly) while non-routine occupations grew about half as fast on average (1.06% vs. 0.54% quarterly). Systematic routine job destruction during recessions almost doubled (from -0.7% to -1.24% quarterly) while non-routine jobs continued to grow during recessions on average, yet at a slower pace (0.4% vs. 0.23% quarterly).

These estimates allow us to assess to what extent these structural breaks can explain the three “jobless recoveries” in the US. Specifically, for each recovery since the 1990 recession, we construct counterfactual paths for occupation-specific per-capita employment under the assumption that the systematic, cluster-specific dynamics had remained unchanged after 1990. We find that aggregate per-capita employment would have recovered significantly more strongly in the absence of the observed structural change around 1990.

While our conclusions align with [Jaimovich and Siu \(2018\)](#), our empirical analysis has several advantages. First, our formal statistical model allows us to draw inference on the estimated effects. Second, while we analyze the effect of a break in occupation-specific stochastic trends, we control for both observed idiosyncratic as well as factor specific shocks. This allows us to disentangle the systematic, structural component from idiosyncratic shocks.

The results from our counterfactual exercises are further consistent with recent work exploiting variation across US states. Both [Jaimovich and Siu \(2018\)](#) and [Burger and Schwartz \(2018\)](#) illustrate that US states with more exposure to declining routine occupations are more likely to experience “jobless recoveries” after 1990. In addition, [Jaimovich and Siu \(2018\)](#) provide descriptive evidence suggesting that the link between job polarization and “jobless recoveries” is also present in other high income countries.⁴

2. Empirical Model

We specify a model for the growth in per-capita employment of occupation $i = 1, \dots, N$ and period $t = 1, \dots, T$, which we denote y_{it} . To capture the notion of “common dynamics” within a set of K distinct “occupation clusters”, we employ a factor model with K factors in which each occupation is exclusively assigned to one factor. We write this model compactly as

$$y_{it} = \sum_{k=1}^K \lambda_{ik} f_{kt} + \varepsilon_{it} \quad (1)$$

$$= \lambda_{i\delta_i} f_{\delta_i t} + \varepsilon_{it} \quad (2)$$

$$\phi_k(L) f_{kt} = \mu_{kS_t} + v_{kt}, \quad v_{kt} \sim N(0, 1) \quad (3)$$

$$\psi_i(L) \varepsilon_{it} = \epsilon_{it}, \quad \epsilon_{it} \sim N(0, \sigma_i^2) \quad (4)$$

where going from (1) to (2) reflects the fact that the latent classification indicator $\delta_i \in \{1, \dots, K\}$ assigns an occupation exclusively to one factor: $\lambda_{ik} \neq 0$ if $\delta_i = k$ and 0 otherwise, with $k = 1, \dots, K$. While each occupation only loads on one factor, the factor loadings may vary across occupations. This restriction is consistent with the procedure in the polarization literature, where each occupation is assigned to only one occupation group (e.g., routine or non-routine). The factor loading then captures the “intensity” with which occupation-specific per-capita employment growth is driven by common factor-specific dynamics.

To capture factor specific dynamics, we assume independent factor-specific $AR(p)$ processes, $\phi_k(L) = 1 - \phi_{k1}L - \dots - \phi_{kp}L^p$ driven by factor specific $N(0, 1)$ “shocks”, v_{kt} . The latent state indicator S_t follows a Markov process and indicates prevailing period-specific business cycle phases. The autoregressive structure in factor dynamics introduces persistence during cycle phases and plays an important role in understanding factor-specific recession-recovery

⁴[Furukawa and Toyoda \(2013\)](#) confirm this result for the case of Japan.

dynamics. While the timing of recessions/expansions is fully synchronized across factors and thus occupations, the conditional, phase-specific employment growth rate, μ_{kS_t} , is factor-specific. The factors capture all common dynamics within each occupation group, while any idiosyncratic, occupation-specific variation is captured by the independent processes ε_{it} , each of which follows an $AR(q)$ process $\psi_i(L) = 1 - \psi_{i1}L - \dots - \psi_{iq}L^q$.

Panel A of Figure 1 suggests that cycle-specific growth rates of routine and non-routine occupations—as defined in the polarization literature (Acemoglu and Autor, 2011)—may have experienced a break around 1990. Thus, to allow for a change in cycle-specific growth rates μ_{kS_t} , we specify four states for the state indicator $S_t = l$, $l = \{1, \dots, 4\}$ with the following interpretations: 1 = pre-break recession, 2 = pre-break expansion, 3 = post-break recession, 4 = post-break expansion. To induce state persistence in S_t , we specify a first-order Markov process with time-varying transition probabilities $\xi_{jl,t}(x_t, t) = P(S_t = j | S_{t-1} = l, x_t, t)$, $l, j = 1, \dots, 4$. The transition probabilities condition on real GDP growth, x_t , and a time trend, t . The former helps identifying business cycle phases while the latter helps identifying the break date. Mostly inspired by panel A of Figure 1, we impose two sets of restrictions on the transition probability matrix,

$$\xi_t = \begin{bmatrix} \xi_{11,t}(x_t, t) & \xi_{12,t}(x_t, t) & 0 & \xi_{14,t}(t) \\ \xi_{21,t}(x_t, t) & \xi_{22,t}(x_t, t) & \xi_{23,t}(t) & 0 \\ 0 & 0 & \xi_{33,t}(x_t) & \xi_{34,t}(x_t) \\ 0 & 0 & \xi_{43,t}(x_t) & \xi_{44,t}(x_t) \end{bmatrix}. \quad (5)$$

First, the structural break may only happen once. That is, state 1 and 2 cannot recur after reaching either state 3 or 4, which we implement by setting the lower-left block of transition probabilities to zero. Second, a potential structural break happens in transition from a recession to an expansion, or vice versa, which is ensured by setting the upper-right diagonal elements of the transition matrix to zero.⁵

Intuitively, we specify an econometric model for occupation-specific growth in per-capita employment and assume that the underlying drift in stochastic growth (μ_{kS_t}) stems from a common factor. Hence, we interpret polarization as a cluster-specific, common phenomenon, which may permanently affect the level of per-capita employment. Occupation-specific per-capita employment contains a stochastic trend, and the switching drifts reflect asymmetric growth rates across business cycle phases. Thus, the time series model implies that the “trend growth rates reflect the cycle” of the series. By postulating a structural break, we allow for changes in phase-specific trend growth around 1990. In addition, the model allows for idiosyncratic occupation-specific stochastic growth without drift (ε_{it}) to capture any other level effects unrelated to common underlying drivers.⁶

2.1. Bayesian Estimation

We estimate the model using Bayesian Markov Chain Monte Carlo (MCMC) methods. To describe our sampling procedure concisely, we introduce the following notation: Denote vectors collecting all values of a specific variable or parameter in bold face, e.g. $\mathbf{y} = \{y_{it} | i = 1, \dots, N, t = 1, \dots, N\}$, $\mathbf{y}_i = \{y_{it} | t = 1, \dots, N\}$, $\boldsymbol{\delta} = \{\delta_i | i = 1, \dots, N\}$, $\mathbf{S} = \{S_t | t = 1, \dots, T\}$ or $\boldsymbol{\psi} = \{\psi_{ij} | i = 1, \dots, N, j = 1, \dots, q\}$. Gather all model parameters in $\theta = \{\boldsymbol{\lambda}, \boldsymbol{\psi}, \boldsymbol{\phi}, \boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\gamma}\}$, and define the augmented parameter vector $\vartheta = \{\theta, \mathbf{f}, \boldsymbol{\delta}, \mathbf{S}\}$.

MCMC estimation provides a sample from the joint posterior distribution of all model parameters and latent variables by combining the likelihood with the prior distribution:

$$\pi(\vartheta | \mathbf{y}) \propto L(\mathbf{y} | \mathbf{f}, \boldsymbol{\delta}, \theta) \pi(\mathbf{f}, \mathbf{S}, \theta) \pi(\mathbf{S} | \mathbf{x}, \mathbf{t}, \boldsymbol{\gamma}) \pi(\boldsymbol{\delta}) \pi(\theta) \quad (6)$$

To obtain draws from (6), we simulate iteratively from the conditional posterior distributions of: the factors, $\pi(\mathbf{f} | \mathbf{y}, \mathbf{S}, \boldsymbol{\delta}, \theta)$; the business cycle indicator $\pi(\mathbf{S} | \mathbf{f}, \boldsymbol{\xi}, \boldsymbol{\mu}, \boldsymbol{\phi})$; the classification indicator $\pi(\boldsymbol{\delta} | \mathbf{y}, \mathbf{f}, \boldsymbol{\psi}, \boldsymbol{\sigma})$; the parameters of the transition distribution $\pi(\boldsymbol{\gamma} | \mathbf{S}, \mathbf{x}, \mathbf{t})$; the remaining parameters $\pi(\theta_{-\boldsymbol{\gamma}} | \mathbf{y}, \mathbf{f}, \mathbf{S}, \boldsymbol{\delta})$, where $\theta_{-\boldsymbol{\gamma}}$ is θ excluding $\boldsymbol{\gamma}$.

All prior and posterior distributions are standard. However, it is worth emphasizing that we use a discrete uniform prior for the latent classification indicator δ_i , which implies that our estimated occupation classification is entirely

⁵In the online appendix, we provide additional details on the parametrization of ξ_t . In addition, we generalize the model to factor-specific state variables and break dates. Based on several generalized specifications, we illustrate that the conclusions presented in the main text remain unchanged.

⁶In the online appendix, we provide an analytical derivation and graphical illustrations for a decomposition of per-capita employment into a common trend, a permanent idiosyncratic, and a cyclical component.

identified from occupation-specific employment dynamics. This is in stark contrast to the polarization literature, in which occupations are grouped according to cross-sectional information in the task content of occupations (e.g., Autor, Levy, and Murnane, 2003). We provide full details for the specification of the likelihood, all priors, and the derivation of the posterior distribution in the online appendix.

Sampler convergence for the state variable is facilitated by the effect of GDP growth on the transition probabilities and by normalizing the trend t to 0 at the expected break date. Convergence for the classification indicator is based on occupation-specific data dynamics. Strongly co-moving occupations will group into clusters and determine posterior inference on cluster-specific parameters. Occupations less correlated with cluster-specific dynamics receive a small or weak factor loading and will mostly be driven by the idiosyncratic component.

3. Data

Our main data source are detailed individual level data from the basic US Current Population Survey (CPS) covering the period 1976q1–2013q3. Based on CPS sampling weights we estimate employment levels at the detailed occupation level on a monthly frequency throughout the entire sample. Since the US Department of Labor’s (DOL) classification of occupations changes several times during our sample period, we aggregate individuals into a panel of 330 consistent occupations as designed by Dorn (2009).⁷ For our preferred specification, summarized in Table 1, we further aggregate the detailed occupations into 21 groups based on Dorn (2009) and take quarterly averages. Following Jaimovich and Siu (2018) we use the share of employment within the US population of age 16 and older to measure each occupation’s employment dynamics over time.

To compare our results to Jaimovich and Siu (2018), we also replicate their dataset, spanning the period 1967–2012, for which data prior to 1983 are taken from the DOL’s *Employment & Earnings* publications and from the Federal Reserve Bank of St. Louis’ FRED database thereafter.⁸ We further seasonally adjust all time series (from both data sources) using the US Census X11 method. Panel A of Figure 1 displays the resulting series for employment as a share of population for the two groups of routine and non-routine occupations, as defined by Acemoglu and Autor (2011).

Finally, we employ real GDP growth as our aggregate measure to help identify business cycles and we draw these data from FRED.

4. Empirical Analysis

We estimate model (1) using the MCMC sampler described in the previous section. In total, we draw 500,000 times out of the posterior distribution and discard the first 300,000 as burn-in. To remove autocorrelation across draws, we retain every fourth of the remaining 200,000 draws.⁹ Our preferred specification is summarized in Table 1. Note that we obtain the most precise factor assignment (see Section 4.1) when we set $K = 2$ and since the ultimate goal of this study is to analyze aggregate labor market dynamics, we choose the specification for which the variance share explained by cluster specific variation is largest. In particular, the online appendix lists this statistic for alternative AR lag lengths, p and q , and shows that a specification with $p = q = 2$ performs best, on average, according to this metric.¹⁰

⁷The DOL has implemented the latest change in their occupation classification system in 2011, and we thank Nir Jaimovich for providing a crosswalk (as used in Cortes, Jaimovich, Nekarda, and Siu, 2014) to extend Dorn’s (2009) panel of occupations beyond this newest change of occupation classifications.

⁸These data contain the level of employment and are already aggregated to about 10 broad occupation groups. Unfortunately, the group definitions are neither fully consistent over time (especially prior to 1983) nor between the aggregates in FRED and in the *Employment & Earnings* publications. Thus, we group occupations into the four broad occupation groups suggested by Acemoglu and Autor (2011): non-routine cognitive (professional, managerial, and technical occupations), routine cognitive (clerical, support, and sales occupations), routine manual (production and operative occupations), non-routine manual (service occupations).

⁹In the online appendix, we plot the retained draws for key model parameters after conditioning on the specific cluster assignment tabulated in Table 2. The plotted series clearly indicate convergence of the MCMC sampler.

¹⁰In the online appendix, we show that our main conclusions are robust to the number of clusters and to state-specific break dates.

Table 1: Preferred Model Specification

A. Specification		
Number of Factors	K	2
Number of States	S	1 = pre-recession, 2 = pre-expansion, 3 = post-recession, 4 = post-expansion
Factor AR lags	p	2
Idiosyncratic AR lags	q	2
B. Sample		
Employment Variable	$y_{i,t}$	Quarterly growth in employment/population (CPS basic files, SA X11)
Aggregate Variable	x_t	Quarterly growth in real GDP
Estimation Sample	T	150 quarters: 1976q2 – 2013q3
Occupation Groups	N	21 Dorn (2009) detailed occupation groups
B. Posterior Sampler		
Total Posterior draws		500,000
Burn-in		300,000
Retained Observations		50,000 (every fourth draw after burn-in)

4.1. Occupation Clusters

Our model identifies two clusters of occupations with distinct cyclical patterns in employment dynamics. Panel B of Figure 1 illustrates that the identifying feature of cluster 2 is the relatively steady average growth in the employment/population ratio throughout the entire period 1976–2013. Per-capita employment in this group grew from less than 20% in 1976 to around 33% in 2013. Moreover, employment of this group does not appear particularly “cyclical”.

On the other hand, cluster 1 groups occupations with employment patterns that differ dramatically from those of cluster 2. First, per-capita employment in this group has declined from around 33% at its high in 1980 to around 24% at its low in 2013. Second, employment in these occupations appears highly “cyclical”. Growth rates obviously differ between recessions and expansions, and there appears to be a visual change in these growth rates around 1990.

Overall, the groups identified by our preferred model specification resemble the patterns in employment dynamics displayed in panel A of Figure 1, which groups occupations according to the polarization literature (e.g., [Acemoglu and Autor, 2011](#)).¹¹ While the aggregate levels in panels A and B of Figure 1 don’t perfectly match up, the long run growth patterns as well as cyclical dynamics are virtually the same. Both aggregation schemes indicate roughly a 50% increase in the employment/population ratio for cluster 2 (non-routine in Panel A) and about a 25% decrease for cluster 1 (routine) over the period 1980 to 2013. Moreover, both aggregation schemes highlight a visual change in the cyclical dynamics of the declining occupation group after 1990.

Our clustering approach at the disaggregated level allows us to further analyze the composition of the two identified occupation groups. The first two columns of Table 2 tabulate the posterior assignment probabilities to the two factors. Almost all of the 21 occupation groups are nearly perfectly assigned to one of the two clusters. Panels A and B respectively group occupations for which the posterior probability of being determined by factors 1 and 2 is larger than 50%.¹² Only a handful of occupations have an assignment probability of less than 2/3. However, these occupations display an employment level nearly constant throughout the entire sample and their share in total employment is very small. Therefore, these occupations do not contain much information about common factor dynamics.

In addition to the assignment probabilities, the last four columns of Table 2 report the posterior mean and median factor loading, $\lambda_{ik}|\bar{\delta}$, as well as the associated 68% posterior coverage interval, conditional on the mean factor assign-

¹¹Note that both panels in Figure 1 simply plot the data, but for different aggregation schemes. Panel A is based on the same data and aggregation as in [Jaimovich and Siu \(2018\)](#), while in panel B we use our detailed CPS dataset and aggregate the employment/population ratios in the two clusters of occupations listed in Table 2.

¹²The classification is robust for all occupations but those not clearly assigned to either cluster to a model specification allowing for state-specific break dates, see the online appendix.

Table 2: Cluster Analysis: Factor Assignment

21 Occupation Groups (Dorn, 2009)	Assignment δ_i		Factor Loading $\lambda_{ik} \bar{\delta}$			
	$Pr[\delta_i = 1 y]$	$Pr[\delta_i = 2 y]$	Mean	Median	68% Coverage	
<i>A. Factor 1 (Routine)</i>						
F.1 Machine Operators, Assemblers, and Inspectors	0.999	0.001	1.107	1.108	0.984	1.239
E.2 Construction Trades	0.993	0.007	0.896	0.910	0.782	1.017
E.4 Precision Production	0.995	0.005	0.855	0.866	0.761	0.957
F.2 Transportation and Material Moving	1.000	0.000	0.690	0.697	0.637	0.766
E.1 Mechanics and Repairers	1.000	0.000	0.636	0.634	0.566	0.702
B.2 Sales	0.954	0.046	0.464	0.462	0.400	0.528
B.3 Administrative Support	0.724	0.276	0.228	0.245	0.164	0.311
C.1 Housekeeping and Cleaning	0.505	0.495	0.196	0.192	0.052	0.324
<i>B. Factor 2 (Non-Routine)</i>						
E.3 Extractive	0.446	0.554	0.816	0.808	0.350	1.253
A.2 Management Related	0.000	1.000	0.777	0.788	0.672	0.878
C.37 Misc. Personal Care and Service	0.182	0.818	0.776	0.772	0.566	0.954
A.1 Executive, Administrative, and Managerial	0.004	0.996	0.557	0.551	0.479	0.627
C.36 Child Care Workers	0.305	0.695	0.494	0.462	0.314	0.649
C.32 Healthcare Support	0.162	0.838	0.432	0.430	0.357	0.517
A.3 Professional Specialty	0.009	0.991	0.394	0.395	0.346	0.440
C.2 Protective Service	0.377	0.623	0.297	0.299	0.217	0.377
C.33 Building/Grounds Cleaning/Maintenance	0.289	0.711	0.263	0.254	0.165	0.375
B.1 Technicians and Related Support	0.238	0.762	0.247	0.246	0.168	0.319
C.34 Personal Appearance	0.446	0.554	0.158	0.164	0.044	0.276
C.31 Food Preparation and Service	0.493	0.507	0.094	0.107	0.035	0.164
C.35 Recreation and Hospitality	0.443	0.557	-0.033	-0.035	-0.217	0.133

Notes: The first two columns report the fraction of posterior draws that classify each occupation into either factor $k = 1$ or $k = 2$. Panel A groups occupations with $Pr[\delta_i = 1|y] > 1/2$ while panel B collects those with $Pr[\delta_i = 2|y] > 1/2$ based on 50,000 retained posterior draws. The last three columns report the posterior mean, median, as well as the upper and lower bound of the 68% posterior coverage region for the factor loading $\lambda_{ik}|\bar{\delta}$, conditional on the mean factor assignment, $\bar{\delta}$. Within each panel, the occupations are sorted in decreasing order by their conditional factor loading $\lambda_{ik}|\bar{\delta}$.

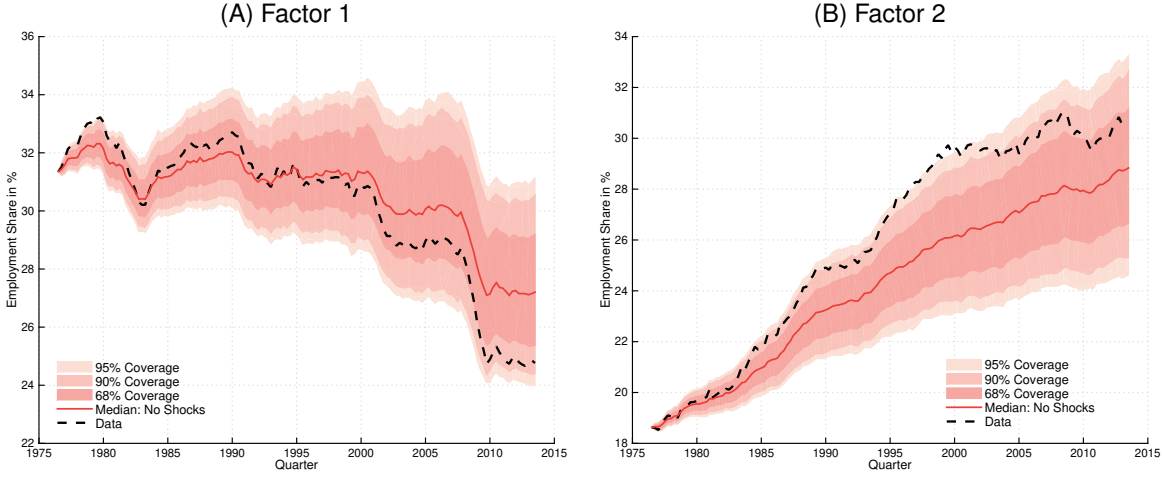
ment, $\bar{\delta}$.¹³ Notice that all but one of these intervals exclude zero. The occupation with an insignificant factor loading is precisely one of the service occupations for which the assignment probability is not decisive.

Within panels A and B of Table 2, we sort occupations in decreasing order of the posterior mean/median factor loading. This provides a measure of the “intensity” with which a given occupation is influenced by the common factor dynamics and we observe that a considerable amount of heterogeneity is present across the factor loadings. This measure indicates that employment dynamics for an occupation with a factor loading close to zero are mostly driven by idiosyncratic dynamics, captured by the mean zero AR(q) process ε_{it} in our model.

The long run patterns (see Figure 1) as well as the composition (see Table 2) of the two identified clusters are consistent with recent evidence presented in Autor and Dorn (2013). They show that polarization in the US labor market over the period 1980-2005 is mainly driven by the growth in service and in “abstract cognitive” occupations combined with the decline in “routine” occupations, which are easily replaceable by technology or offshoring. In line with their evidence, panel B of Table 2 illustrates that cluster 2 largely consists of managerial and professional specialty occupations and of a number of service occupations. However, our classification also suggests that some

¹³That is, these moments are computed from joint posterior draws for which each occupation is assigned exactly as grouped in Table 2.

Figure 2: Structural Employment Dynamics



Notes: Conditional on $Pr[\hat{\delta}_i|y] > 0.5$, the figure shows the cluster-specific observed aggregate employment/population ratio and the estimated implied aggregate employment/population ratio, setting $\varepsilon_{it} = 0$. We construct these aggregates by summing occupation-specific implied levels. The posterior coverage region is obtained from quantiles of the empirical posterior implied by all MCMC draws, conditional on the mean factor assignment, $\bar{\delta}$.

service occupations, like “recreation and hospitality” or “housekeeping and cleaning”, are not decisively assigned to either factor.

4.2. Cluster-Specific Cyclical Dynamics

Conditional on the mean factor assignment from Table 2, we now shift our focus to the decomposition of employment dynamics into a “structural” (factor specific) and an “idiosyncratic” (occupation-specific) component. Figure 2 illustrates the cluster-specific per-capita employment *levels*, implied by the estimated structural factor dynamics. Notice that we estimated the model in growth rates and the displayed level dynamics reflect the cumulative effect of all shocks encountered throughout the entire sample. To construct these graphs, we first recursively compute the occupation-specific implied level series (setting $\varepsilon_{it} = 0$ for all t) for each of the retained joint posterior draws, conditional on the mean posterior factor assignment probability $\bar{\delta} = \{\hat{\delta}_1, \dots, \hat{\delta}_N\}$:

$$\tilde{y}_{i,t}^{(m)} = \left(1 + \hat{\lambda}_{i,\hat{\delta}_i}^{(m)} \hat{f}_{\hat{\delta}_i,t}^{(m)}\right) \tilde{y}_{i,t-1}^{(m)} \quad \text{for all draws } m|\bar{\delta}, \quad (7)$$

where m is the respective MCMC draw, $\tilde{y}_{i,t}^{(m)}$ denotes the implied level in period t , $\hat{f}_{\hat{\delta}_i,t}^{(m)}$ is the estimated factor specific growth, and $\tilde{y}_{i,0}^{(m)} = y_{i,0}$ is the observed initial level of per-capita employment in occupation i . Based on these occupation-specific implied level series we then aggregate across occupations within each cluster:

$$\hat{Y}_{kt}^{(m)} = \sum_{i|\hat{\delta}_i=k} \tilde{y}_{i,t}^{(m)} \quad \text{for all draws } m|\bar{\delta} \text{ and } k = 1, 2 \quad (8)$$

These implied cluster-specific level series capture the predictive ability of the common factor dynamics, while the deviations from observed employment levels are explained by the accumulation of the occupation-specific idiosyncratic component, ε_{it} . We utilize the quantiles of the empirical posterior distribution of the cluster-specific level series to compute posterior coverage regions.¹⁴

Figure 2 shows that, for factor one, the data lie within the 95% posterior coverage region throughout the entire sample. This implies that the common dynamics of occupations in cluster 1 capture a large portion of the aggregate

¹⁴Note that, while inference for this simulation is conditional on the mean factor assignment, $\bar{\delta}$, we fully account for conditional uncertainty in all other parameters.

dynamics in employment. In fact, the structural component of factor 1 explains about 61% of the variation in the data. The structural component of factor 2 captures less of the observed aggregate employment dynamics, given that the data series lies marginally outside the 95% coverage region for some quarters in the second half of the 1990s (see panel B of Figure 2). Nevertheless, the explained variance share of the structural component in cluster 2 amounts to about 59%.

Figure 2 further reveals that there is a significant amount of variation originating from idiosyncratic, occupation-specific variation, captured by ε_{it} in our framework. The ability to distinguish between structural and idiosyncratic dynamics is an important contribution of our statistical framework. Note that the grouping of occupations in panel A of Figure 1 is derived from cross sectional information on the task content of each occupation. While the polarization literature has documented common long-run polarizing trends in these groups (e.g., Acemoglu and Autor, 2011), this does not necessarily imply that the underlying occupations share the same cyclical dynamics. For example, it is a-priori possible that the aggregate short-run dynamics of routine occupations are almost entirely driven by machine operators, assemblers, and inspectors (F.1 in Table 2), an occupation group that is highly concentrated in the manufacturing sector. Likewise, the aggregate cyclical dynamics of non-routine workers could entirely be accounted for by management occupations. While we would not be able to detect this in an illustration like Figure 1, our statistical approach explicitly separates such idiosyncratic variation from the cyclical dynamics common to all occupations within each cluster.¹⁵

In Sections 4.3 and 4.4, we will utilize this aspect of our model to draw inference about the importance of structural change in the common component for explaining aggregate employment dynamics after the 1990 recession, while simultaneously controlling for the observed idiosyncratic variation in ε_{it} .

4.3. A Structural Break in the Systematic Cluster Dynamics

Visual inspection of panel A of Figure 1 suggests that the marked change in the business cycle dynamics of “routine” occupations around 1990 may constitute a structural break. In a series of counterfactuals based on descriptive statistics, Jaimovich and Siu (2018) illustrate that this apparent break may have the potential to explain a substantial part of the three “jobless recoveries” since 1990.

Our econometric approach allows us to revisit this question more formally. Specifically, our Markov switching specification allows for a structural break in the cyclical dynamics of both factors and our posterior estimates provide evidence for a break having occurred at the end of the long-lasting expansion during the 1980s (see panel A of Figure 3).

Panel A of Figure 3 depicts the estimated posterior state probabilities and reveals an almost perfect match with the NBER’s business cycle dating committee’s classification. The state probabilities are jointly identified from variation in occupation-specific per-capita employment and variation in real GDP growth. In particular, GDP growth mainly helps to identify expansions and recessions, while the factor dynamics—inferred from dynamics in occupation specific per-capita employment—identify the structural break. Specifically, panel A of Figure 3 reveals that the structural break occurred during the 1990/91 recession, as our posterior estimates assign this recession to state 3.

To visualize the nature of this break, panel B of Figure 3 illustrates the systematic element of the estimated factor specific growth in per-capita employment. To construct this figure, we first compute the factor specific systematic dynamics dynamics

$$\tilde{f}_{t,k}^{(m)} = \hat{f}_{k,t}^{(m)} - \hat{v}_{k,t} = \hat{\mu}_{k,S_t}^{(m)} + \sum_{j=1}^p \hat{\phi}_{k,j}^{(m)} \hat{f}_{k,t-j}^{(m)} \quad \text{for all draws } m|\bar{\delta} \text{ and } k \in \{1, 2\} \quad (9)$$

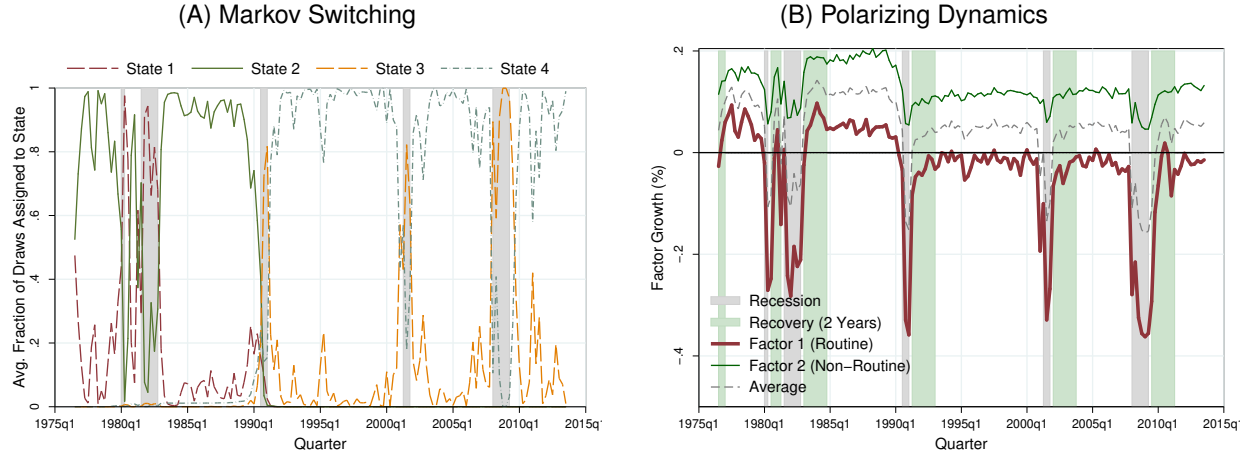
where we initialize this “filtered” series with the estimated factor means $\hat{\mu}_{k,S_j}$ for the first p periods $j = 0, \dots, p - 1$. To illustrate the aggregate growth contribution of these systematic dynamics, panel B of Figure 3 plots the posterior median of the weighted average

$$\tilde{F}_{t,k}^{(m)} = \sum_{i|\hat{\delta}_i=k} w_{i,t} \hat{\lambda}_{i,\hat{\delta}_i}^{(m)} \tilde{f}_{t,k}^{(m)}, \quad (10)$$

where $\hat{\lambda}_{i,\hat{\delta}_i}^{(m)}$ is the estimated factor loading and $w_{i,t}$ is the relative size of occupation i at time t .

¹⁵In fact, in the online appendix, we derive a decomposition that allows us to compute the influence of the common and idiosyncratic component

Figure 3: Markov Switching & Polarizing Dynamics



Notes: Panel A illustrates the estimated assignment probabilities (mean of retained MCMC draws) for the four latent states of the economy: pre-break recession (state 1), pre-break expansion (state 2), post-break recession (state 3), post-break expansion (state 4). Panel B illustrates the posterior median of the systematic, factor-specific aggregate growth contribution $\bar{F}_{t,k}^{(m)}$, as computed in equation (10), for the two factors $k \in \{1, 2\}$. The dashed line is the average of the systematic factor dynamics.

This figure reveals several interesting aspects of the systematic cyclical dynamics within each occupation cluster: first, both clusters feature substantial systematic cyclical variation, which is not immediately visible in Figure 1. Second, both clusters experience a structural break in the dynamics of recessions and expansions. Specifically, we find that systematic routine employment growth, during expansions completely vanished (from $\hat{\mu}_{1,2} = 0.16\%$ to $\hat{\mu}_{1,4} = -0.03\%$ quarterly) while non-routine occupations grew about half as fast on average ($\hat{\mu}_{2,2} = 1.06\%$ vs. $\hat{\mu}_{2,4} = 0.54\%$ quarterly). Most strikingly however, systematic routine job destruction during recessions almost doubled (from $\hat{\mu}_{1,1} = -0.7\%$ to $\hat{\mu}_{1,3} = -1.24\%$ quarterly) while non-routine jobs continued to grow systematically during recessions, yet at a slower pace ($\hat{\mu}_{2,1} = 0.4\%$ vs. $\hat{\mu}_{2,3} = 0.23\%$ quarterly). Panel B of Figure 3 makes clear that these underlying breaks in mean factor growth are preserved in the aggregate.

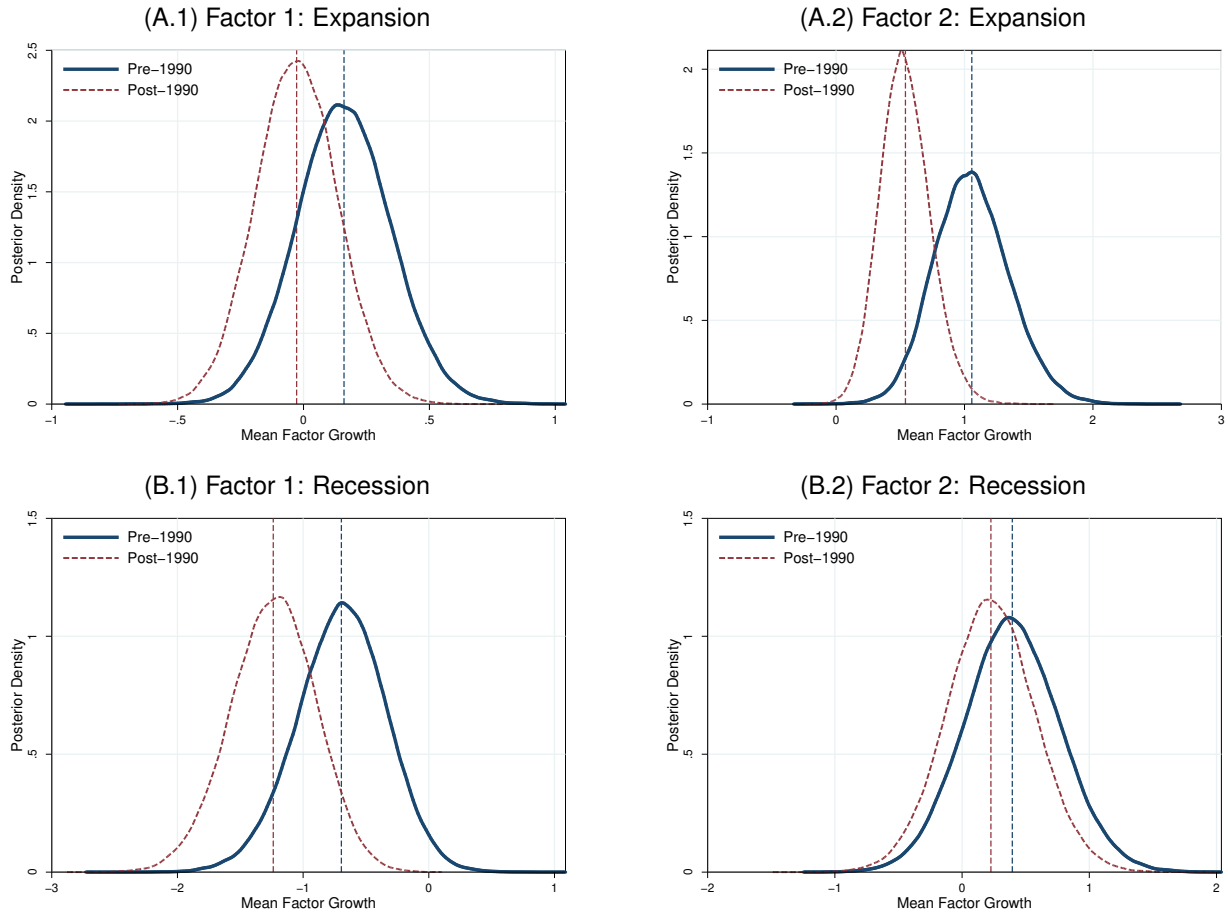
Moreover, these state-specific breaks imply that routine occupations have become “more cyclical” after 1990, while non-routine occupations became “less cyclical”. That is, the difference between the median growth rates during expansions and recessions has increased (decreased) for routine (non-routine) occupations.

Perhaps the most striking feature is the “polarizing” nature of this structural break. Despite the fact that the *aggregate* growth contribution during recessions has remained roughly unchanged—or if anything slightly decreased—after 1990, it is about half as large for non-routine occupations during recessions, while it slightly increased for routine occupations. This suggests that the much more substantial aggregate job destruction within routine occupations during recessions—apparent in Figure 1—is likely a structural, cluster-specific phenomenon and not merely a feature of more pronounced aggregate cyclical shocks.

Somewhat surprisingly, we find that the structural break in the cluster-specific median employment growth during expansions was not as “polarizing” as one would expect from visual inspection of Figure 1. In fact, the difference in systematic expansion growth between routine and non-routine occupations slightly shrunk rather than widened after 1990 (see panel B of Figure 3). This implies that the large difference in post-1990 average expansion growth apparent in Figure 1 must largely be driven by specific occupations within the two clusters. One source for this phenomenon can easily be seen in Figure 2: a large portion of employment growth in cluster 2 is explained by idiosyncratic variation. Hence, the observed average growth rate within this occupation group is much higher than predicted by the systematic

at the occupation level. Unsurprisingly, this decomposition reveals that the variance share explained by the common component is largest for occupations with large factor loadings.

Figure 4: Structural Breaks in Estimated Factor Means



Notes: The graphs illustrate kernel density estimates of μ_{k,S_t} based on our sample of posterior draws.

component alone.

Finally, to illustrate uncertainty in the magnitude of the estimated structural break we plot the posterior distributions of the estimated factor- and state-specific means, $\hat{\mu}_{k,S_t}$, in Figure 4. Panels A.1 and B.1 compare the posterior distribution of pre-1990 and post-1990 growth rates for factor 1 during expansions and recessions, respectively. Panels A.2 and B.2 plot the same comparisons for factor 2. Except for the modest change in median recession growth for factor 2, all of the breaks are statistically significant.

In sum, our analysis clearly illustrates that the long run trends documented in the polarization literature (see [Acemoglu and Autor, 2011](#)) have a significant *systematic*, cyclical component that is predominantly concentrated in recessions. This finding confirms that the strongly polarizing aggregate cyclical dynamics of routine and non-routine occupations (panel A of Figure 1) are indeed driven by a systematic feature that is common to the underlying detailed occupations. Nonetheless, a non-negligible portion of the polarizing dynamics, especially during post-1990 expansions, appear to be substantially amplified by idiosyncratic variation within a number of underlying occupations.

4.4. Can the Structural Break Explain Jobless Recoveries?

The last three recoveries in the US were “jobless”. That is, output and other measures of real activity started to recover at the NBER recession trough, while jobs did not. In earlier recessions, employment started recovery within two quarters after the NBER trough while it continued to decline for at least six quarters in the three recessions since

1990 (see the dashed lines in Figures 5 and 6). Moreover, job recovery was very modest thereafter.¹⁶ Jaimovich and Siu (2018) have recently linked this phenomenon to job polarization and we revisit their arguments here.

The results in Section 4.3 clearly show that the *systematic* component of job “polarization”—i.e. the gradual disappearance of routine jobs and the simultaneous rise of non-routine jobs—has become highly cyclical after 1990. Most importantly, routine job destruction in recessions has increased, while routine job creation has completely vanished—in fact, routine jobs systematically continue to slightly decline rather than recover during post-1990 expansions.

Can this highly cyclical, systematic disappearance of routine occupations after 1990 explain why jobs did not recover “as usual” *in the aggregate* after the last three recessions? To address this question we consider the following thought experiment: If state specific average factor growth in routine occupations had remained unchanged after 1990 (i.e., $\mu_{1,3} = \mu_{1,1}$ and $\mu_{1,2} = \mu_{1,4}$), how would per-capita employment have evolved in the aftermath of the last three recessions?

While this thought experiment is similar to the one analyzed by Jaimovich and Siu (2018), ours differs in several respects. First, they consider a structural break in post-1990 average *recovery* growth within routine occupations, where they define recoveries as the 24 months following the official NBER trough. In contrast, we analyze a joint break in the posterior distributions of state (expansion/recession) and cluster (routine/non-routine) specific growth. Since there is both uncertainty about the business cycle state (S_t) and because we allow for autoregressive factor dynamics, our model “endogenously” allows for different systematic growth rates during “recovery periods”. To visualize this aspect of our model, the gray and green areas in panel B of Figure 3 highlight NBER recessions and recovery periods (8 quarters) respectively. This figure clearly shows that recoveries in routine job-growth were significantly slower after 1990, in the sense that it takes almost two years to return to “normal” expansion growth. In contrast, non-routine occupations appear to recover almost instantly, just as they did before 1990. Thus, despite the fact that the median routine/non-routine growth gap during *expansions* (overall) has not seen a significant break, *recoveries* have indeed become polarizing after 1990.

Second, although we have thus far documented significant structural change in the *common* component of routine and non-routine jobs, we also detect a substantial amount of idiosyncratic variation in employment growth (see Figure 2), captured by ε_{it} in our model. Thus, a clean test should control for this idiosyncratic variation to isolate the *net* contribution of structural change. Finally, our empirical framework allows us to draw formal posterior inference about all estimated effects.

To construct our counterfactual, we simulate factor dynamics for the post-1990 period under the assumption that $\tilde{\mu}_{k,3} = \hat{\mu}_{k,1}$ (recessions) and $\tilde{\mu}_{k,4} = \hat{\mu}_{k,2}$ (expansions), where $\hat{\mu}_{k,1}$ and $\hat{\mu}_{k,2}$ denote the estimated pre-1990 factor-specific growth in the respective business cycle states.¹⁷ Based on the resulting hypothetical factor series, $\tilde{f}_{k,t}^c = \tilde{\mu}_{k,S_t} + \hat{\phi}_k \tilde{f}_{k,t-1}^c + \hat{y}_{k,t}$, and conditional on the estimated occupation classification, we then compute two versions of occupation-specific employment growth: one in which we assume that $\varepsilon_{it} = 0$ for all t (“no shocks”),

$$y_{it}^{NS} = \hat{\lambda}_{it} \tilde{f}_{\delta_{i,t}}^c \text{ for all } i, \quad (11)$$

and one in which we postulate that $\varepsilon_{it} = \hat{\varepsilon}_{it}$ (“shocks”), that is

$$y_{it}^S = y_{it}^{NS} + \hat{\varepsilon}_{it} \text{ for all } i, \quad (12)$$

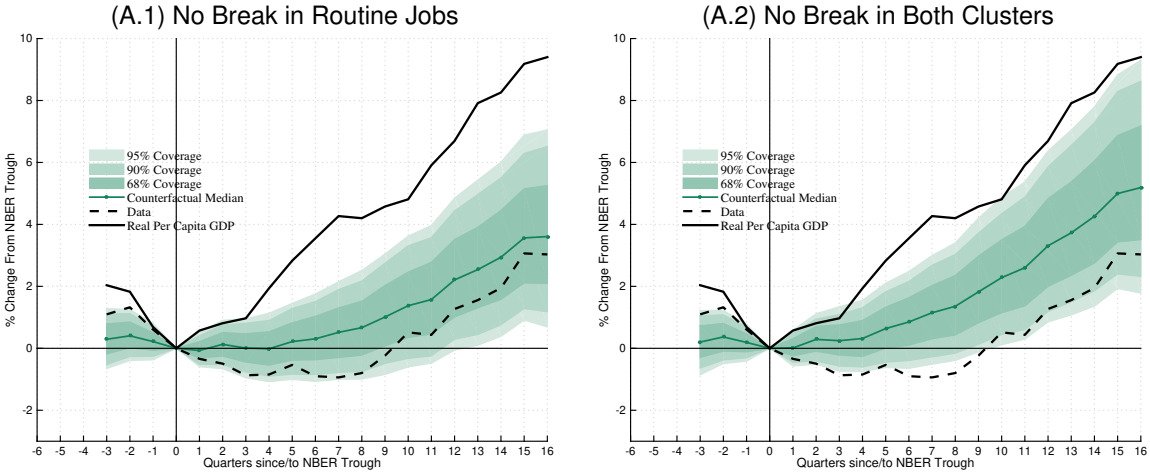
where $\hat{\varepsilon}_{it} = y_{it} - \hat{\lambda}_{it} \hat{f}_{\delta_{i,t}}$ captures the idiosyncratic, occupation-specific variation implied by the estimated factor dynamics. In analogy to the derivations in equations (7) and (8) we then compute the implied level series and sum across all occupations to obtain aggregate level series.

To assess the resulting counterfactual employment dynamics following NBER dated recessions, we simulate three counterfactual paths, respectively starting at the beginning of the 1990/91, the 2001, and the 2008/09 recessions. We then normalize both the counterfactual path as well as the observed data to equal zero at the trough of each respective recession. Since we start our simulated paths at the peak before each respective recession, we also account for the delayed effects of the much more polarizing recessions after 1990, which carry over into the recovery period through the autoregressive structure in the factor dynamics.

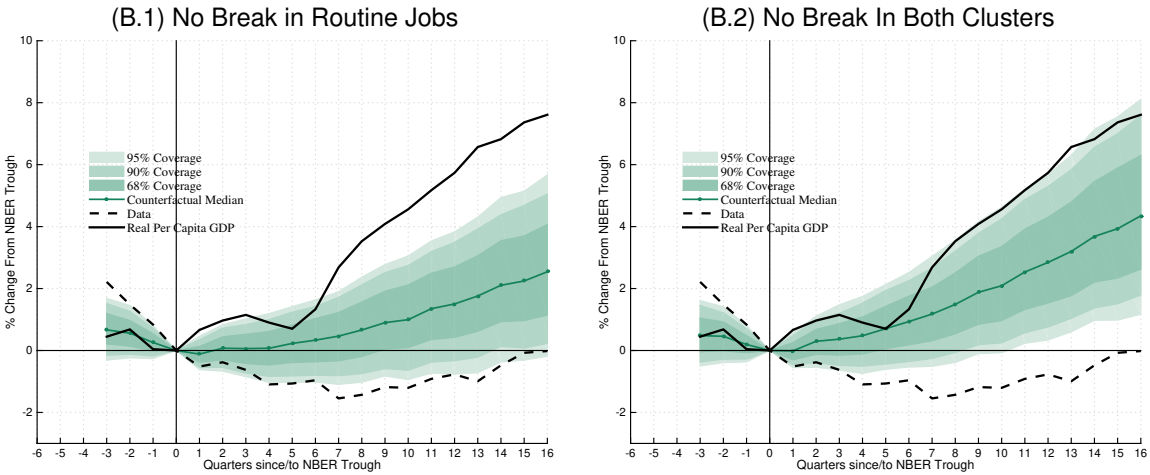
¹⁶For further discussion of this phenomenon see Schreft and Singh (2003), Groshen and Potter (2003), and Jaimovich and Siu (2018)

¹⁷Generally, we denote MCMC estimates with a “hat”. To draw proper inference on the counterfactuals, we compute implied employment levels for each of the MCMC draws $m|\delta^{50}$, and use the resulting empirical posteriors to construct coverage regions. For notational convenience, we omit m henceforth.

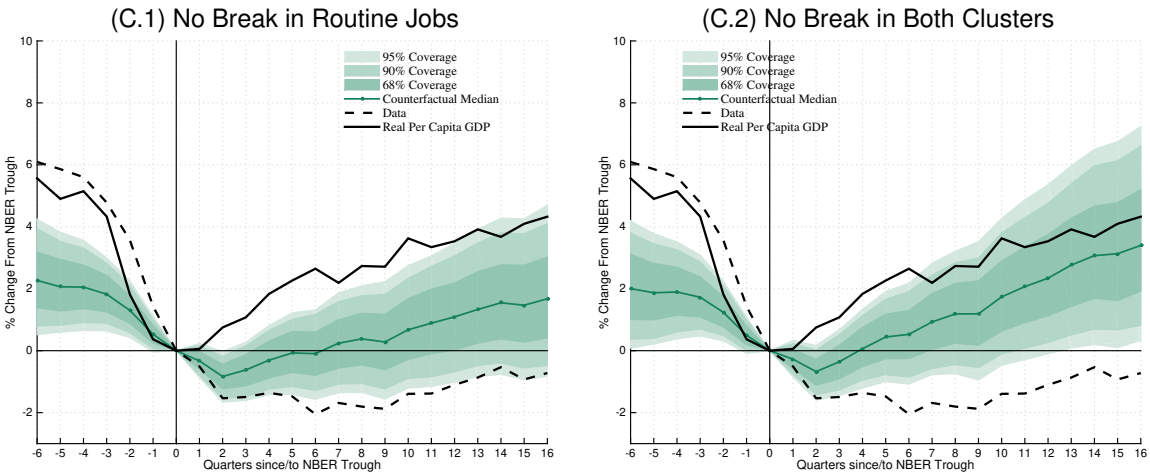
Figure 5: Counterfactual Experiment: Structural Dynamics Only
 A. 1990/91 Recession (Trough: 1991m3)



B. 2001 Recession (Trough: 2001m11)

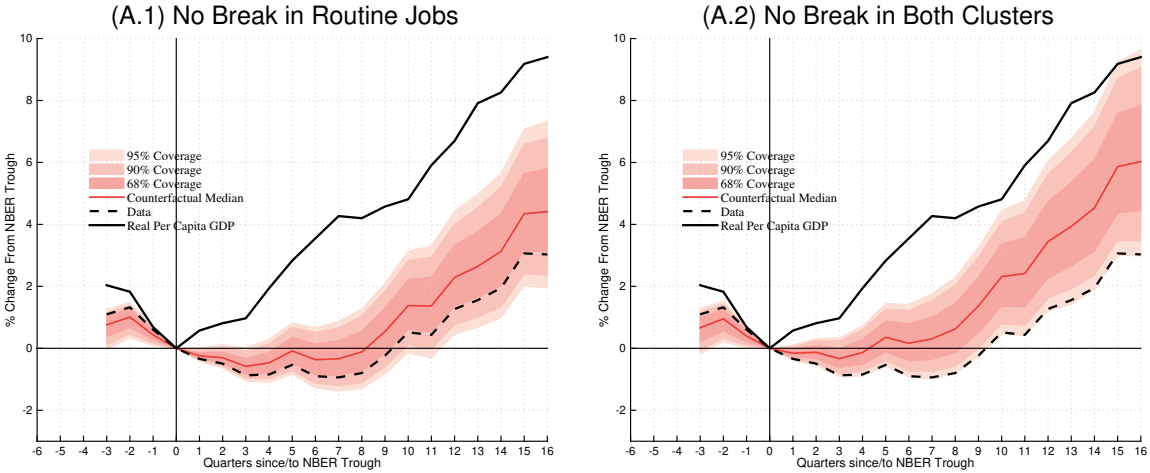


C. 2008/09 Recession (Trough: 2009m6)

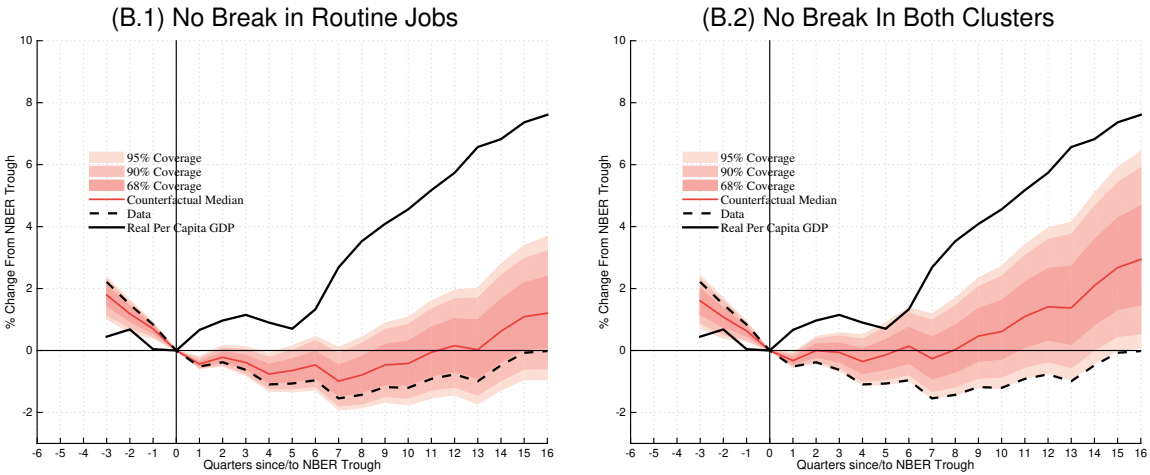


Notes: All graphs illustrate counterfactual experiments in which the factors $f_{k,t}$ are governed by the systematic component of the pre-break dynamics only, i.e. with $\varepsilon_{i,t} = 0$ for all t . Panels A.1, B.1, and C.1 illustrate the counterfactual in which we “undo” the break in routine jobs only, while panels A.2, B.2, and C.2 illustrate the implied recovery paths if we “undo” the break in both factors.

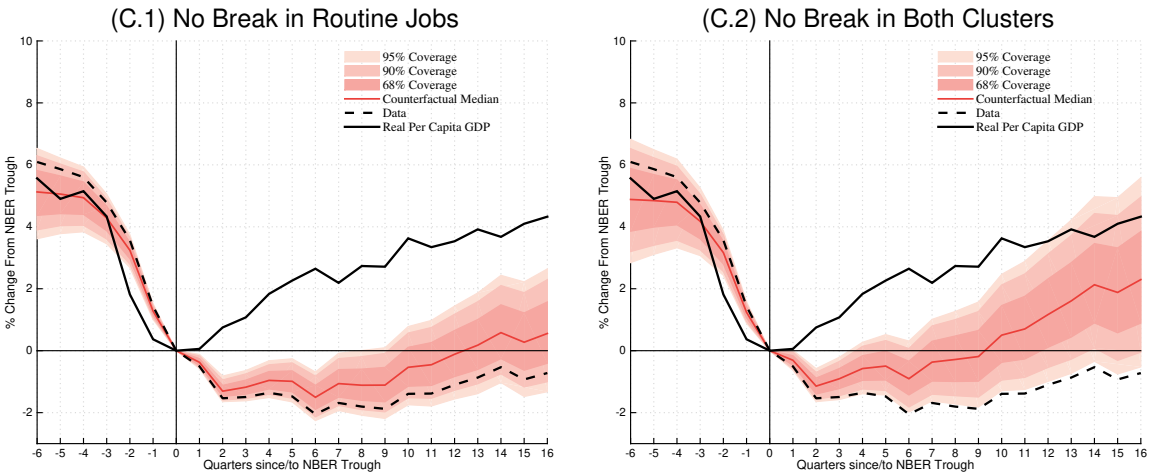
Figure 6: Counterfactual Experiment: Structural & Idiosyncratic Dynamics
 A. 1990/91 Recession (Trough: 1991m3)



B. 2001 Recession (Trough: 2001m11)



C. 2008/09 Recession (Trough: 2009m6)



Notes: All graphs illustrate counterfactual experiments in which the factors $f_{k,t}$ are governed by the systematic component of the pre-break dynamics and we control for occupation specific variation, i.e., $\varepsilon_{i,t} = \hat{\varepsilon}_{it}$. Panels A.1, B.1, and C.1 illustrate the counterfactual in which we “undo” the break in routine jobs only, while panels A.2, B.2, and C.2 illustrate the implied recovery paths if we “undo” the break in both factors.

If the structural change in routine employment dynamics truly played an important role for aggregate recovery dynamics, then we would expect the counterfactual path to significantly diverge from that observed in the data, even when accounting for the observed idiosyncratic variation. Specifically, we would expect the counterfactual paths during recoveries to lie significantly above the data.

Figure 5 illustrates the resulting counterfactual paths under the assumption that employment dynamics are entirely driven by the *systematic* component (i.e., y_{it}^{NS} with $\varepsilon_{it} = 0$). Panels A.1, B.1, and C.1 display the effects of “removing” the break in routine jobs, while panels A.2, B.2, and C.2 illustrate the impact of “undoing” the break in both clusters for each of the three recessions since 1990. This figure highlights a number of interesting results: first, all counterfactual paths indicate a significantly stronger recovery than observed in the data, as the data lie at least outside the 68% coverage region during the first two years of recovery for all counterfactual experiments. Second, in all counterfactuals, employment starts recovery at least after two quarters and returns to its respective trough value in less than two years, even after the 2008/09 recession. Third, as expected, when we remove both breaks we see a slightly stronger recovery than the one predicted by undoing the break in routine jobs only.

However, our estimates suggest that idiosyncratic variation (captured by ε_{it} in our model) is non-negligible for the observed aggregate employment dynamics since 1990. Thus, Figure 6 illustrates an alternative set of counterfactual paths, in which we account for the estimated idiosyncratic variation in the underlying occupation dynamics (based on y_{it}^S with $\varepsilon_{it} = \hat{\varepsilon}_{it}$).

Like in the baseline, all counterfactual paths display significantly stronger recoveries than observed in the data. However, after controlling for idiosyncratic variation, the breaks in the systematic component have a much weaker impact on the speed of aggregate employment recovery. Specifically, when we only undo the break in routine occupations, the counterfactual paths recover to their trough value at the earliest after 8 quarters. Nevertheless, even after the 2008/09 recession, the counterfactual paths predict recovery to the trough value within 13 quarters. When we remove both breaks, for routine and non-routine occupations, we see recovery to the trough value within 5, 6, and 10 quarters after the 1990/91, 2001, 2008/09 recessions, respectively.

5. Conclusions

The primary contribution of this paper is a statistical framework that allows us to disentangle group-specific (common) employment dynamics from occupation-specific (idiosyncratic) ones and simultaneously identify clusters of jobs that share common cyclical patterns. Based on detailed occupation level data from the CPS, we find that our model fits best when occupations are grouped into two clusters that almost perfectly coincide with occupation groups that Autor et al. (2003) label “routine” and “non-routine” jobs, respectively. Moreover, it detects a significant structural break in the cluster-specific dynamics of both routine and non-routine occupations around the 1990/91 recession. We then assess the impact of this structural break in the common group dynamics on employment growth in the three recoveries since 1990. Our main finding is that, in the absence of this structural break, aggregate employment in the US would have recovered significantly more strongly than observed in the data during these “jobless recoveries”.

References

- Acemoglu, D., December 1999. Changes in unemployment and wage inequality: An alternative theory and some evidence. *American Economic Review* 89 (5), 1259–1278.
URL <http://ideas.repec.org/a/aea/aecrev/v89y1999i5p1259-1278.html>
- Acemoglu, D., Autor, D., 2011. Skills, Tasks and Technologies: Implications for Employment and Earnings. Vol. 4 of *Handbook of Labor Economics*. Elsevier, Ch. 12, pp. 1043–1171.
URL <http://ideas.repec.org/h/eee/labchp/5-12.html>
- Akerman, A., Gaarder, I., Mogstad, M., 2015. The Skill Complementarity of Broadband Internet. *The Quarterly Journal of Economics* 130 (4), 1781–1824.
URL <https://ideas.repec.org/a/oup/qjecon/v130y2015i4p1781-1824.html>
- Autor, D. H., 2013. The “task approach” to labor markets : an overview. *Journal for Labour Market Research* 46 (3), 185–199.
URL <http://ideas.repec.org/a/iab/iabjlr/v2013i3p185-199.html>
- Autor, D. H., Dorn, D., August 2013. The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review* 103 (5), 1553–97.
URL <http://ideas.repec.org/a/aea/aecrev/v103y2013i5p1553-97.html>

- Autor, D. H., Katz, L. F., Kearney, M. S., May 2008. Trends in u.s. wage inequality: Revising the revisionists. *The Review of Economics and Statistics* 90 (2), 300–323.
URL <http://ideas.repec.org/a/tpr/restat/v90y2008i2p300-323.html>
- Autor, D. H., Levy, F., Murnane, R. J., November 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118 (4), 1279–1333.
URL <http://ideas.repec.org/a/tpr/qjecon/v118y2003i4p1279-1333.html>
- Bachmann, R., 2012. Understanding the jobless recoveries after 1991 and 2001. Working paper, University of Michigan.
- Berger, D., 2012. Countercyclical Restructuring and Jobless Recoveries. 2012 Meeting Papers 1179, Society for Economic Dynamics.
URL <https://ideas.repec.org/p/red/sed012/1179.html>
- Bernanke, B. S., 2003. The jobless recovery.
URL <https://www.federalreserve.gov/boarddocs/speeches/2003/200311062/default.htm>
- Bernanke, B. S., 2009. On the outlook for the economy and policy.
URL <http://www.federalreserve.gov/newsevents/speech/bernanke20091116a.htm>
- Burger, J. D., Schwartz, J. S., 2018. Jobless recoveries: Stagnation or structural change? *Economic Inquiry* 56 (2), 709–723.
URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/ecin.12535>
- Cortes, G. M., Jaimovich, N., Nekarda, C. J., Siu, H. E., July 2014. The micro and macro of disappearing routine jobs: A flows approach. Working Paper 20307, National Bureau of Economic Research.
URL <http://www.nber.org/papers/w20307>
- Dorn, D., 2009. Essays on inequality, spatial interaction, and the demand for skills. Dissertation 3613, University of St. Gallen.
- Furukawa, Y., Toyoda, H., Jun. 2013. Job polarization and jobless recoveries in Japan: Evidence from 1984 to 2010. KIER Working Papers 874, Kyoto University, Institute of Economic Research.
URL <https://ideas.repec.org/p/kyo/wpaper/874.html>
- Gaggl, P., Wright, G. C., July 2017. A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive. *American Economic Journal: Applied Economics* 9 (3), 262–294.
URL <https://ideas.repec.org/a/aea/aejapp/v9y2017i3p262-94.html>
- Goos, M., Manning, A., February 2007. Lousy and lovely jobs: The rising polarization of work in Britain. *The Review of Economics and Statistics* 89 (1), 118–133.
URL <http://ideas.repec.org/a/tpr/restat/v89y2007i1p118-133.html>
- Goos, M., Manning, A., Salomons, A., May 2009. Job polarization in Europe. *American Economic Review* 99 (2), 58–63.
URL <http://ideas.repec.org/a/aea/aecrev/v99y2009i2p58-63.html>
- Gordon, R. J., Baily, M. N., 1993. The jobless recovery: Does it signal a new era of productivity-led growth? *Brookings Papers on Economic Activity* 1993 (1), pp. 271–316.
URL <http://www.jstor.org/stable/2534606>
- Groshen, E. L., Potter, S., 2003. Has structural change contributed to a jobless recovery? *Current Issues in Economics and Finance* 9 (Aug).
URL <http://ideas.repec.org/a/fip/fednci/y2003iaugnv.9no.8.html>
- Jaimovich, N., Siu, H. E., 2018. Job polarization and jobless recoveries. *The Review of Economics and Statistics*, forthcoming.
- Koenders, K., Rogerson, R., 2005. Organizational dynamics over the business cycle: a view on jobless recoveries. *Review* (Jul), 555–580.
URL <https://ideas.repec.org/a/fip/fedlrv/y2005ijulp555-580nv.87no.4.html>
- Schreft, S. L., Singh, A., 2003. A closer look at jobless recoveries. *Economic Review* (Q II), 45–73.
URL <http://ideas.repec.org/a/fip/fedker/y2003iqiip45-73nv.88no.2.html>