

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

March 7, 2015

Abstract: We analyze quarterly occupation-level data from the US Current Population Survey for 1976-2013. Based on common cyclical employment dynamics, we 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, factor model

Paul Gaggl¹

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

Sylvia Kaufmann¹

Study Center Gerzensee
Dorfstrasse 2, P.O. Box 21
CH-3115 Gerzensee
Switzerland

Email: sylvia.kaufmann@szgerzensee.ch

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

1. Introduction

A number of recent studies document that, 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 \(2012\)](#) use an aggregated mapping of skills into jobs and document that 92% of the decline in “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.³ 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.

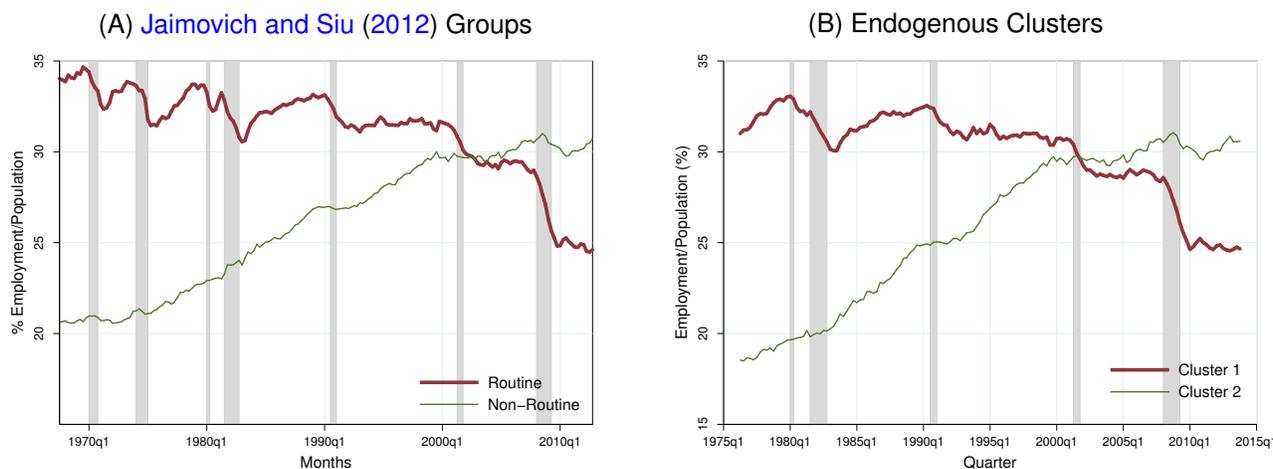
If these stark cyclical patterns are truly due to a distinguishing feature, that is common to occupations within each broad category considered by [Jaimovich and Siu \(2012\)](#), then we should be able to identify this group-specific characteristic as well as potential structural breaks from high frequency 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

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, 2013](#); [Gaggl and Wright, 2014](#)) 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.

³This mapping of aggregate job categories into three broad skill groups is based on the work surveyed in [Acemoglu and Autor \(2011\)](#).

Figure 1: Group Employment Trends



Notes: Panel A plots employment trends as reported in [Jaimovich and Siu \(2012\)](#). 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 3. 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.

formally test [Jaimovich and Siu’s \(2012\)](#) hypothesis that labor market polarization may play an important role in explaining jobless recoveries in the US. 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 employment per-capita, 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 1976-2013. Even though our identification is based entirely on cyclical employment dynamics within 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 metrics (see [Autor, 2013](#), for a detailed discussion of these difficulties). Thus, we find it reassuring that our “agnostic” approach, based on completely separate identifying variation, delivers an almost identical grouping of occupations. In particular, the estimated occupation clusters suggest 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 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*, as emphasized by [Jaimovich and Siu \(2012\)](#). A visual comparison of panels A and B in [Figure 1](#) clearly reveals the similarity in the aggregate dynamics between the two classification schemes.⁴

To capture group-specific business cycle dynamics, we adopt a Markov switching structure that accounts for asymmetric dynamics across expansions and recessions as well as potential breaks in these dynamics. This flexible specification allows us to identify a significant structural break in the group-specific cyclical dynamics around 1990. This result is very much in concert with [Jaimovich and Siu \(2012\)](#). Specifically, we find that systematic routine employment growth during expansions completely vanished (from 0.19% to -0.02% quarterly) while non-routine occupations grew half as fast on average (1% vs. 0.5% quarterly). Second, systematic routine job destruction during recessions almost doubled (from -0.7% to -1.3% quarterly) while non-routine jobs per-capita continued to grow during recessions on average, yet at a slower pace (0.38% vs. 0.27% quarterly).

These estimates allow us to revisit a question first posed by [Jaimovich and Siu \(2012\)](#): can these structural breaks explain the three “jobless recoveries” in the US? Specifically, for each recovery since the 1990 recession, we construct counterfactual paths for employment per-capita within each detailed occupation, under the assumption that the systematic, cluster-specific dynamics had remained unchanged after 1990. As descriptively suggested by [Jaimovich and Siu \(2012\)](#), we find that aggregate employment per-capita would have recovered significantly more strongly in the absence of the observed structural change around 1990. However, while our conclusions align with theirs, 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 last point is especially important in light of the results presented by [Foote and Ryan \(2012\)](#). Contrary to [Jaimovich and Siu \(2012\)](#), these authors argue that large aggregate shocks caused job-

⁴In [Section 3](#) we discuss in detail why the levels in panels A and B of [Figure 1](#) don’t match up. We argue that the information about the dynamics of aggregate employment contained in either aggregation scheme is essentially the same.

less recoveries in the US and that labor market polarization played only a minor role in this context. They reach this conclusion by observing that the variance share explained by industry-skill-specific variation is low compared to the variance share explained by a single unobserved common component. Our results refine this picture. We endogenously identify two distinct clusters of occupations with common cyclical dynamics, which jointly explain at least 60% of the variation in the data. At the same time we account for aggregate shocks through our Markov switching structure and allow for idiosyncratic, occupation-specific shocks. The key difference between the two models is that we allow for cluster specific dynamics, *conditional* on the aggregate state—i.e., recessions and expansions. Thus, even though we find that recessions became “more severe” for labor markers after 1990—in the sense that aggregate net job destruction during recessions was larger than before—we mostly attribute the slow employment recoveries to disproportionately more pronounced structural breaks in the dynamics of “routine” occupations—in one word, polarization.

2. Empirical Model

We specify a model for per-capita employment growth in 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

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

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

$$\phi_k(L) f_{kt} = \mu_k S_t + \nu_{kt}, \nu_{kt} \sim N(0, 1) \quad (3)$$

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

where going from (1) to (2) reflects the fact that the latent classification indicator $\delta_i \in \{1, \dots, K\}$ exclusively assigns one factor to each occupation: $\lambda_{ik} \neq 0$ if $\delta_i = k$ and 0 otherwise, with $k = 1, \dots, K$. Note that, while each occupation only loads on one factor, the factor loadings may vary across occupations.

To capture factor specific dynamics, we assume that each factor follows an $AR(p)$ process, $\phi_k(L) = 1 - \phi_{k1}L - \dots - \phi_{kp}L^p$. Moreover, the factors are affected by independent factor specific “shocks”, ν_{kt} , and by the latent Markov process S_t , which captures the aggregate business cycle phase. Notice further that the autoregressive structure in the factor dynamics introduces persistence in the effects of recessionary episodes, which plays an important role in understand-

ing factor-specific recovery dynamics. The timing of recessions/expansions is fully synchronized across all factors and thus occupations, yet mean, phase-specific employment growth, μ_{k,S_t} , varies across factors. While the factors capture all common dynamics within each occupation group, any idiosyncratic, occupation-specific variation is absorbed 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 the dynamics of routine and non-routine occupations—as defined in the polarization literature (Acemoglu and Autor, 2011)—may have experienced a break around 1990 and we capture this possibility in our specification of the latent state indicator S_t . In each period, the economy can be in one of four states, $S_t \in \{1, \dots, 4\}$, which ex-post have the following interpretation: 1 = pre-break recession, 2 = pre-break expansion, 3 = post-break recession, 4 = post-break expansion.

To estimate the time path of S_t we assume the following structure: first, we postulate that S_t follows a time-varying first-order Markov process with transition probabilities, $\xi_{jl,t} = P(S_t = l | S_{t-1} = j, x_t, t)$, $l, j = 1, \dots, 4$, where the Markov property introduces some state persistence. Second, we condition the state transition probabilities on real GDP growth, x_t , and a time trend, t . The former helps identify business cycle phases while the latter helps identify the break date.⁵ Mostly inspired by panel A of Figure 1, we then impose two sets of restrictions on the transition probabilities, summarized in the following time specific transition 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, we only allow a structural break to happen once. That is, once state 3 or 4 is reached, the economy cannot go back to either 1 or 2. This is implemented by setting the lower-left block of transition probabilities to 0. Second, a potential structural break must happen in transition from a recession to an expansion, or vice versa. For example, we don't allow the economy to switch from a pre-break expansion into a post-break expansion. This is implemented by setting the upper-right diagonal elements of the transition matrix to 0.

⁵See Appendix A for more details on the parametrization of the state transition matrix.

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:

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

To obtain draws from (6), we simulate iteratively from the conditional posterior distributions of:

1. the factors, $\pi(\mathbf{f}|\mathbf{y}, \mathbf{S}, \boldsymbol{\delta}, \theta)$.
2. the business cycle indicator $\pi(\mathbf{S}|\mathbf{f}, \boldsymbol{\xi}, \boldsymbol{\mu}, \boldsymbol{\phi})$.
3. the classification indicator $\pi(\boldsymbol{\delta}|\mathbf{y}, \mathbf{f}, \boldsymbol{\psi}, \boldsymbol{\sigma})$.
4. (a) the parameters of the transition distribution $\pi(\boldsymbol{\gamma}|\mathbf{S}, \mathbf{x}, \mathbf{t})$.
 (b) the remaining parameters $\pi(\theta_{-\boldsymbol{\gamma}}|\mathbf{y}, \mathbf{f}, \mathbf{S}, \boldsymbol{\delta})$, where $\theta_{-\boldsymbol{\gamma}}$ is θ without $\boldsymbol{\gamma}$.

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 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 show in [Appendix B](#) that the imposed structure allows us to sample from a normal posterior and we provide full details on the specification of the likelihood, all priors and the derivation of the posterior in [Appendix B](#).

3. Data

Our main data source are detailed individual level data from the basic US Current Population Survey (CPS) covering the period 1976m1-2013m12. 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 \(2009\)](#).⁶ For our baseline specification, summarized in [Table 1](#), we further aggregate the detailed occupations into 21 groups, as also designed by [Dorn \(2009\)](#). Following [Jaimovich and Siu \(2012\)](#) we use the share of employment within the US population of age 16 and older to obtain each occupation’s labor market dynamics over time.

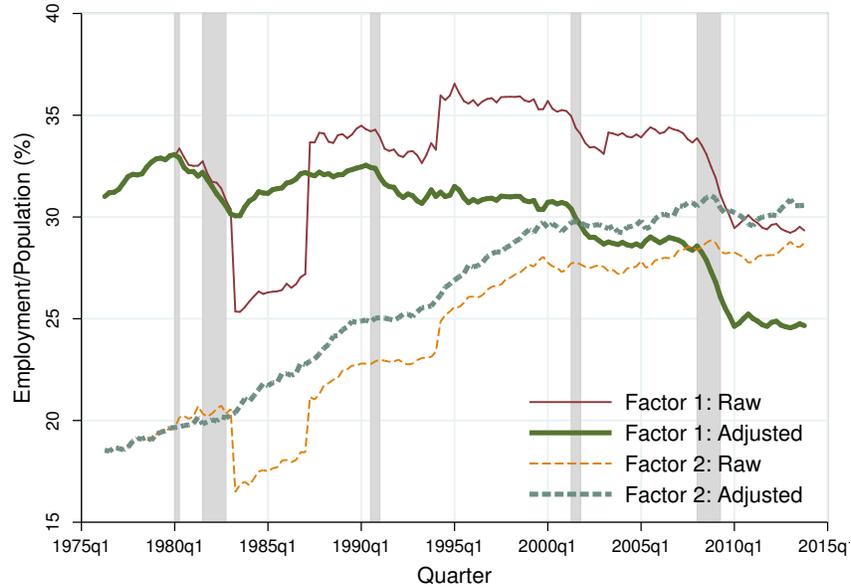
To compare our results to [Jaimovich and Siu \(2012\)](#) we also replicate their dataset, spanning the period 1967-2012, for which data prior to 1983 are taken from the DOL’s *Employment & Earnings* publications and from the Federal Reserve Bank of St. Louis’ FRED database thereafter. These data contain 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. However, consistent with [Jaimovich and Siu \(2012\)](#), we are able to 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). We further seasonally adjust all time series (from both data sources) using the US Census X11 method. Based on this dataset, panel A of [Figure 1](#) displays employment as a share of population for the two groups of non-routine and routine occupations. As expected, this figure resembles [Figure 4](#) in [Jaimovich and Siu \(2012\)](#).

One of the biggest challenges in working with the detailed CPS data are the frequent changes in the DOL’s system for classifying occupations. Even [Dorn’s \(2009\)](#) consistent panel features many jumps in the level of employment since various occupations “jump” from one group to another, new occupations are introduced, or old ones disappear. These jumps are not readily visible in long run comparisons (e.g. across decades) but they become immediately apparent at higher frequencies. To avoid this problem, [Foote and Ryan \(2012\)](#), who also study the cyclicity of labor market polarization, decide to use industry-skill cells as a proxy for jobs/tasks instead of occupations specified by the DOL.

However, since the level jumps are due to purely administrative changes, they always happen in a single month. Therefore, one way to accommodate the level jumps, is to use growth rates instead of levels and average out the jumps for the occupations in which administrative changes happen. [Figure 2](#) shows the levels implied by our adjusted growth rate series. It is obvious that any adjustment procedure introduces some measurement error, but [Figure 1](#) illustrates that the

⁶The 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.

Figure 2: Employment Trends Based on Growth Rates



Notes: The figure illustrates the cumulative growth of employment/population in each occupation assigned to factors 1 and 2 in model (1), which are tabulated in Table 3. The imputed level series start with the level of employment/population in 1976q1 and illustrate the variation in growth rates used in our estimation procedure. The series labeled “raw” are based on the unadjusted growth rates in the occupation level series while those labeled “adjusted” are based on a growth rate series in which the administrative “jumps” were interpolated based on the median January growth (all administrative jumps happen in January). Data for this graph are directly constructed from the monthly basic CPS files for the consistent panel of occupations compiled by Dorn (2009).

dynamic patterns in the level of routine and non-routine jobs implied by these adjusted growth rates is virtually the same as in the level series employed by Jaimovich and Siu (2012). In fact, our approach to adjust in growth rates is very similar in spirit to the “flows approach” of Cortes et al. (2014).

Ultimately, it should be clear that all four approaches, broad aggregation as in Jaimovich and Siu (2012), forming industry-skill cells as in Foote and Ryan (2012), the “flows approach” by Cortes et al. (2014), and our adjustment in growth rates, are an imperfect solution and introduce some form of measurement error. However, given the nature of administrative changes in the DOL’s definition of occupations, these are the best options available.⁷

Finally, we employ real GDP growth as our aggregate measure to help identify business cycles

⁷We obtain the same qualitative results when we estimate our model with 9 occupation groups assembled as in Jaimovich and Siu (2012).

Table 1: Preferred Model Specification

<i>A. Specification</i>		
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
<i>B. Sample</i>		
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
<i>B. Posterior Sampler</i>		
Total Posterior draws		500,000
Burn-in		300000
Retained Observations		50,000 (every fourth draw after burn-in)

and we draw these data from FRED.

4. Empirical Analysis

We estimate model (1) using the Markov Chain Monte Carlo (MCMC) posterior sampler described in [Appendix B](#). 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.⁸ 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, [Table 2](#) 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.

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 more than 30% in 2013.

⁸In [Appendix C](#), we graphically illustrate the retained draws for selected model parameters after conditioning on the specific cluster assignment tabulated in [Table 3](#).

Table 2: Model Selection

q	p	$Var(\hat{Y}_k)/Var(Y_k)$		
		Cluster 1	Cluster 2	Avg.
1	1	0.345	0.353	0.349
1	2	0.286	0.216	0.251
2	1	0.138	0.138	0.138
2	2	0.607	0.588	0.597
2	3	0.306	0.213	0.260
3	2	0.484	0.620	0.552
3	3	0.223	0.199	0.211

Notes: The table reports the fraction of the variation in aggregate employment/population that is explained by common cluster dynamics, conditional on the median factor assignment. Maxima are highlighted.

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 peak in 1980 to about 25% at its trough in 2013. Second, employment in these occupations appears highly “cyclical”. Growth rates obviously differ between recessions and expansions, and there seems to be a change in these growth rates around 1990.

Overall, the groups identified by our preferred model specification resemble the patterns in employment dynamics presented by [Jaimovich and Siu \(2012\)](#), as displayed in panel A of Figure 1.⁹ While the aggregate levels in panels A and B of Figure 1 don’t perfectly match up, the long run growth patterns 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. There are two major reasons for why the aggregate level series don’t match up exactly. First, our consistent panel does not cover all occupations, since some occupations are simply not available consistently over the entire sample period (see [Dorn, 2009](#), for details). Second, our model is based on adjusted growth rates instead of the actual level series. This comes at the expense of losing some information about the “true” level of employment but allows for a straightforward way to accommodate the administrative changes in the occupation classification. However, since the focus of our analysis are the *dynamics* in group-specific employment, this choice does not dramatically affect our inference or our conclusions (see

⁹Note 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 \(2012\)](#), 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 3.

Table 3: Cluster Analysis: Factor Assignment

21 Occupation Groups (Dorn, 2009)	Assignment δ_i		Factor Loading $\lambda_{ik} \delta^{50}$			
	$Pr[\delta_i = 1 \mathbf{y}]$	$Pr[\delta_i = 2 \mathbf{y}]$	Mean	Median	68% Coverage	
<i>A. Factor 1 (Routine)</i>						
F.1 Machine Operators, Assemblers, and Inspectors	0.999	0.001	1.069	1.072	0.950	1.191
E.2 Construction Trades	0.999	0.001	0.883	0.890	0.770	0.991
E.4 Precision Production	0.999	0.001	0.840	0.844	0.750	0.939
F.2 Transportation and Material Moving	1.000	0.000	0.686	0.684	0.617	0.754
E.1 Mechanics and Repairers	1.000	0.000	0.626	0.622	0.552	0.695
B.2 Sales	0.975	0.025	0.466	0.463	0.413	0.516
B.3 Administrative Support	0.760	0.240	0.243	0.250	0.170	0.322
C.31 Food Preparation and Service	0.608	0.392	0.168	0.167	0.100	0.240
<i>B. Factor 2 (Non-Routine)</i>						
E.3 Extractive	0.229	0.771	0.941	0.940	0.544	1.361
A.2 Management Related	0.000	1.000	0.818	0.821	0.729	0.918
C.37 Misc. Personal Care and Service	0.031	0.969	0.800	0.785	0.574	1.013
A.1 Executive, Administrative, and Managerial	0.001	0.999	0.599	0.598	0.532	0.675
C.36 Child Care Workers	0.130	0.870	0.586	0.567	0.411	0.750
C.32 Healthcare Support	0.076	0.924	0.444	0.439	0.359	0.536
A.3 Professional Specialty	0.005	0.995	0.374	0.379	0.329	0.425
C.2 Protective Service	0.335	0.665	0.329	0.325	0.256	0.411
C.33 Building/Grounds Cleaning/Maintenance	0.105	0.895	0.296	0.290	0.198	0.388
C.34 Personal Appearance	0.326	0.674	0.241	0.245	0.128	0.363
B.1 Technicians and Related Support	0.144	0.856	0.239	0.240	0.171	0.307
C.35 Recreation and Hospitality	0.476	0.524	-0.061	-0.061	-0.251	0.142
C.1 Housekeeping and Cleaning	0.418	0.582	-0.257	-0.259	-0.376	-0.141

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 $Pr[\delta_i = 1|\mathbf{y}] > 1/2$ while panel B collects those with $Pr[\delta_i = 2|\mathbf{y}] > 1/2$ 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 $\lambda_{ik}|\delta^{50}$, conditional on the median factor assignment, δ^{50} . Within each panel, the occupations are sorted in decreasing order by their conditional factor loading $\lambda_{ik}|\delta^{50}$.

the strong similarity in the dynamic patterns between panel A and panel B of Figure 1).

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 3 tabulate the posterior assignment probabilities for the two factors. Notice that almost all of the 21 Dorn (2009) occupation groups are nearly perfectly assigned to one of the two clusters. Only a handful of occupations have an assignment probability of less than 2/3. Panels A and B respectively group occupations for which the posterior probability of being determined by factors 1 and 2 is larger than 50%.¹⁰

¹⁰Except for the particular assignment of the few occupations that are not decisively associated with either cluster, this classification is robust across all other model specifications with $K = 2$ factors that we considered. As discussed in the text, the “unassigned” occupations are quantitatively small and show little growth throughout the entire sample.

Notice that there are two service occupations (C.35 and C.1) that essentially have a 50/50 chance of belonging to either factor. It turns out that employment in these occupations is essentially constant (in levels) throughout the entire sample and that their share in total employment is very small. Therefore, these occupations do not contain much information about the factor dynamics, which is also reflected in an insignificant factor loading for the recreation and hospitality group (C.1).

In addition to the assignment probabilities, the last four columns of Table 3 report the posterior mean and median factor loading, $\lambda_{ik}|\delta^{50}$, as well as the associated 68% posterior coverage interval, conditional on the median factor assignment, δ^{50} .¹¹ Notice that all but one of these intervals exclude zero. The occupation with an insignificant factor loading is precisely one of the two service occupations for which the assignment probability is not decisive. If we estimate specifications with $K > 2$, the assignment classification deteriorates considerably for all occupations, which further suggests that the “unassigned” occupations simply experience very idiosyncratic employment dynamics rather than being influenced by an additional factor.

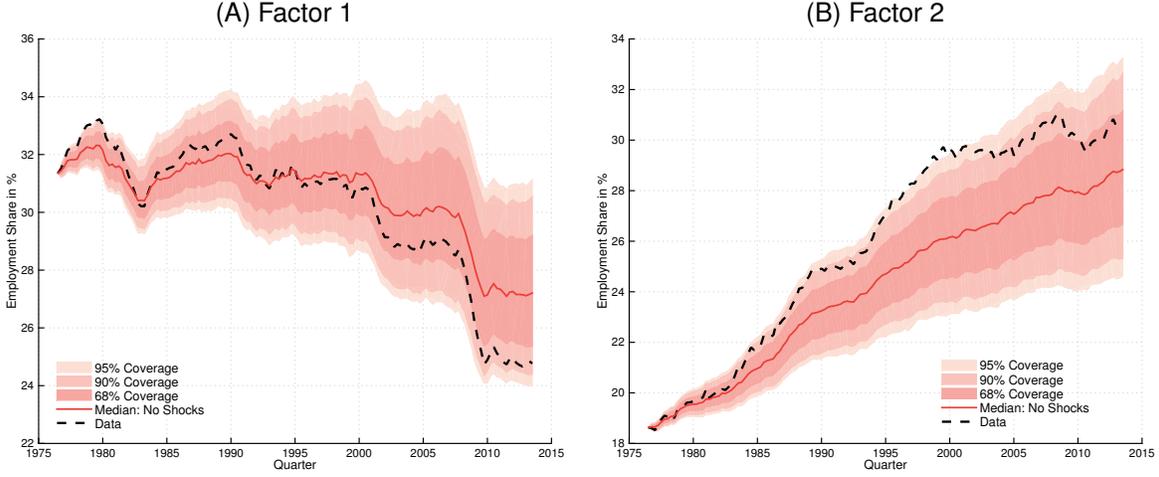
Within panels A and B of Table 3, we sort occupations in decreasing order of the posterior mean/median factor loading. This provides a measure of the “intensity” with which a given occupation is influenced by the common factor dynamics and we observe that a considerable amount of heterogeneity is present 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 3) 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 3 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 service occupations, like “recreation and hospitality” and “housekeeping and cleaning”, are not decisively assigned to either factor, and that one service occupation (“food preparation and service”) is assigned to cluster 1 with a 60% probability. On the other hand, one traditional “blue collar” occupation (“extractive”) is decisively assigned to factor 2. This may very well be driven

Therefore, they are not likely to have much influence on aggregate employment dynamics.

¹¹That is, these moments are computed from joint posterior draws for which $Pr[\delta_i = 2|\mathbf{y}] > 1/2$ and for which each occupation, i , is assigned to one of the two clusters exactly as in Table 3.

Figure 3: Structural Employment Dynamics



Notes: Conditional on $Pr[\hat{\delta}_i | \mathbf{y}] > 0.5$, 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 median factor assignment, δ^{50} .

by a rising demand for highly skilled engineers in areas such as “fracking”, off-shore drilling, and other recent high-tech resource extraction techniques.

4.2. Cluster-Specific Cyclical Dynamics

Conditional on the median factor assignment from Table 3, we now shift our focus to the decomposition of employment dynamics into a “structural” (factor specific) and an “idiosyncratic” (occupation-specific) component. Figure 3 illustrates the cluster-specific employment per capita *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 median posterior factor assignment probability $\delta^{50} = \{\hat{\delta}_1, \dots, \hat{\delta}_N\}$:

$$\tilde{y}_{i,t}^{(m)} = \left(1 + \hat{\lambda}_{i,\hat{\delta}_i}^{(m)} \hat{f}_{\hat{\delta}_i,t}^{(m)}\right) \tilde{y}_{i,t-1}^{(m)} \quad \text{for all draws } m | \delta^{50}, \quad (7)$$

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 employment per capita in occupation i . Based on these occupation-specific implied level series we then aggregate

across occupations within each cluster:

$$\hat{Y}_{kt}^{(m)} = \sum_{i|\hat{\delta}_i=k} \tilde{y}_{i,t}^{(m)} \quad \text{for all draws } m|\delta^{50} \text{ and } k = 1, 2, \quad (8)$$

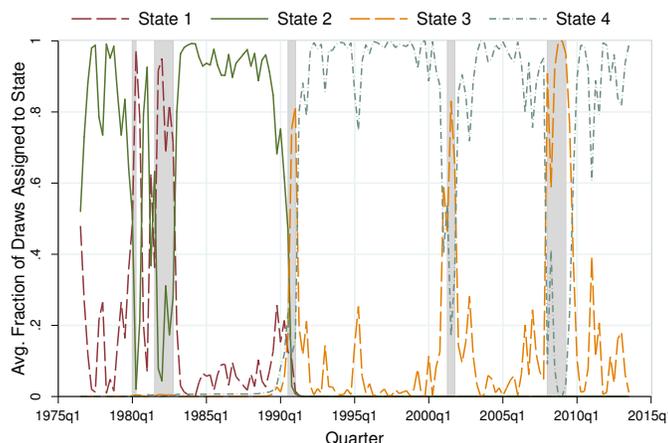
These implied cluster-specific level series capture the predictive ability of the common factor dynamics, while the deviations from observed employment levels are explained by the accumulation of the occupation-specific idiosyncratic component, ε_{it} . We utilize the quantiles of the empirical posterior distribution of the cluster-specific level series to compute posterior coverage regions.¹²

Figure 3 shows that, for factor one, the data lie within the 95% posterior coverage region almost throughout the entire sample. This implies that the common dynamics of occupations in cluster 1 capture a large portion of the aggregate dynamics in employment. In fact, Table 2 reveals that 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 3). Nevertheless, the explained variance share of the structural component in cluster 2 amounts to about 59% (see Table 2). Finally, it is worth emphasizing that Figure 3 is a very conservative test of our model. Recall that we estimate our model in growth rates, yet the fit in levels (based on a single starting value for each occupation) is remarkable. Therefore, as suggested by Jaimovich and Siu (2012), a structural break in the common component of employment dynamics has the potential to significantly affect aggregate employment dynamics.

Figure 3 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 (see Acemoglu and Autor, 2011) has documented common long-run polarizing trends in these groups, 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 3), 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 occu-

¹²Note that, while inference for this simulation is conditional on the median factor assignment, δ^{50} , we fully account for conditional uncertainty in all other parameters.

Figure 4: Markov Switching



Notes: The figure 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).

pations. 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. 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} .

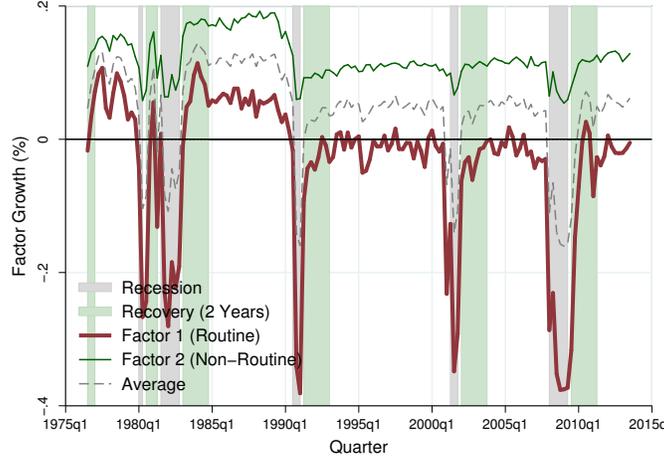
4.3. A Structural Break in the Systematic Cluster Dynamics

Inspired by the patterns apparent in panel A of Figure 1, Jaimovich and Siu (2012) recently hypothesized 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 they then 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 such a break at the end of the long-lasting expansion during the 1980s (see Figure 4).

Figure 4 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 employment per capita and variation in

Figure 5: Polarizing Dynamics



Notes: The figure illustrates the posterior median of the systematic, factor-specific aggregate growth contribution $\tilde{F}_{t,k}^{(m)}$, as computed in equation (10), for the two factors $k \in \{1, 2\}$. The dashed line is the average of the systematic factor dynamics.

real GDP growth. In particular, GDP growth mainly helps to identify expansions and recessions, while the factor dynamics—inferred from dynamics in occupation level employment per capita—identify the structural break. Specifically, Figure 4 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, Figure 5 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

$$\tilde{f}_{t,k}^{(m)} = \hat{f}_{k,t}^{(m)} - \hat{\nu}_{k,t} = \hat{\mu}_{k,S_t}^{(m)} + \sum_{j=1}^p \hat{\phi}_{k,j}^{(m)} \hat{f}_{k,t-j}^{(m)} \quad \text{for all draws } m | \delta^{50} \text{ and } k \in \{1, 2\} \quad (9)$$

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, Figure 5 plots the posterior median of the weighted average

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

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 .

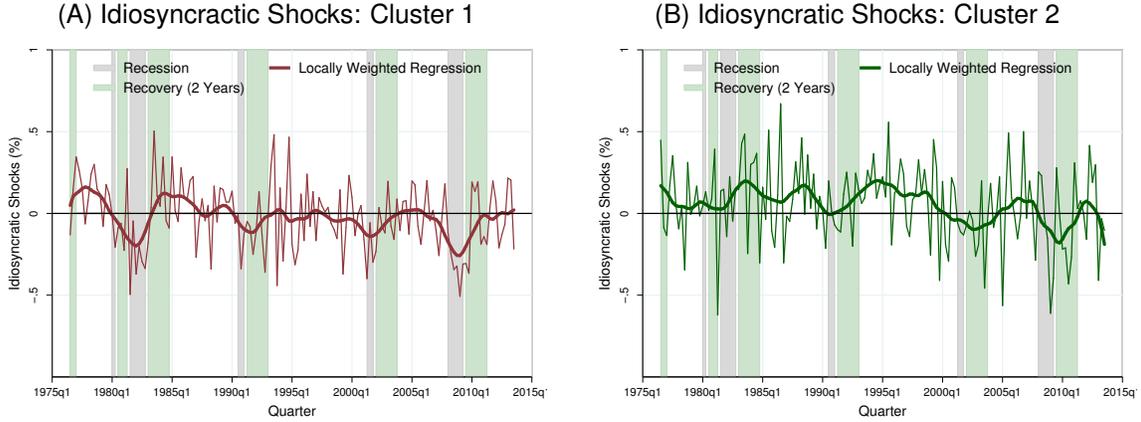
This figure reveals several interesting aspects of the systematic cyclical dynamics within each occupation cluster: first, both clusters feature substantial systematic cyclical variation, which is not immediately visible in Figure 1. Second, both clusters experience a structural break in the dynamics of recessions and expansions. Specifically, we find that systematic routine employment growth, during expansions completely vanished (from $\hat{\mu}_{1,2} = 0.19\%$ to $\hat{\mu}_{1,4} = -0.02\%$ quarterly) while non-routine occupations grew half as fast on average ($\hat{\mu}_{2,2} = 1\%$ vs. $\hat{\mu}_{2,4} = 0.5\%$ quarterly). Most strikingly however, systematic routine job destruction during recessions almost doubled (from $\hat{\mu}_{1,1} = -0.7\%$ to $\hat{\mu}_{1,3} = -1.3\%$ quarterly) while non-routine jobs per-capita continued to grow systematically during recessions, yet at a slower pace ($\hat{\mu}_{2,1} = 0.38\%$ vs. $\hat{\mu}_{2,3} = 0.27\%$ quarterly). Figure 5 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 gap during recessions has remained roughly unchanged—or if anything slightly decreased—after 1990, this gap 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. This finding confirms the results by [Jaimovich and Siu \(2012\)](#) while it contrasts [Foote and Ryan \(2012\)](#). Moreover, it further reinforces the importance of grouping individuals by occupation, rather than industry or industry-skill cells (as suggested by [Foote and Ryan, 2012](#)), when trying to understand the potential structural changes in labor market dynamics over the past 30 years.

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—or if anything remained roughly the same—rather than widened after 1990 (see Figure 5). 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 3: 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 component alone.

Figure 6: Idiosyncratic Dynamics



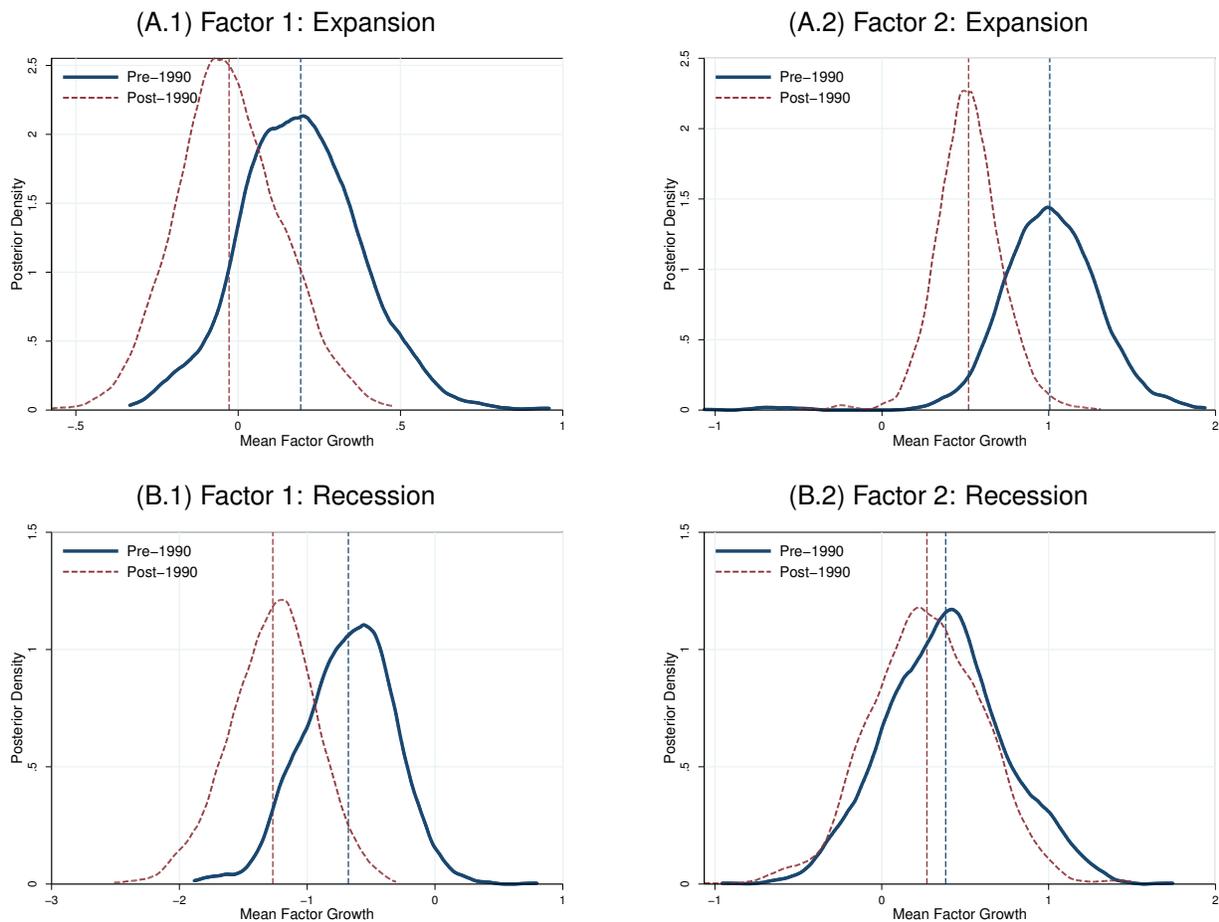
Notes: Panels A and B illustrate the estimated aggregate growth contribution of occupation specific idiosyncratic shocks (posterior median) within each cluster, $\hat{\varepsilon}_{kt} = \sum_{i|\hat{\delta}_i=k} w_{i,t} \hat{\varepsilon}_{i,t}$ for $k \in \{1, 2\}$, and where $w_{i,t}$ is the relative size of occupation i at time t . To visually illustrate the low frequency component in the idiosyncratic shocks we also report a locally weighted regression (thick).

To illustrate this point more precisely, Figure 6 depicts the estimated idiosyncratic variation, aggregated over all occupations according the median posterior assignment probability within each cluster, $\hat{\varepsilon}_{k,t} = \sum_{i|\hat{\delta}_i=k} w_{i,t} \hat{\varepsilon}_{i,t}$ for $k \in \{1, 2\}$, and where $w_{i,t}$ is the relative size of occupation i at time t .¹³ Panel B makes clear that idiosyncratic variation had a clear positive contribution to aggregate post-break growth within non-routine occupations, especially during the 1990s expansion. In a similar vein, idiosyncratic variation within routine occupations had a non-negligible influence on pre-1990 expansion growth, especially in the immediate aftermath of recessions. Taken together, these two observations make clear why the break in the systematic component of expansion growth does not appear to be polarizing, despite the large differential in observed aggregate post-1990 expansion growth (see Figure 1).

In sum, Figures 5 through 6 clearly illustrate 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, as emphasized by Jaimovich and Siu (2012), 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

¹³To visually illustrate the low frequency component in the idiosyncratic shocks we also report a locally weighted regression (thick line).

Figure 7: 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.

during post-1990 expansions, appear to be substantially amplified by idiosyncratic variation within a number of underlying occupations.

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 7. 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.

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 8 and 9). Moreover, job recovery was very modest thereafter.¹⁴ Jaimovich and Siu (2012) 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” 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 employment per capita have evolved in the aftermath of the last three recessions?

While this thought experiment is similar to the one analyzed by Jaimovich and Siu (2012), ours differs in several respects. First, Jaimovich and Siu (2012) 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 state of nature (S_t) and because we allow for autoregressive factor dynamics, our model “endogenously” allows for different systematic growth rates during recovery periods, as defined by Jaimovich and Siu (2012). To visualize this aspect of our model, the gray and green areas in Figure 5 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*

¹⁴For further discussion of this phenomenon see Schreft and Singh (2003), Groshen and Potter (2003), and Jaimovich and Siu (2012)

component of routine and non-routine jobs, we also detect a substantial amount of idiosyncratic variation in employment growth (see Figures 3 and 6), 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 $\tilde{\mu}_{1,3} = \hat{\mu}_{1,1}$ (recessions) and $\tilde{\mu}_{1,4} = \hat{\mu}_{1,2}$ (expansions), where $\hat{\mu}_{1,1}$ and $\hat{\mu}_{1,2}$ denote the estimated pre-1990 factor-specific growth in the respective business cycle states.¹⁵ Based on the resulting hypothetical factor series, $\tilde{f}_{k,t}^c = \tilde{\mu}_{k,S_t} + \hat{\phi}_k \tilde{f}_{k,t-1}^c + \hat{\nu}_{k,t}$, 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”),

$$y_{it}^{NS} = \hat{\lambda}_{it} \tilde{f}_{\delta_i,t}^c \text{ for all } i,$$

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

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

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.¹⁶ Analogous 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.

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

¹⁵Generally, we denote all MCMC estimates for the pre and post-break period 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 the ease of notation we omit m for the rest of this section and we ask the reader to keep this in mind.

¹⁶The cluster-specific (size-weighted) aggregate of these “occupation-specific shocks” is illustrated in panels A and B of Figure 6.

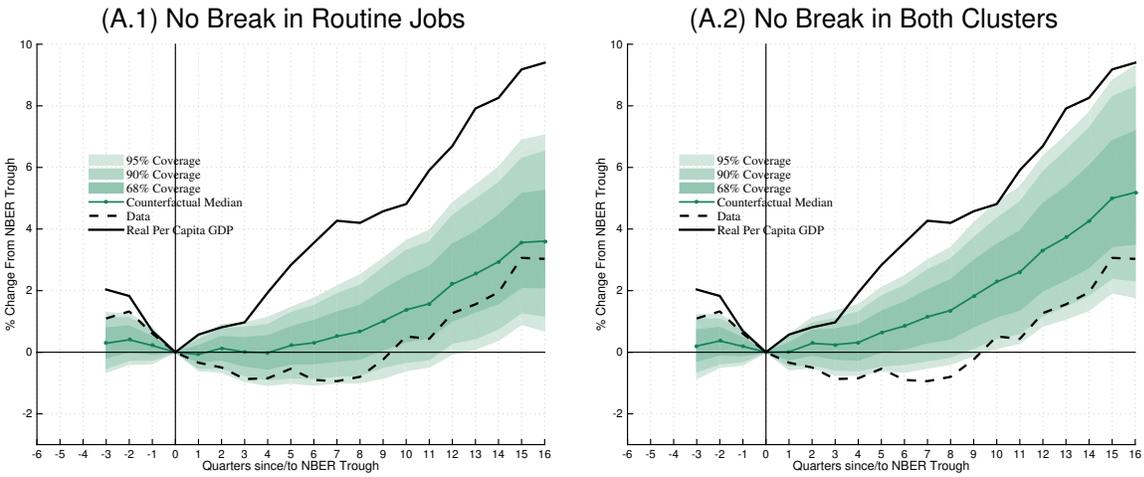
from that observed in the data, even when accounting for the observed idiosyncratic variation. Specifically, we would expect the counterfactual paths during recoveries to lie significantly above the data. Note that these baseline counterfactual paths assume that non-routine jobs *did* experience the break in both recession and expansion growth. Thus, in a final set of counterfactuals we “undo” this break as well, and we expect an even stronger recovery than in the baseline counterfactual experiment.

Figure 8 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 recovery period 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. These results are generally in line with [Jaimovich and Siu \(2012\)](#) and suggest that a “polarizing” structural break in the dynamics of routine (and non-routine) occupations explains a large portion of the slow employment recoveries since 1990.

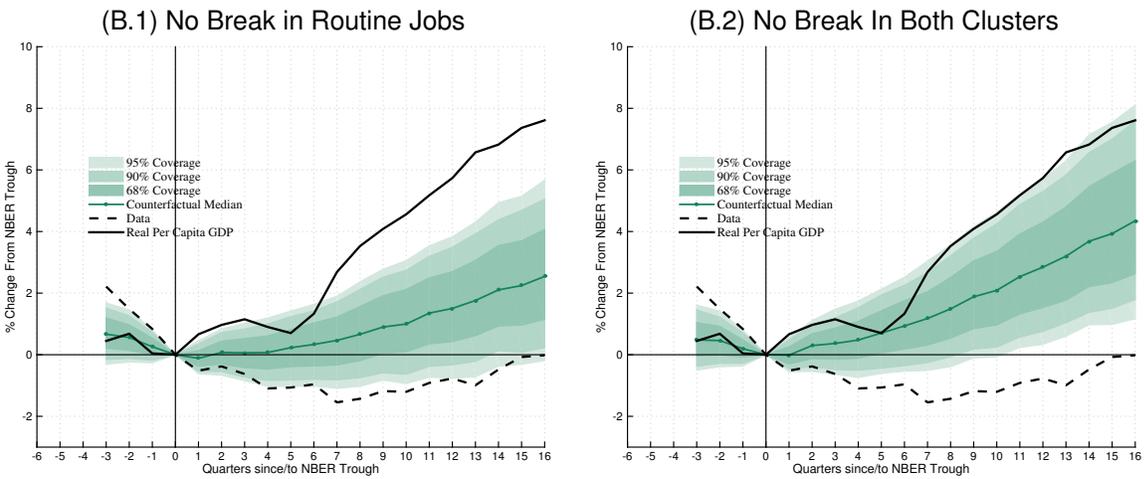
However, our estimates suggest that idiosyncratic variation (captured by ε_{it} in our model) is non-negligible for the observed aggregate employment dynamics since 1990 (see Figure 6). Thus, Figure 9 illustrates the counterfactual paths when we account for the estimated idiosyncratic variation in the underlying occupation dynamics (based on y_{it}^S with $\varepsilon_{it} = \hat{\varepsilon}_{it}$). Like in the baseline counterfactual, 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—an event yet to be seen in the data. 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.

Figure 8: Counterfactual Experiment: Structural Dynamics Only

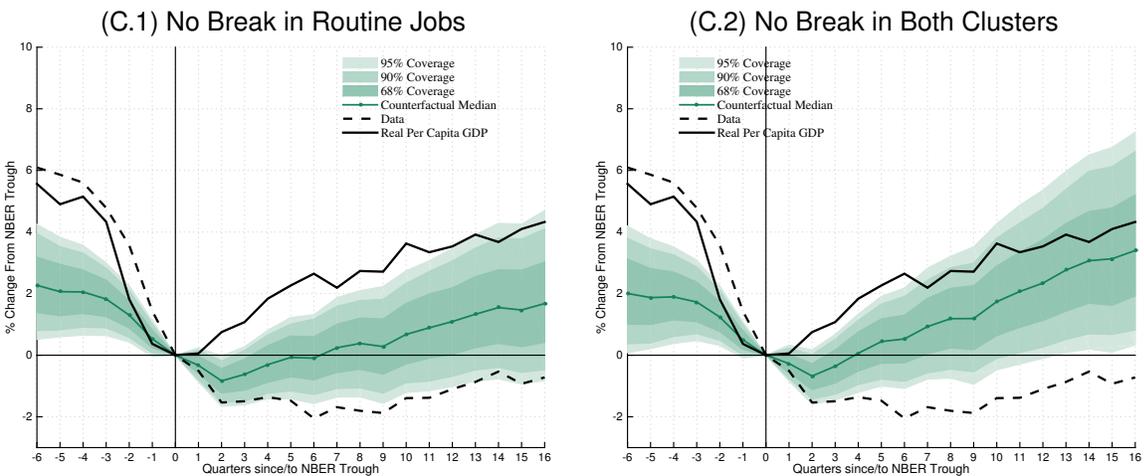
A. 1990/91 Recession (Trough: 1991m3)



B. 2001 Recession (Trough: 2001m11)



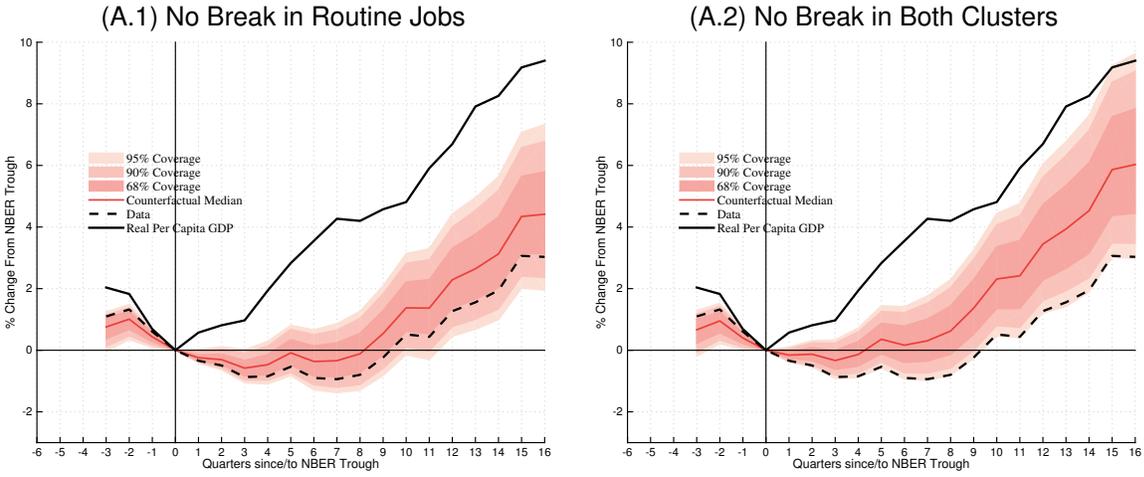
C. 2008/09 Recession (Trough: 2009m6)



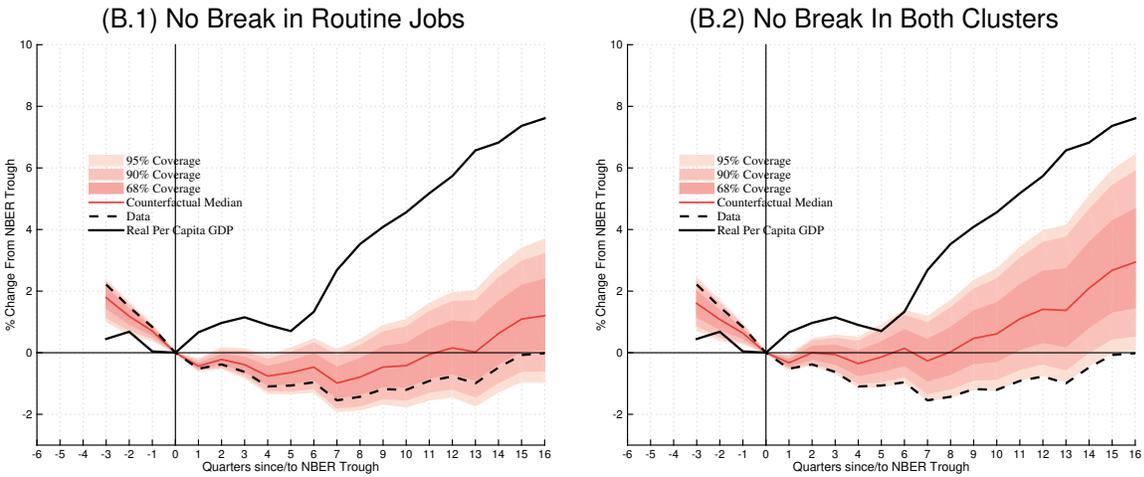
Notes: All graphs illustrate counterfactual experiments in which the factors $f_{k,t}$ are governed by the systematic component of the pre-break dynamics only, i.e. with $\varepsilon_{i,t} = 0$ for all t . Panels A.1, B.1, and C.1 illustrate the counterfactual in which we “undo” the break in routine 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 9: Counterfactual Experiment: Structural & Idiosyncratic Dynamics

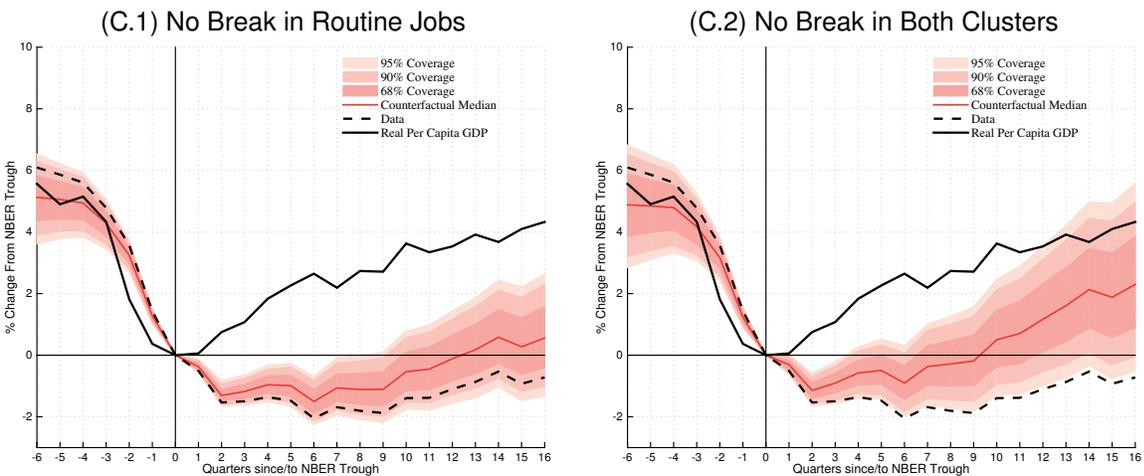
A. 1990/91 Recession (Trough: 1991m3)



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.

5. Polarization and Jobless Recoveries in Developed Countries

Our analysis for the US provides evidence for a systematic connection between job polarization and jobless recoveries in the US. However, while polarization has been documented for a number of developed countries (see for example [Goos et al., 2009](#); [Michaels, Natraj and van Reenen, 2014](#)) we know very little about its relationship to jobless recoveries outside the US. In this section we provide some suggestive evidence for a more global connection between job polarization and jobless recoveries.

We start with documenting the timing of polarization for 16 developed countries using the EU KLEMS database.¹⁷ This dataset provides annual employment shares for low-, middle-, and high-skilled individuals for a number of developed countries over the period 1980-2005. Since detailed occupation level data that are comparable to the US CPS are not readily available for other countries we use these three broad skill groups to proxy the employment share of routine and non-routine occupations, respectively. Specifically, we define middle-skill occupations as routine while we consider low- and high-skill occupations as non-routine.¹⁸ As a first step, we look at the time path of the routine employment share and determine its peak for each country.¹⁹ It turns out that 12 out of the 16 countries in the dataset observe a peak in the routine employment share between 1980 and 2004. The remaining four countries did not reach a turning point in routine employment prior to 2005, the end of our sample.

The last column in [Table 4](#) reports these peak dates. The first two columns of this table report the average annual growth rate in the routine and non-routine employment share after this peak. For the four countries without a decisive turning point in the routine employment share we report the annual growth rate in 2005, the final year in our sample. Two important observations are worth noting: first, notice that ten out of the twelve countries with turning points display clear polarization after the turning point—that is, the routine employment share is declining and the non-routine employment share is rising, on average. Second, there is a fair amount of cross-country heterogeneity in terms of the “polarization start”. If polarization is ultimately caused by the “computer revolution” (or by increased off-shoring activities) then this timing heterogeneity simply suggests that some countries were faster to adopt these new technologies (or business practices) than others. Notice further that the EU KLEMS data suggest that the US labor market started polarizing in 1994 while our detailed occupation level analysis suggests a significant structural

¹⁷All data employed here are publicly available at <http://www.euklems.net>.

¹⁸This is the same mapping between tasks and skills as used in previous studies (see for example [Goos et al., 2009](#); [Michaels et al., 2014](#)).

¹⁹Figures [D.14](#) through [D.16](#) in [Appendix D](#) illustrate these time paths and also show the identified turning points in routine employment.

Table 4: Job Polarization in Developed Countries

Country	Avg. Growth in Employment Share (%)		Polarization Start
	Routine (Middle Skill)	Non-Routine (Low+High Skill)	
Italy	-0.31	2.30	1991
United States	-0.56	0.86	1994
Germany	-0.52	0.90	1997
South Korea	-0.69	0.56	1994
Netherlands	-0.17	0.81	1993
United Kingdom	-0.26	0.60	1998
Hungary	-0.20	0.37	1999
Japan	-0.19	0.37	2003
Denmark	-0.13	0.23	2003
Slovenia	-0.10	0.16	2004
Austria	0.05	-0.11	1997
Czech Republic	0.14	-0.56	2004
Belgium	1.33	-1.70	
Finland	1.81	-1.51	
Australia	4.17	-2.59	
Spain	6.96	-3.15	

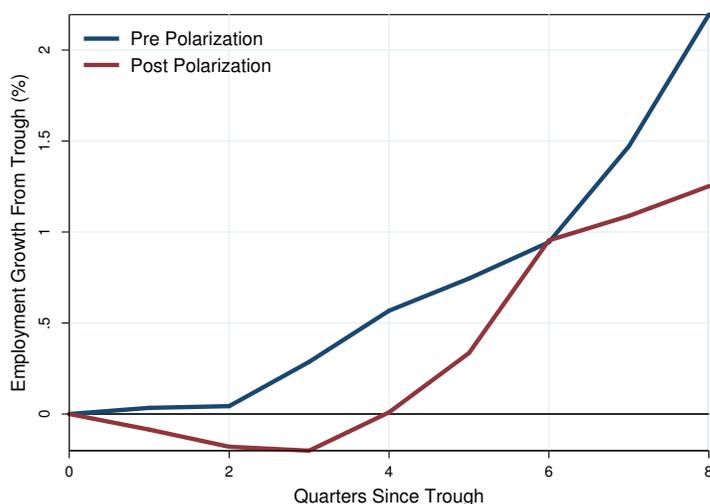
Notes: The table reports the average annual growth rate in employment share within routine (middle skill) and non-routine (high+low skill) occupations since the country started “polarizing”. The polarization start year is defined as the peak in the log routine employment share. All data for this table are taken from the EU KLEMS database (available at <http://www.euklems.net>) and span the years 1980-2005 for countries with the longest time series.

break in employment dynamics around 1990. Given the substantial difference in the approximation of the mapping from jobs/skills into tasks we consider this timing evidence broadly consistent. In sum, Table 4 suggests that the majority of the 16 countries in the EU KLEMS sample started to experience job polarization at some point after 1990, which is in line with the existing evidence (see for example [Goos et al., 2009](#); [Michaels et al., 2014](#)).

As a second step, we provide suggestive evidence for the hypothesis that polarization and jobless recoveries are systematically related. To accomplish this, we use quarterly data on real GDP and employment provided by the OECD to relate the paths of employment and output growth during recoveries. Unfortunately, there are no “official” business cycle dates—like provided by the NBER in the US—for the remaining countries listed in Table 4. To define recessions, we therefore use the dates provided by [Berge \(2012\)](#), who uses a combination of real GDP, industrial production, and the unemployment rate to date the turning points in economic activity in a way that is consistent across countries.²⁰ Equipped with this consistent set of turning points we then define a “recovery”

²⁰Note that his dating procedure matches the NBER dates almost perfectly. Figures D.14 through D.16 in Appendix

Figure 10: Job Recoveries in Developed Countries (Excl. US)



Notes: The figure plots the average employment growth from the recession trough during the eight quarters following a recession trough, both before and after the start of polarization for the countries listed in Table 4. The pre- and post-polarization cutoff for each country is the fourth quarter of the year listed in the last column of Table 4. The countries included are the ones for which the OECD provides sufficient quarterly data on both employment and real GDP: Austria, Germany, Denmark, the UK, Italy, Japan, South Korea, and the Netherlands.

as the eight quarters of expansion immediately after the business cycle trough, consistent with our analysis of jobless recoveries in the previous sections.²¹

In analogy to the analysis of Section 4.4, Figure 10 plots the average employment recovery before and after the start of polarization for the countries listed in Table 4, excluding the US.²² The start of the post polarization era is the fourth quarter of the year listed in the last column of Table 4. This figure provides suggestive evidence that, like in the US, employment recoveries in the pre-polarization era were markedly stronger than in the post-polarization era.

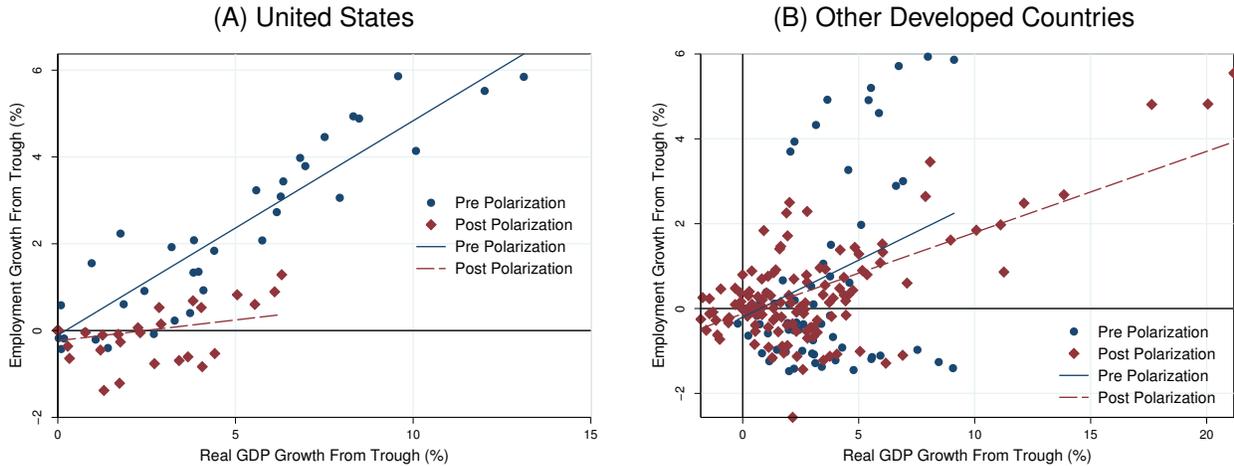
As a final piece of suggestive evidence we explore the relationship between employment and output growth during recoveries. This is an important question to ask, as a “jobless recovery” requires that output recovers but jobs do not. It is therefore possible that the paths of employment recovery in Figure 10 may be misleading if output recovery was also systematically slower in the

D illustrate these recessions and also plot the available data on quarterly real GDP and employment provided by the OECD.

²¹In case there is a recession in less than eight quarters following the previous recession we cut the recovery period at the peak of the business cycle.

²²The average employment growth since each recession trough in Figure 10 is taken across recoveries and countries.

Figure 11: Jobless Recoveries and Polarization



Notes: Panel A relates employment growth to real GDP growth relative to the recession trough for the US. The sample is split into the period before and after the polarization start (1990 for the US). The two lines are fitted regression lines to illustrate the relationship between employment and output recovery. Panel B repeats this exercise for countries with a peak in routine employment before 2005 as listed in Table 4. The pre- and post-polarization cutoff for each country is the fourth quarter of the year listed in the last column of Table 4. The countries included are the ones for which the OECD provides sufficient quarterly data on both employment and real GDP: Austria, Germany, Denmark, the UK, Italy, Japan, South Korea, and the Netherlands. Note that, for consistency, all recession dates (including the US) are taken from Berge (2012).

post-polarization era.²³ To assess this possibility, Figure 11 plots employment growth since the last recession trough against real GDP growth during the same period, both for the US (panel A) and the remaining countries (panel B). Panel A clearly illustrates that employment recoveries were much “slower”, on average, relative to output recoveries after the start of polarization.²⁴ Panel B suggests similar patterns for other developed economies, yet the slowdown in recoveries does not seem quite as pronounced as in the US.

Thus, despite the fact that we were not able to replicate our detailed occupation-level analysis from the US in other countries, we take the evidence provided in this section as suggestive of the idea that jobless recoveries and job polarization may indeed be systematically related on a more global scale.

²³It is worth noting at this point, that Berge’s (2012) turning points take into account the unemployment rate. This means that the recession trough may be dated slightly later than if he had only considered real GDP and industrial production. Therefore, our suggestive evidence provided here is on the more conservative side since the turning point does not necessarily mark the turning point in output, but a combination of the turning point in output and the unemployment rate.

²⁴Notice that the analysis in Figure 11 is in principle also prone to bias since the dynamics of recovery might have changed over time. For example, it is possible that post-polarization employment dynamics are systematically more sluggish at first but recover more rapidly after say one year. Then comparing employment vs. GDP growth after say eight quarters would be the more relevant statistic. However, since our goal is to present some basic suggestive evidence we do not further investigate this possibility here.

6. Conclusions

We provide 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, we find a significant structural break in the cluster-specific dynamics of both routine and non-routine occupations around the 1990/91 recession. Motivated by [Jaimovich and Siu \(2012\)](#), we then assess the impact of this structural break in the common group dynamics on employment growth in the three recoveries since 1990. We find 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”. While we cannot conduct our detailed analysis for other countries we provide some suggestive evidence that this connection between job polarization and jobless recoveries is likely present in other developed countries who have recently started to experience job polarization.

To appreciate the implications of these results we conclude with a discussion of the potential mechanisms that may lead to this relationship. The most common explanation for polarization are massive efficiency gains in information and communication technology (ICT), which have been documented to directly complement non-routine workers (e.g., [Akerman et al., 2013](#); [Gaggl and Wright, 2014](#); [Michaels et al., 2014](#)) and to substitute for routine workers (e.g., [Eden and Gaggl, 2014](#); [Gaggl and Wright, 2014](#)). A second, and likely interrelated factor, is the rising importance of offshoring and international trade (e.g., [Autor, Dorn and Hanson, 2013](#)) in the production of goods and services. Both phenomena have eventually lead developed nations to gradually specialize in production tasks that use non-routine labor more intensively, which is reflected in a rising expenditure share on non-routine labor ([Eden and Gaggl, 2014](#)).

Taken together, these tendencies may be interpreted as “routine biased technological change” (RBTC). One way to measure this form of technological change is to look at the evolution of the price for ICT capital. For example, the relative price of ICT capital (relative to the price of output) in the US was roughly constant up until around 1980 and has fallen precipitously since then (see for example Figure 6 in [Eden and Gaggl, 2014](#)), which suggests continuous efficiency gains in ICT over the past three decades.

But if ICT prices started falling around 1980, why do we see a polarizing structural break around 1990? While this is still an open question, at least two contributions provide theoretical explanations for a causal relationship between ICT driven RBTC and the stark cyclical employment

dynamics observed in Figure 1. Both of these theoretical explanations have one common thread: while the efficiency gains in ICT gradually increase the incentive to restructure, recessions amplify this incentive and cause “lumpy” labor re-allocation. Jaimovich and Siu (2012) provide a model in which the incentive for workers to “retrain” from a middle-skill worker to a high-skill worker gradually rises due to RBTC. In their framework, aggregate productivity shocks amplify this incentive and encourage workers to “lumpily” switch from middle-skill to high-skill employment during recessions.²⁵ While there is no capital in their model, and the connection to ICT is therefore somewhat loose, Morin (2014) presents a model in which firms may employ two types of labor (routine and non-routine) and two types of capital (ICT and non-ICT).²⁶ The two key ingredients in his model are that firms are facing labor adjustment costs (for both routine and non-routine workers) and that ICT is a substitute for routine workers. As the price of ICT precipitously falls, firms will gradually choose to substitute computers for routine workers. Like in Jaimovich and Siu’s (2012) model, this incentive is particularly high during recessions. Specifically, if the economy is hit by a temporary recessionary shock, it is optimal to temporarily reduce all factor inputs and re-hire these factors after the shock has vanished. Since ICT and routine labor are substitutes and re-hiring workers is more costly than re-purchasing ICT, the firm will choose to predominantly invest in ICT instead of routine labor during the recovery. In addition, the firm has an incentive to fire fewer non-routine workers (relative to routine workers) in a first place, since it anticipates re-hiring them due to their complementarity with ICT.

In sum, these theories are consistent with the interpretation that the structural break identified here is likely caused by routine biased technological change (RBTC). This further suggests that our counterfactual exercises point to an important structural component of the recent jobless recoveries in the US. Finally, even though evidence on the connection between polarization and jobless recoveries is thus far weak outside the US, it is conceivable that this phenomenon may become globally relevant as more and more countries start to systematically substitute ICT for routine tasks.

References

- Acemoglu D. 1999. Changes in unemployment and wage inequality: An alternative theory and some evidence. *American Economic Review* **89**: 1259–1278.
URL <http://ideas.repec.org/a/aea/aecrev/v89y1999i5p1259-1278.html>

²⁵Note that this lumpy reallocation is mutually beneficial to workers and firms.

²⁶Note that similar production structures have also been used by Krusell, Ohanian, Rios-Rull and Violante (2000), Autor and Dorn (2013), or Eden and Gaggl (2014).

- Acemoglu D, Autor D. 2011. *Skills, Tasks and Technologies: Implications for Employment and Earnings*, volume 4 of *Handbook of Labor Economics*, chapter 12. Elsevier, 1043–1171.
URL <http://ideas.repec.org/h/eee/labchp/5-12.html>
- Akerman A, Gaarder I, Mogstad M. 2013. The Skill Complementarity of Broadband Internet. IZA Discussion Papers 7762, Institute for the Study of Labor (IZA).
URL <http://ideas.repec.org/p/iza/izadps/dp7762.html>
- Autor D, Dorn D, Hanson GH. 2013. Untangling Trade and Technology: Evidence from Local Labor Markets. IZA Discussion Papers 7329, Institute for the Study of Labor (IZA).
URL <http://ideas.repec.org/p/iza/izadps/dp7329.html>
- Autor DH. 2013. The “task approach” to labor markets : an overview. *Journal for Labour Market Research* **46**: 185–199.
URL <http://ideas.repec.org/a/iab/iabjlr/v2013i3p185-199.html>
- Autor DH, Dorn D. 2013. The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review* **103**: 1553–97.
URL <http://ideas.repec.org/a/aea/aecrev/v103y2013i5p1553-97.html>
- Autor DH, Katz LF, Kearney MS. 2008. Trends in u.s. wage inequality: Revising the revisionists. *The Review of Economics and Statistics* **90**: 300–323.
URL <http://ideas.repec.org/a/tpr/restat/v90y2008i2p300-323.html>
- Autor DH, Levy F, Murnane RJ. 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* **118**: 1279–1333.
URL <http://ideas.repec.org/a/tpr/qjecon/v118y2003i4p1279-1333.html>
- Berge TJ. 2012. Has globalization increased the synchronicity of international business cycles? *Economic Review* **2012(Q3)**.
URL <http://ideas.repec.org/a/fip/fedker/y2012iqiinv.97no.3x3.html>
- Chan JC, Jeliaskov I. 2009. Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation* **1**: 101–120.

- Chib S. 1996. Calculating posterior distributions and modal estimates in markov mixture models. *Journal of Econometrics* **75**: 79–97.
URL <http://ideas.repec.org/a/eee/econom/v75y1996i1p79-97.html>
- Cortes GM, Jaimovich N, Nekarda CJ, Siu HE. 2014. The micro and macro of disappearing routine jobs: A flows approach. Unpublished manuscript, Duke Univeristy.
- Dorn D. 2009. Essays on inequality, spatial interaction, and the demand for skills. Dissertation 3613, University of St. Gallen.
- Eden M, Gaggl P. 2014. The substitution of ICT capital for routine labor: Transitional dynamics and long-run implications. SSRN Scholarly Paper ID 2432313, Social Science Research Network, Rochester, NY.
URL <http://papers.ssrn.com/abstract=2432313>
- Foot CL, Ryan RW. 2012. Labor-market polarization over the business cycle. Public Policy Discussion Paper 12-8, Federal Reserve Bank of Boston.
URL <http://ideas.repec.org/p/fip/fedbpp/12-8.html>
- Frühwirth-Schnatter S. 2010. *Finite Mixture and Markov Switching Models*. Springer Series in Statistics. Springer. ISBN 9781441921949.
URL <http://books.google.com/books?id=WvJ5cgAACAAJ>
- Frühwirth-Schnatter S, Frühwirth R. 2010. Data augmentation and mcmc for binary and multinomial logit models. In Kneib T, Tutz G (eds.) *Statistical Modelling and Regression Structures*. Physica-Verlag HD. ISBN 978-3-7908-2412-4, 111–132.
URL http://dx.doi.org/10.1007/978-3-7908-2413-1_7
- Gaggl P, Wright GC. 2014. A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive. Economics Discussion Papers 752, University of Essex, Department of Economics.
URL <http://ideas.repec.org/p/esx/essedp/752.html>
- Goos M, Manning A. 2007. Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics* **89**: 118–133.
URL <http://ideas.repec.org/a/tpr/restat/v89y2007i1p118-133.html>
- Goos M, Manning A, Salomons A. 2009. Job polarization in europe. *American Economic Review* **99**: 58–63.
URL <http://ideas.repec.org/a/aea/aecrev/v99y2009i2p58-63.html>

Groschen EL, Potter S. 2003. Has structural change contributed to a jobless recovery? *Current Issues in Economics and Finance* **9**.

URL <http://ideas.repec.org/a/fip/fednci/y2003iaugnv.9no.8.html>

Jaimovich N, Siu HE. 2012. The trend is the cycle: Job polarization and jobless recoveries. Working Paper 18334, National Bureau of Economic Research.

URL <http://www.nber.org/papers/w18334>

Kaufmann S. 2014. K-state switching models with time-varying transition distributions – Does credit growth signal stronger effects of variables on inflation? Working Papers 14.04, *re-submitted*, Swiss National Bank, Study Center Gerzensee.

URL <http://ideas.repec.org/p/szg/worpap/1404.html>

Krusell P, Ohanian LE, Rios-Rull JV, Violante GL. 2000. Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica* **68**: 1029–1054.

URL <http://ideas.repec.org/a/ecm/emetrp/v68y2000i5p1029-1054.html>

Michaels G, Natraj A, van Reenen J. 2014. Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years. *The Review of Economics and Statistics* **96**: 60–77.

URL <http://ideas.repec.org/a/tpr/restat/v96y2014i1p60-77.html>

Morin M. 2014. Computer adoption and the changing labor market. Unpublished Manuscript (Job Market Paper), Columbia University.

URL http://www.columbia.edu/~mm3509/Miguel_Morin_Jobless_Recoveries.pdf

Schreft SL, Singh A. 2003. A closer look at jobless recoveries. *Economic Review* : 45–73.

URL <http://ideas.repec.org/a/fip/fedker/y2003iqiip45-73nv.88no.2.html>

Appendix A. Parametrization of the transition distribution

Normalizing the transition to state 1 as the reference state, i.e. $\gamma_{j1,\cdot} = 0$, $j = 1, 2$, the explicit parametrization of (5) is:

$$\xi_t = \begin{bmatrix} \frac{1}{1 + \sum_{s=\{2,4\}} \exp(X'_t \gamma_{1s})} & \frac{\exp(\gamma_{12,0} + \gamma_{12,1}x_t + \gamma_{12,2}t)}{1 + \sum_{s=\{2,4\}} \exp(X'_t \gamma_{1s})} & 0 & \frac{\exp(\gamma_{14,2}t)}{1 + \sum_{s=\{2,4\}} \exp(X'_t \gamma_{1s})} \\ \frac{1}{1 + \sum_{s=\{2,3\}} \exp(X'_t \gamma_{2s})} & \frac{\exp(\gamma_{22,0} + \gamma_{22,1}x_t + \gamma_{22,2}t)}{1 + \sum_{s=\{2,3\}} \exp(X'_t \gamma_{2s})} & \frac{\exp(\gamma_{23,2}t)}{1 + \sum_{s=\{2,3\}} \exp(X'_t \gamma_{2s})} & 0 \\ 0 & 0 & \frac{\exp(\gamma_{33,0} + \gamma_{33,1}x_t)}{\sum_{s=3}^4 \exp(\gamma_{3s,0} + \gamma_{3s,1}x_t)} & \frac{\exp(\gamma_{34,0} + \gamma_{34,1}x_t)}{\sum_{s=3}^4 \exp(\gamma_{3s,0} + \gamma_{3s,1}x_t)} \\ 0 & 0 & \frac{\exp(\gamma_{43,0} + \gamma_{43,1}x_t)}{\sum_{s=3}^4 \exp(\gamma_{4s,0} + \gamma_{4s,1}x_t)} & \frac{\exp(\gamma_{44,0} + \gamma_{44,1}x_t)}{\sum_{s=3}^4 \exp(\gamma_{4s,0} + \gamma_{4s,1}x_t)} \end{bmatrix} \quad (\text{A.1})$$

The vector $X'_t = (1, x_t, t)$ contains a constant, GDP growth (x_t) and a time trend t . GDP growth provides additional information for transitions into business cycle phases (recovery or recession), and time t introduces prior information on the break date. The parameters $\gamma_{jl,m}$, with $j, l = 1, \dots, 4$, and $m = 0, 1, 2$, correspond to the state-dependent, state-specific effects of the variables in X_t . So, $\gamma_{jl,m}$ represents either the constant transition ($m = 0$) effect, the effect of GDP growth ($m = 1$) or the trend effect ($m = 2$) on the transition probability to switch from state j to state l . The denominators are written in a general form, but note that appropriate elements of γ_{14} and γ_{23} are restricted to zero.

Time enters the transition distribution of states 1 and 2, to include prior information on the break date around 1990. We normalize t to be zero in the third quarter of 1990, which corresponds to the peak of the 1980s expansion. The effect of time t should be decreasing for $\xi_{j2,t}$, $j = 1, 2$ and increasing for $\xi_{14,t}$ and $\xi_{23,t}$. Therefore, we expect to estimate $(\gamma_{12,2}, \gamma_{22,2}) \leq 0$ and $(\gamma_{14,2}, \gamma_{23,2}) > 0$. These expectations can be included as information into the prior distribution. In the empirical application, we are less informative and set $\pi(\gamma_{12,2}, \gamma_{22,2}) = N(0, 0.16 \cdot I_2)$ and $\pi(\gamma_{14,2}, \gamma_{23,2}) = N(1, 0.16 \cdot I_2)$, i.e. we do not truncate the distributions.

Normalizing t to zero in the third quarter of 1990 favors a break after the expansion of the 1980s into the early 1990s recession. We explicitly make this choice based on the stylized patterns in Figure 1. Our prime interest is to identify the existence, magnitude, and potential *effects* of a structural break around 1990, rather than the timing of the break itself. The current specification provides a convenient framework to conduct posterior inference on the structural component of employment dynamics before and after 1990. Nevertheless, the framework is general enough to conduct inference on the break date itself in future research.

Appendix B. Bayesian Setup and Estimation

To expose the setup in a concise way, we cast model (1)-(4) into a condensed state-space framework (see also [Chan and Jeliazkov \(2009\)](#)). First, we stack all filtered units in period t into the $N \times 1$ vector y_t^* :

$$\Psi(L)y_t = y_t^* = \lambda f_t - \lambda \odot (\psi_{\cdot 1} \otimes \mathbf{1}_{1 \times k}) f_{t-1} - \dots - \lambda \odot (\psi_{\cdot q} \otimes \mathbf{1}_{1 \times k}) f_{t-q} + \varepsilon_t \quad (\text{B.1})$$

$$\varepsilon_t \sim N(0, \Sigma_\varepsilon), \Sigma_\varepsilon \text{ diagonal}$$

$$f_t = \mu_{S_t} + \Phi_1 f_{t-1} + \dots + \Phi_p f_{t-p} + \eta_t, \quad \eta_t \sim N(0, I_k) \quad (\text{B.2})$$

where \odot and \otimes represent the Hadamar and the Kronecker product, respectively. The $N \times 1$ vector $\psi_{\cdot j}$, $j = 1, \dots, q$, stacks the coefficient at lag j of the idiosyncratic dynamics (see (4)) of all units. The $1 \times k$ row vector $\mathbf{1}_{1 \times k}$ is filled with 1s. The $K \times K$ matrices Φ_j are diagonal and each row of the $N \times K$ matrix λ , λ_i , contains only one non-zero element, i.e. $\lambda_{ik} \neq 0$ if $\delta_i = k$ and $\lambda_{ik} = 0$ otherwise. We stack all observations to obtain the matrix representation:

$$\mathbf{y}^* = \mathbf{\Lambda} \mathbf{f} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(0, I_{T-q} \otimes \Sigma_\varepsilon) \quad (\text{B.3})$$

$$\mathbf{\Phi} \mathbf{f} = \boldsymbol{\mu} + \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim N(0, \boldsymbol{\Omega}) \quad (\text{B.4})$$

where $\mathbf{y}^* = (y_{q+1}^{*'}', \dots, y_T^{*'}')$ and $\mathbf{f} = (f'_{q+1-\max(p,q)}, \dots, f'_{q+1}, \dots, f'_T)'$ stacks all unobserved factors, including initial states. The matrices $\mathbf{\Lambda}$ and $\mathbf{\Phi}$ are, respectively, of dimension $(T-q)N \times (T+d)k$ and square $(T+d)k$, with $d = (p-q)I_{\{p>q\}}$. Typically, these matrices are sparse and band-diagonal:

$$\mathbf{\Lambda} = \begin{bmatrix} \mathbf{0}_{(T-q)N \times dk} & -\lambda \odot (\psi_{\cdot q} \otimes \mathbf{1}_{1 \times k}) & \dots & \lambda & 0 \dots & 0 \\ & & \ddots & \ddots & \ddots & \vdots \\ & 0 \dots & 0 & -\lambda \odot (\psi_{\cdot q} \otimes \mathbf{1}_{1 \times k}) & \dots & \lambda \end{bmatrix}$$

$$\mathbf{\Phi} = \begin{bmatrix} I_p \otimes I_k & 0 & \dots \\ -\Phi_p & \dots & -\Phi_1 & I_k & 0 & \dots \\ & & & & \ddots & \\ & \dots & 0 & -\Phi_p & \dots & -\Phi_1 & I_k \end{bmatrix}, \quad \boldsymbol{\Omega} = \begin{bmatrix} I_p \otimes \Sigma_\eta^0 & 0 & \dots \\ 0 & & \\ \vdots & I_{T+d-p} \otimes I_k & \end{bmatrix}$$

where Σ_η^0 represents the variance of the initial states (see below). The vector $\boldsymbol{\mu}$ includes the state-dependent intercept,

$$\boldsymbol{\mu}^* = \left[\mathbf{0}_{1 \times \max(p,q)k}, \mu'_{S_{q+1}}, \dots, \mu'_{S_T} \right]'$$

Appendix B.1. Likelihood

Given the representation in (B.3)-(B.4), the complete data likelihood has a normal distribution:

$$L(\mathbf{y}^* | \mathbf{f}, \mathbf{S}, \boldsymbol{\delta}, \theta) \sim N(\boldsymbol{\Lambda} \mathbf{f}, I_{T-q} \otimes \Sigma_\varepsilon) \quad (\text{B.5})$$

Appendix B.2. Prior Distributions

For the unobserved factors, from (B.4) we obtain the normal prior distribution:

$$\begin{aligned} \mathbf{f} | \mathbf{S}, \theta &\sim N(\mathbf{f}_0, F_0^{-1}) \\ \mathbf{f}_0 &= \boldsymbol{\Phi}^{-1} \boldsymbol{\mu}^*, \quad F_0 = \boldsymbol{\Phi}' \boldsymbol{\Omega}^{-1} \boldsymbol{\Phi} \end{aligned} \quad (\text{B.6})$$

In $\boldsymbol{\Omega}$, the variance of the initial states, Σ_η^0 , may be chosen to be diffuse. Here, we will choose Σ_η^0 to be a multiple of the identity matrix, $\Sigma_\eta^0 = \kappa I_k$.

The prior for the unobserved state indicator factorizes into

$$\pi(\mathbf{S} | \boldsymbol{\xi}) = \prod_{t=q+1}^T \pi(S_t | S^{t-1}, \boldsymbol{\xi}_t) \pi(S_q)$$

where the initial state distribution $\pi(S_q)$ is assumed to be uniform across state 1 and 2, $P(S_q = s) = 0.5$, for $s = 1, 2$ and $P(S_q = s) = 0$ for $s = 3, 4$.

The prior for the classification indicator is assumed to be uniform discrete, $P(\delta_i = k) = 1/K, \forall i$.

To complete the model, we assume that the parameters are block-independent a priori, $\pi(\theta) = \pi(\boldsymbol{\lambda} | \boldsymbol{\delta}) \pi(\boldsymbol{\psi}) \pi(\boldsymbol{\phi}) \pi(\boldsymbol{\mu}) \pi(\boldsymbol{\sigma}) \pi(\boldsymbol{\gamma})$, with standard distributions:

1. $\pi(\boldsymbol{\lambda} | \boldsymbol{\delta}) = \prod_{i=1}^N \pi(\lambda_i \delta_i) = \prod_{i=1}^N N(\mathbf{l}_0, \mathbf{L}_0)$
2. $\pi(\boldsymbol{\psi}) = \prod_{i=1}^N \pi(\psi_{i1}, \dots, \psi_{iq}) = \prod_{i=1}^N N(q_0, Q_0) I_{\{Z(\psi_i) > 1\}}$
where $I_{\{\cdot\}}$ is the indicator function and $Z(\varphi) > 1$ means that the characteristic roots of the process $\varphi(L)$ lie outside the unit circle.
3. $\pi(\boldsymbol{\phi}) = \prod_{k=1}^K \pi(\phi_{k1}, \dots, \phi_{kp}) = \prod_{k=1}^K N(p_0, P_0) I_{\{Z(\phi_k) > 1\}}$
4. $\pi(\boldsymbol{\mu}) = \prod_{k=1}^K \pi(\mu_{k1}, \dots, \mu_{k4}) = \prod_{k=1}^K N(m_0, M_0)$
5. $\pi(\boldsymbol{\sigma}) = \pi(\sigma_1^2, \dots, \sigma_N^2) = \prod_{i=1}^N IG(\mathbf{e}_0, \mathbf{E}_0)$

6. $\pi(\boldsymbol{\gamma}) = \prod_{s=2}^4 \pi(\boldsymbol{\gamma}_s) = \prod_{s=2}^4 N(g_{0s}, G_{0s})$
 where $\boldsymbol{\gamma}_s = (\gamma'_{1s}, \dots, \gamma'_{4s})'$, $s = 2, \dots, 4$, and g_{0s} and G_{0s} have appropriate dimensions.

Appendix B.3. Posterior Distributions

Combining the prior with the likelihood, we obtain the posterior for

1. the factors, $\pi(\mathbf{f}|\mathbf{y}^*, \mathbf{S}, \boldsymbol{\delta}, \theta) = N(\mathbf{f}, F^{-1})$, with $F = F_0 + \boldsymbol{\Lambda}'(I_{T-q} \otimes \Sigma_\varepsilon^{-1}) \boldsymbol{\Lambda}$ and $\mathbf{f} = F^{-1}(\boldsymbol{\Lambda}'(I_{T-q} \otimes \Sigma_\varepsilon^{-1}) \mathbf{y}^* + F_0 \mathbf{f}_0)$.
 To avoid the full inversion of F we take the Cholesky decomposition, $F = LL'$, then $F^{-1} = L^{-1}L^{-1'}$. We obtain a draw \mathbf{f} by setting $\mathbf{f} = \mathbf{f} + L^{-1}\boldsymbol{\nu}$, where $\boldsymbol{\nu}$ is a $(T+d)k$ vector of independent draws from the standard normal distribution.
2. the state indicator, $\pi(\mathbf{S}|\mathbf{f}, \boldsymbol{\xi}, \boldsymbol{\mu}, \phi)$. To obtain a draw, we adapt the forward-filtering, backward-sampling procedure described in [Chib \(1996\)](#) to the time-varying Markov structure, see also [Frühwirth-Schnatter \(2010\)](#), Algorithm 11.1 and 11.2.
3. the classification indicator, $\pi(\boldsymbol{\delta}|\mathbf{y}, \mathbf{f}, \boldsymbol{\psi}, \boldsymbol{\sigma}) = \prod_{i=1}^N \pi(\delta_i|\mathbf{y}_i, \mathbf{f}, \boldsymbol{\psi}_i, \sigma_i^2)$. To obtain a draw, we compute the posterior classification probabilities

$$P(\delta_i = k|\mathbf{y}_i, \mathbf{f}, \boldsymbol{\psi}_i, \sigma_i^2) \propto \exp \left\{ - \sum_{t=q+1}^T \left(y_{it}^* - \frac{\sum y_{it}^* f_{ikt}^*}{\sum f_{ikt}^*} f_{ikt}^* \right)^2 \right\} P(\delta_i = k), \quad k = 1, \dots, K \quad (\text{B.7})$$

where y_{it}^* and f_{ikt}^* represent the filtered values $y_{it}^* = y_{it} - \psi_{i1}y_{i,t-1} - \dots - \psi_{iq}y_{i,t-q}$ and $f_{ikt}^* = f_{kt} - \psi_{i1}f_{k,t-1} - \dots - \psi_{iq}f_{k,t-q}$, respectively. The indicator δ_i is set equal to

$$k = \left(\sum_{l=1}^K I \left\{ \left(\sum_{j=1}^l P(\delta_i = j|\cdot) \right) \leq U \right\} \right) + 1$$

where $I\{\cdot\}$ is the indicator function, $P(\delta_i = j|\cdot)$ are the normalized posterior indicator probabilities obtained from (B.7) and $U \sim U(0, 1)$ is drawn from the uniform distribution.

4. the parameters in the state transition distribution, $\pi(\boldsymbol{\gamma}|\mathbf{S}, \mathbf{x}, \mathbf{t})$. Conditional on two layers of data augmentation, the posterior turns out to be a normal distribution (see [Frühwirth-Schnatter and Frühwirth \(2010\)](#) and for additional details [Kaufmann \(2014\)](#)).

- The first layer expresses the latent state utilities $S_{st}^u, \forall s \in \{2, \dots, 4\}$, in difference to the maximum of all other *relevant* latent state utilities, and defines the binary observation $D_t^{(s)} = I\{S_t = s\}$. We obtain a linear, non-normal model:

$$S_{st}^* := S_{st}^u - S_{-s,t}^u = c + \mathbf{X}_t' \boldsymbol{\gamma}_s + v_{st}, \quad v_{st} \text{ i.i.d. Logistic}$$

Table B.5: Relevant states in \mathcal{S}_{-s}^* for $S_t = s$ given S_{t-1}

$S_t =$	2		3			4		
$S_{t-1} =$	1	2	2	3	4	1	3	4
$\mathcal{S}_{-s}^* =$	{1,4}	{1,3}	{1,2}	{4}	{4}	{1,2}	{3}	{3}

where

$$S_{st}^u = \mathbf{X}'_t \boldsymbol{\gamma}_s + \nu_{st}^u, \nu_{st}^u \text{ i.i.d. Type I Extreme Value}$$

$$S_{-s,t}^u = \max_{j \in \mathcal{S}_{-s}^*} S_{jt}^u$$

with c being a constant, $\mathbf{X}'_t = (X'_t D_{t-1}^{(1)}, \dots, X'_t D_{t-1}^{(4)})$. The elements of $\boldsymbol{\gamma}_s$ are restricted appropriately to obtain the specification in (A.1). The *relevant* other latent utilities are those corresponding to states to which the transition probability is not restricted to zero in (A.1), see table B.5.

- In the second layer, we approximate the Logistic distribution by a mixture of normals with M components, $\mathbf{R}_s = (R_{s,q+1}, \dots, R_{sT})$, $\forall s = 2, \dots, 4$. Conditional on $\mathcal{S}_s^* = (S_{s,q+1}^*, \dots, S_{sT}^*)$ and \mathbf{R}_s , we obtain the normal posterior (see Kaufmann (2014) for additional details on the moments):

$$\boldsymbol{\gamma}_s | \mathbf{S}, \mathbf{X} \sim N(g_s(\mathcal{S}_s^*, \mathbf{R}_s), G_s(\mathbf{R}_s))$$

5. the remaining parameters, which can be sampled out of standard distributions:

- (a) $\pi(\boldsymbol{\lambda} | \mathbf{y}, \mathbf{f}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\psi}) = \prod_{i=1}^N N(l_i, L_i)$, where

$$L_i = \left(\sigma_i^{-2} \sum f_{i\delta_{it}}^{*2} + L_0^{-1} \right)^{-1}, \quad l_i = L_i \left(\sigma_i^{-2} \sum y_{it}^* f_{i\delta_{it}}^* + L_0^{-1} l_0 \right)$$

- (b) $\pi(\boldsymbol{\psi} | \mathbf{y}, \mathbf{f}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\lambda}) = \prod_{i=1}^N N(q_i, Q_i) I_{\{Z(\psi_i) > 1\}}$, with

$$Q_i = \left(\sigma_i^{-2} \boldsymbol{\varepsilon}'_{i,-1} \boldsymbol{\varepsilon}_{i,-1} + Q_0^{-1} \right)^{-1}, \quad q_i = Q_i \left(\sigma_i^{-2} \boldsymbol{\varepsilon}'_{i,-1} \boldsymbol{\varepsilon}_i + Q_0^{-1} q_0 \right) I_{\{Z(\psi_i) > 1\}}$$

and where $\boldsymbol{\varepsilon}_i$ and $\boldsymbol{\varepsilon}_{i,-1}$ are, respectively, the appropriately designed left- and right-hand side matrices of the regression model:

$$\boldsymbol{\varepsilon}_{it} = \psi_1 \boldsymbol{\varepsilon}_{i,t-1} + \dots + \psi_q \boldsymbol{\varepsilon}_{i,t-q} + \boldsymbol{\varepsilon}_{it}, \quad \boldsymbol{\varepsilon}_{it} \sim N(0, \sigma_i^2)$$

(c) $\pi(\boldsymbol{\phi}, \boldsymbol{\mu} | \mathbf{f}, \mathbf{S}) = \prod_{k=1}^K N(p_k, P_k) I_{\{Z(\phi_k) > 1\}} I_{\{\mu_{k1} < \mu_{k2}, \mu_{k3} < \mu_{k4}\}}$, with

$$P_k = ([f_{k,-1} \ D]' [f_{k,-1} \ D] + \text{diag}(P_0, M_0)^{-1})^{-1} \quad (\text{B.8})$$

$$p_k = P_k ([f_{k,-1} \ D]' f_k + \text{diag}(P_0, M_0)^{-1} \text{vec}(p_0, m_0)) \quad (\text{B.9})$$

and where $f_k, f_{k,-1}, D$ are respectively, the appropriate matrices of the regression model:

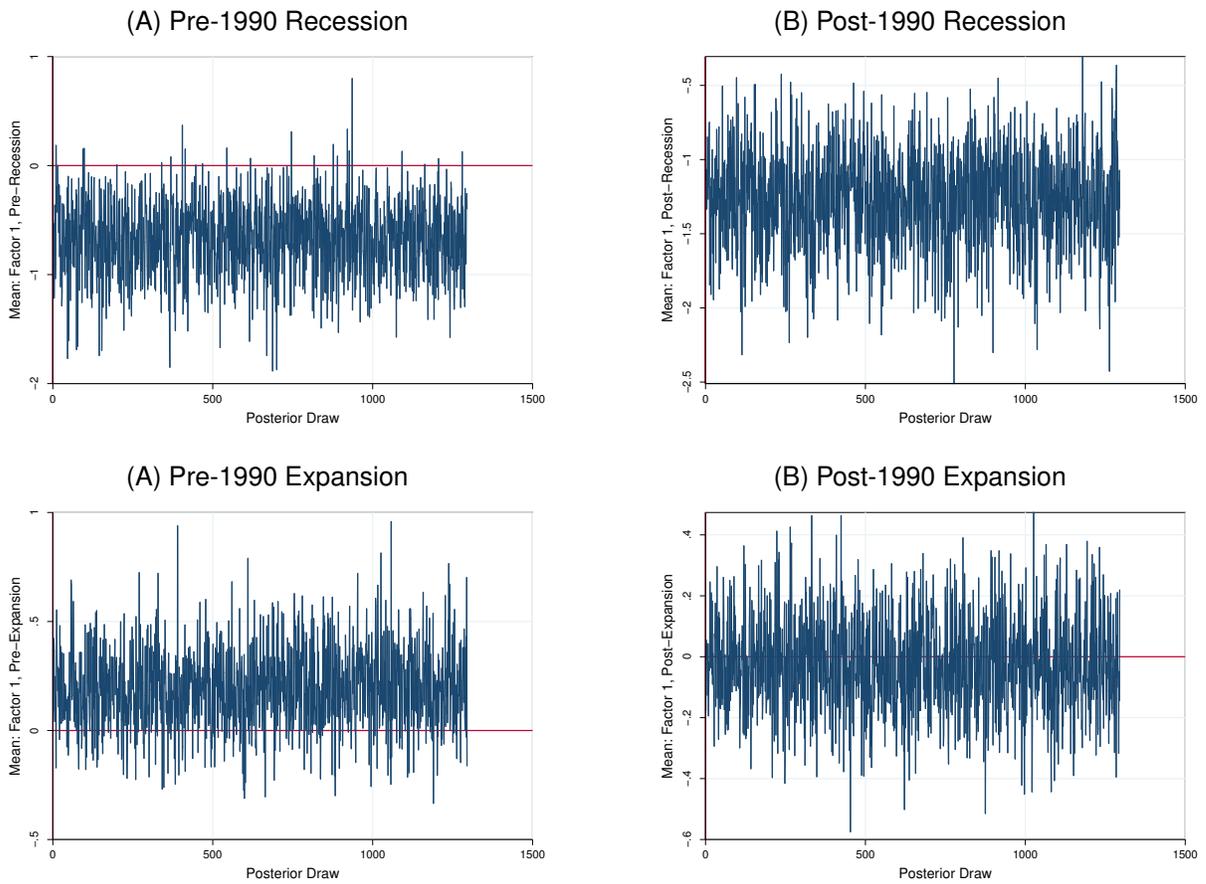
$$f_{kt} = \phi_1 f_{k,t-1} + \cdots + \phi_p f_{k,t-p} + \mu_1 D_t^{(1)} + \cdots + \mu_4 D_t^{(4)} + \nu_t, \nu_t \sim N(0, 1)$$

(d) $\pi(\boldsymbol{\sigma} | \mathbf{y}, \mathbf{f}, \boldsymbol{\delta}, \boldsymbol{\psi}, \boldsymbol{\lambda}) = \prod_{i=1}^N IG(\mathbf{e}_i, \mathbf{E}_i)$ where

$$\mathbf{e}_i = \mathbf{e}_0 + 0.5(T - q), \quad \mathbf{E}_i = \mathbf{E}_0 + 0.5 \sum_{t=q+1}^T \epsilon_{it}^2$$

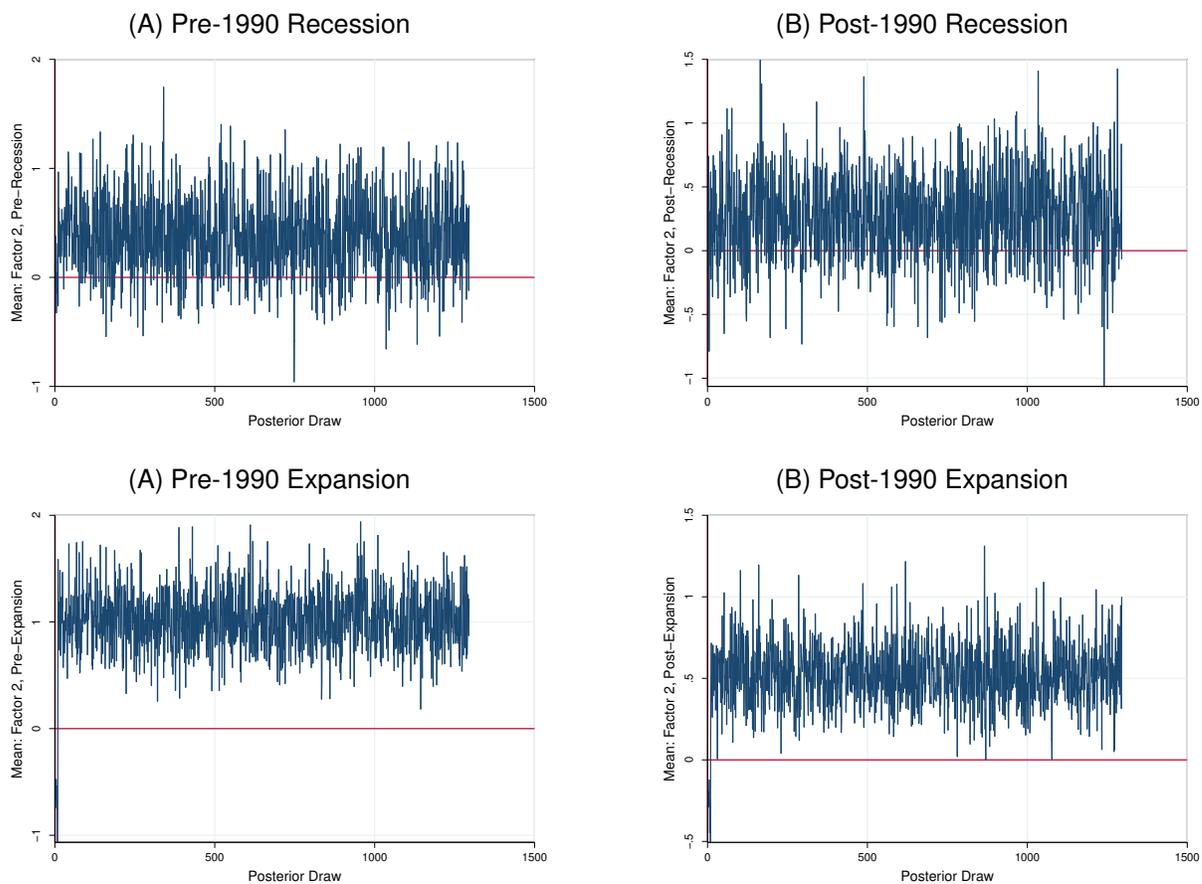
Appendix C. Sampler Convergence

Figure C.12: State Dependent Mean of Factor 1



Notes: The graphs illustrate the retained sample of posterior draws for $\mu_{k|S_t}$ in which the median cluster assignment is sampled, i.e., when $\delta = \delta_{50}$.

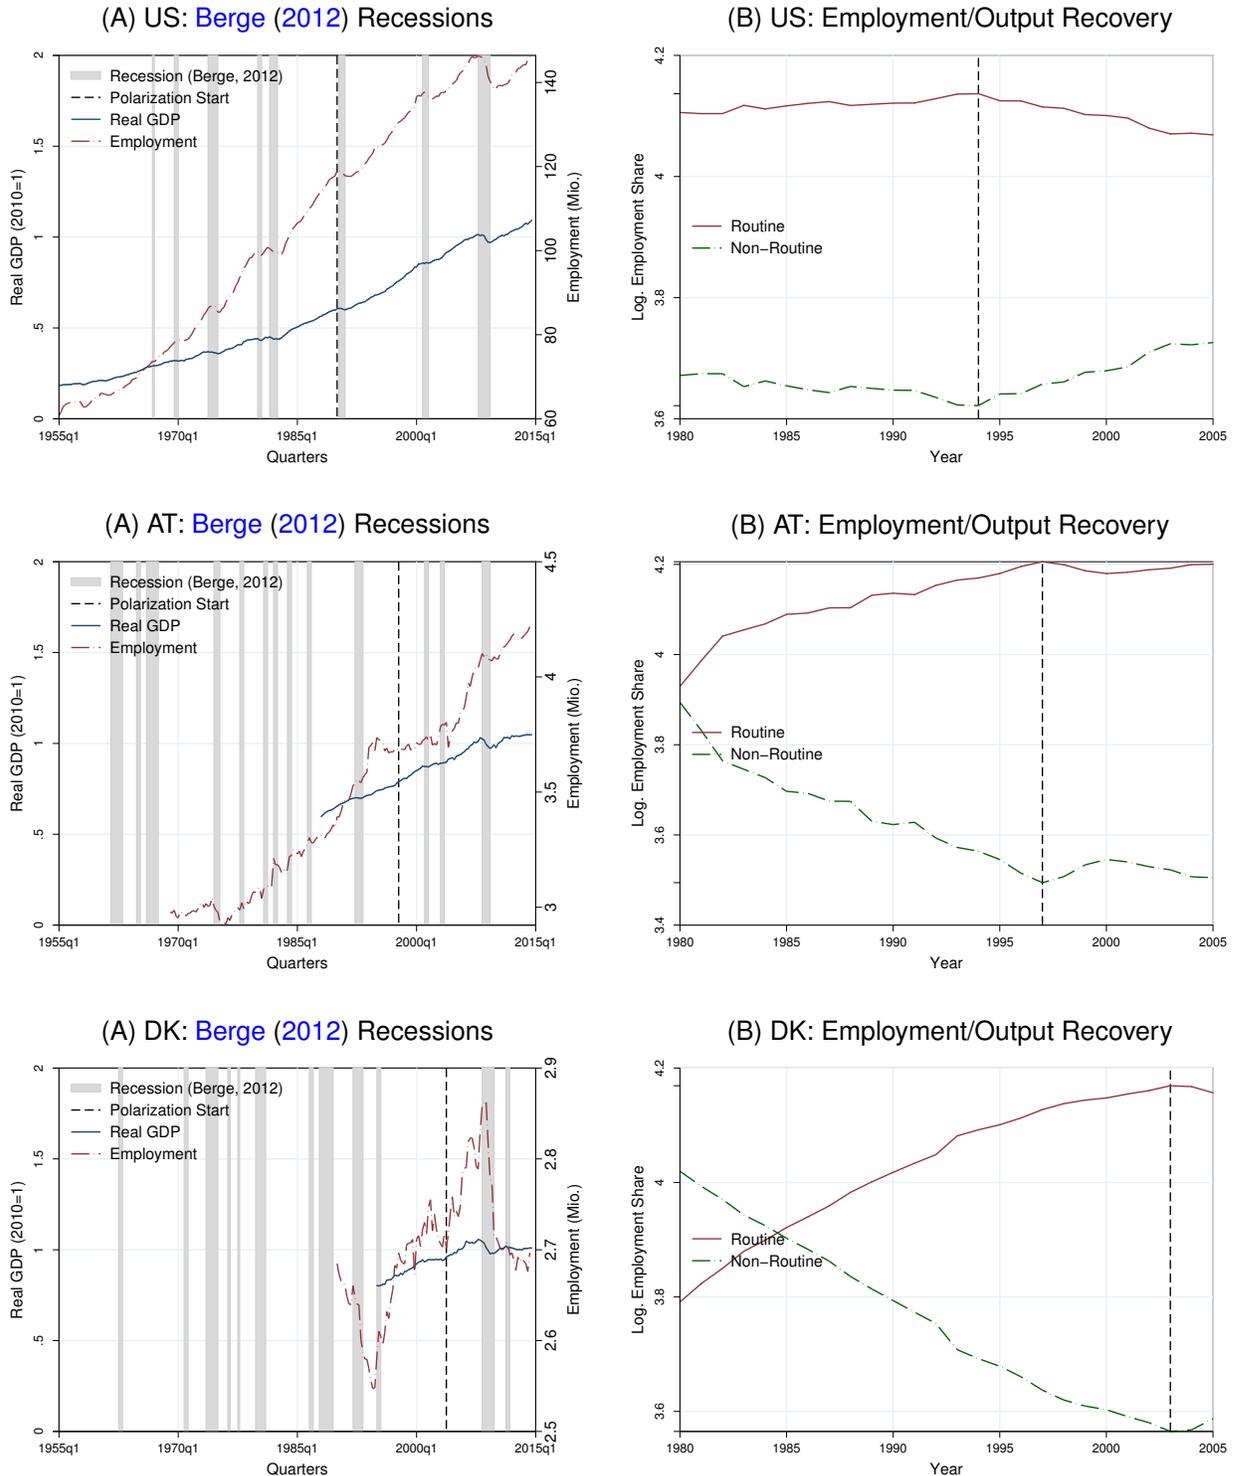
Figure C.13: State Dependent Mean of Factor 2



Notes: The graphs illustrate the retained sample of posterior draws for μ_{kS_t} in which the median cluster assignment is sampled, i.e., when $\delta = \delta_{50}$.

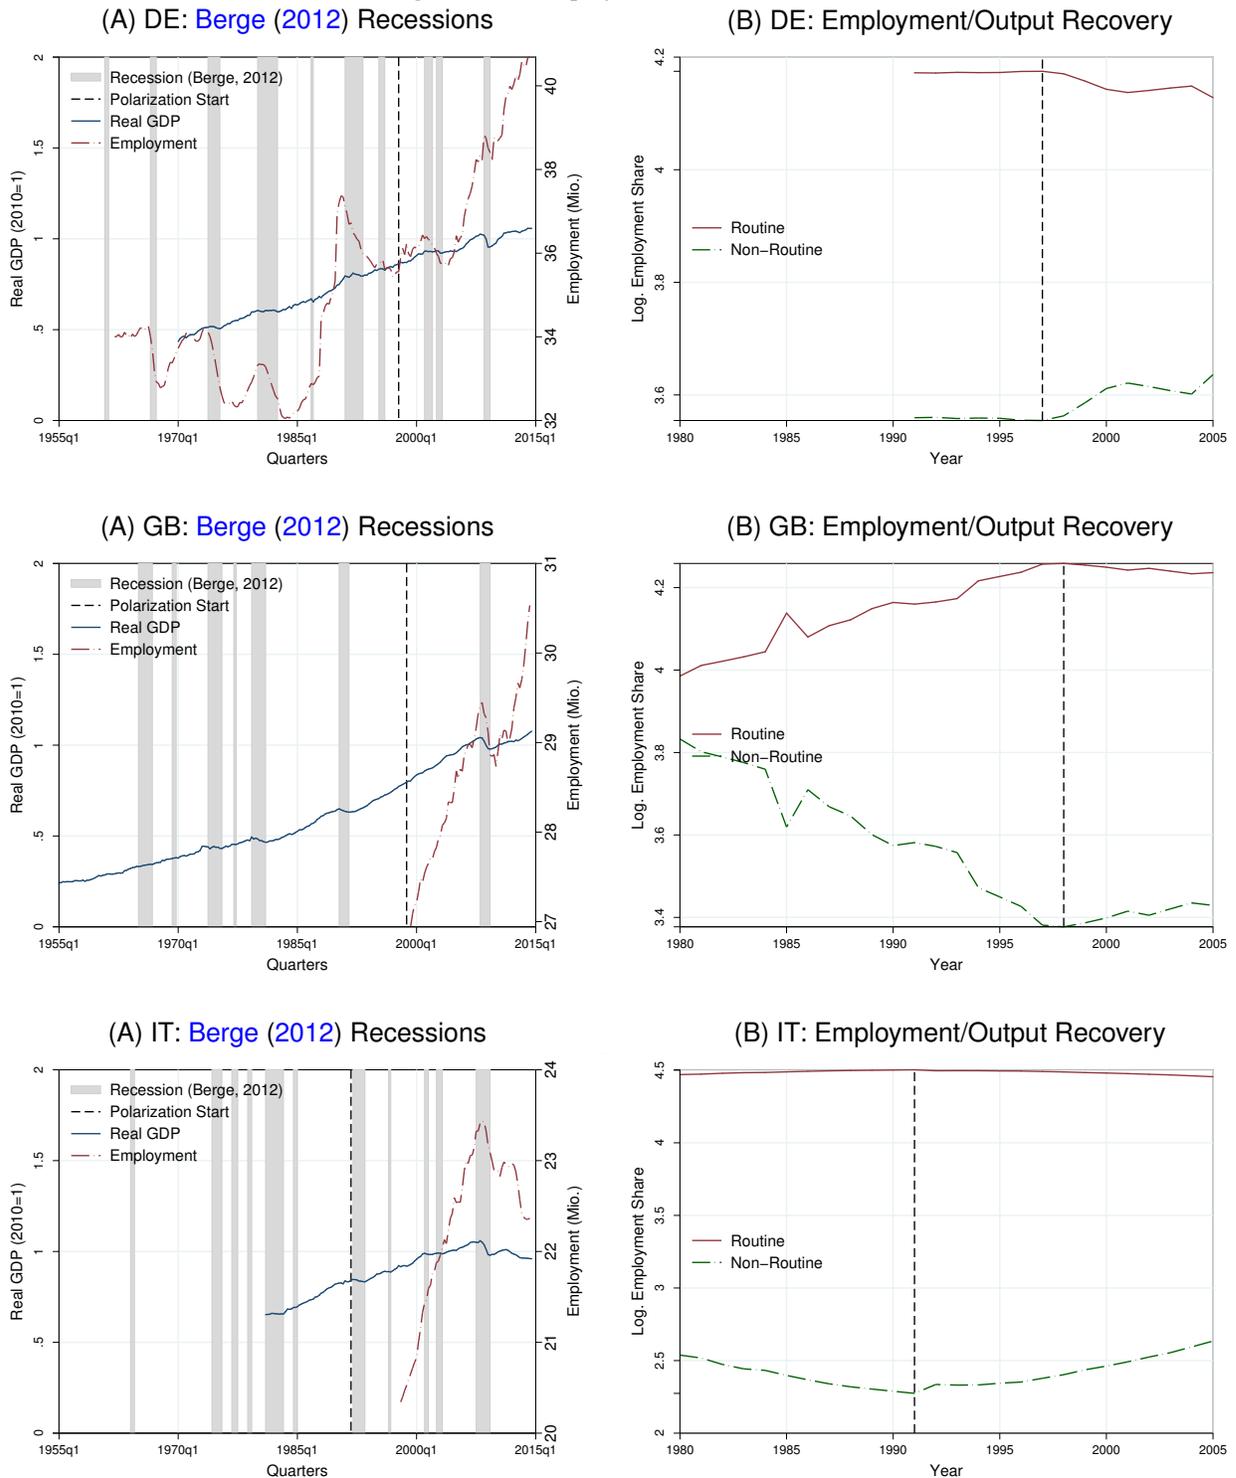
Appendix D. Recessions and Polarization in Developed Countries

Figure D.14: Recessions and Polarization



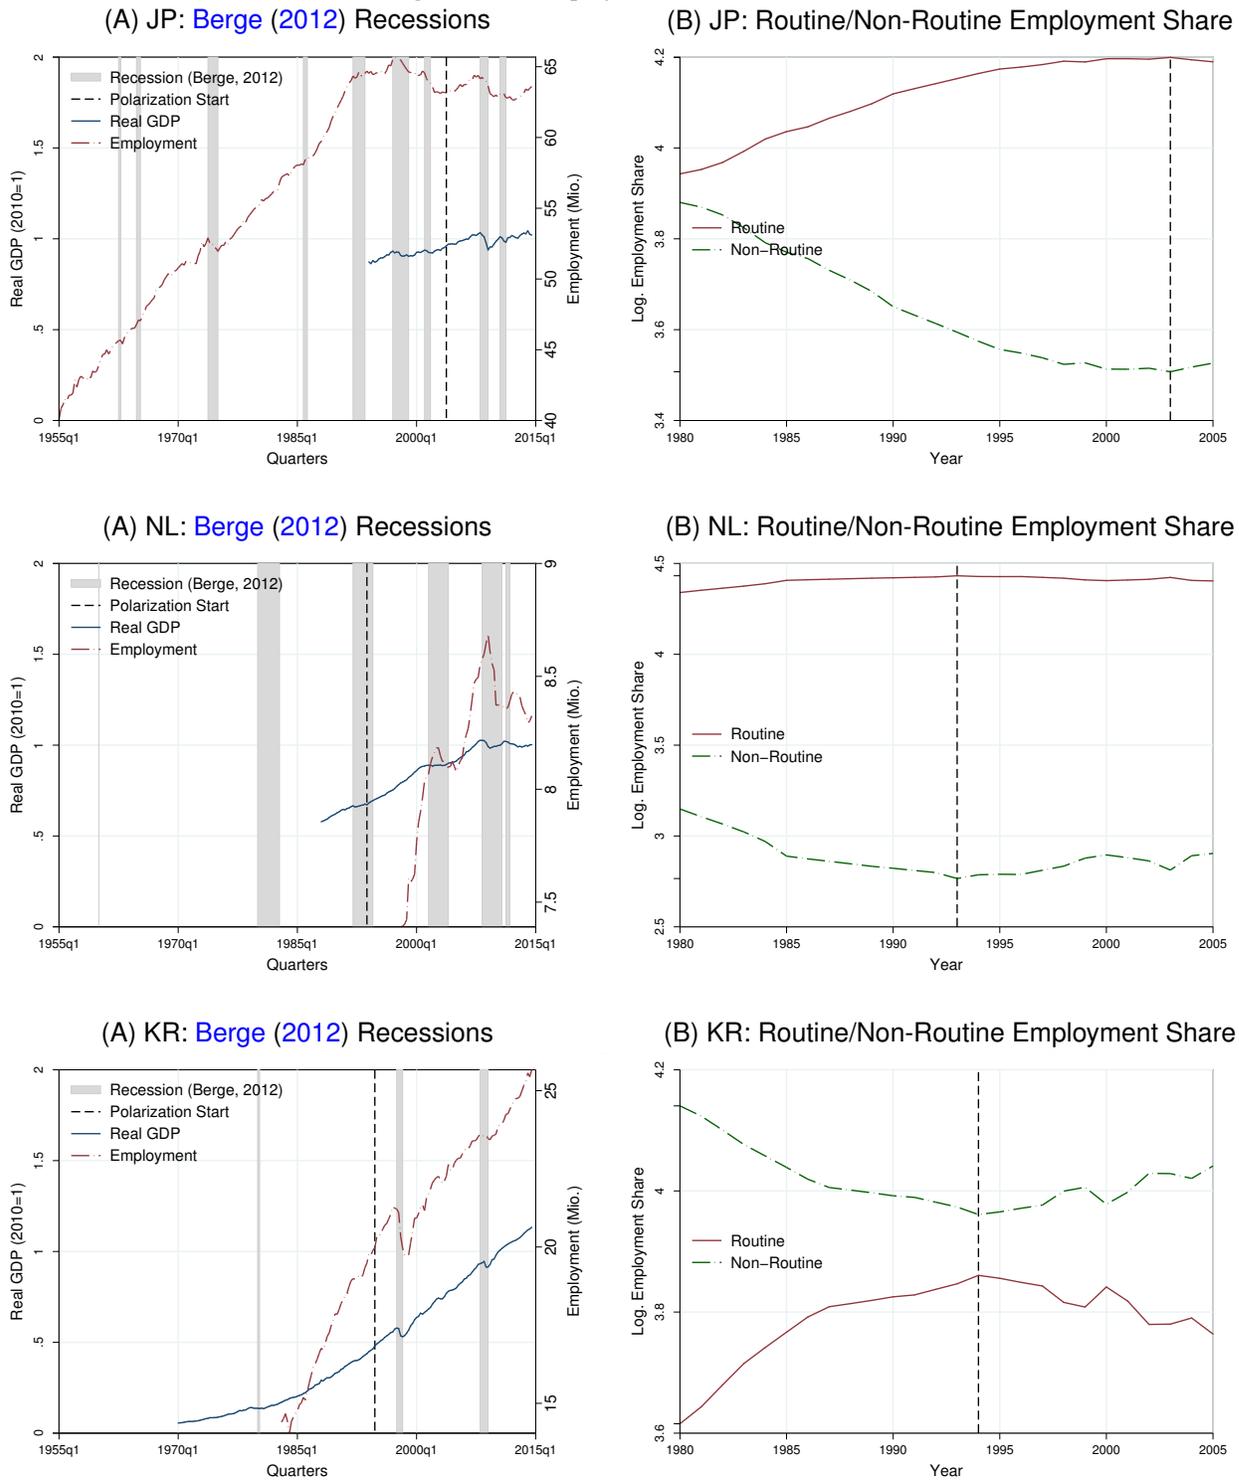
Notes: Panel A illustrates recession dates based on Berge (2012) as well as the path of employment and real GDP for countries with available data from both the OECD and EUKLEMS over the relevant time horizon. Panel B plots the log employment share in routine (middle skill) and non-routine (high+low skill) occupations in the EUKLEMS database. The dashed line indicates the peak in the log routine employment share.

Figure D.15: Employment Recoveries



Notes: Panel A illustrates recession dates based on [Berge \(2012\)](#) as well as the path of employment and real GDP for countries with available data from both the OECD and EUKLEMS over the relevant time horizon. Panel B plots the log employment share in routine (middle skill) and non-routine (high+low skill) occupations in the EUKLEMS database. The dashed line indicates the peak in the log routine employment share.

Figure D.16: Employment Recoveries



Notes: Panel A illustrates recession dates based on Berge (2012) as well as the path of employment and real GDP for countries with available data from both the OECD and EUKLEMS over the relevant time horizon. Panel B plots the log employment share in routine (middle skill) and non-routine (high+low skill) occupations in the EUKLEMS database. The dashed line indicates the peak in the log routine employment share.