The Sectoral Origins of Regional and Aggregate Business Cycles*

Sabrina Considine† Andrew Foerster‡ Sanjay Singh§

September 25, 2026

Abstract

We document that while employment growth at the national level fluctuates with business cycles, there is substantial variation in outcomes across US states. States vary in both their volatility and comovement with the nation. The differences in outcomes, volatility, and comovement holds even when considering industries at the state level. Industrial composition alone therefore cannot explain the variation in state outcomes. To interpret these facts we build a quantitative dynamic spatial general-equilibrium model with multiple industries, capital accumulation, production networks, and forward-looking worker mobility. We calibrate the model to the data, and then use the model to infer both the shocks that have driven employment across industries and states, and the underlying spatial correlation structure of these shocks. Shocks depend on aggregate, state, and industry factors. While the aggregate factor explains the bulk of national fluctuations, it only explain about a third of the variation seen across states. State and industry factors matter less for national fluctuations, but more for cross-state variability. We quantify the effects of economic mechanisms such as capital accumulation, production networks, migration, and trade in coordinating outcomes across states.

Keywords: regional business cycles; employment growth; spatial general equilibrium; production networks; labor migration

JEL Codes: E20, E32, R12, R15

*We thank Pierre-Daniel Sarte, Dave Donaldson, Johannes Wieland, Sylvain Leduc, and Michael Owyang for helpful conversations and comments, plus participants at the Federal Reserve System Committee on Regional Economics conference. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

†UC Berkeley. Email: sabrinaconsidine@berkeley.edu.

‡Federal Reserve Bank of San Francisco. Email: andrew.foerster@sf.frb.org.

§Federal Reserve Bank of San Francisco and UC Davis. Email: sanjay.singh@sf.frb.org.

1 Introduction

Business cycles in the United States are shared nationally yet experienced unevenly across places. While employment growth at the national level follows familiar business cycle patterns – growing during expansions and falling during recessions – this pattern masks heterogeneity across states. In any given quarter, the variation in employment growth rates across states is substantial relative to the magnitude of fluctuations in the national economy.

A natural starting explanation for regional heterogeneity in exposure to national business cycles is their sectoral or industrial composition. States have different makeups of their underlying compositions: some are weighted relatively more towards cyclical industries such as Manufacturing while others rely more heavily on less cyclical ones such as Education and Health; some may tend to be more diversified with employment across a broad range of industries, while others may be concentrated in one industry or another.

Documenting and understanding how sectors originate both regional and aggregate business cycles has implications for both the knowledge of how the macroeconomy operates, and for stabilization policies at the state and national level. But this understanding is complicated by the fact that business cycles are equilibrium outcomes driven by different types of shocks, with economic mechanisms such as trade, migration, capital accumulation, and production linkages that drive interdependence across industries and states.

In this paper, we investigate the sources of employment volatility across U.S. states and industries. Using detailed employment data at the state-industry level, we first document significant heterogeneity in employment fluctuations across states. While employment growth at the national level is a well-known feature of business cycles, the cross-state variability is less appreciated but still meaningful: the cross-sectional standard deviation is about 70% of the time-series standard deviation of national employment. Moreover, there is sizable dispersion in states' volatilities – with some states more variable than others – and in their correlation with the national series – with some states comoving with the nation much more than others.

This heterogeneity at the state level is present when we focus on individual industries as well. Not only is there a large dispersion in employment volatility within industry across states, we also find that there is substantial dispersion across states in the correlation of within-industry employment growth with national industry level employment growth. These facts imply a challenge for calculations relying on logic from Bartik (1991) where state-level conditions are instrumented by using state shares and national industry dynamics, since there is often a significant discrepancy between industry dynamics at the state and national level. Indeed, we compute two different counterfactual series for state employment – one using national industry weights and one using national industry growth rates – and show the former matches the data better than the latter. The implication is then that dispersion in state outcomes cannot be explained by industry composition directly. Instead, our evidence indicates that differential movements in underlying industry dynamics is the key to understanding state level variation. This finding in turn motivates a framework that nests aggregate, sectoral, and regional sources of business-cycle risk.

Motivated by these empirical observations, we develop a spatial dynamic general equilibrium model designed to capture both regional and sectoral dynamics. Specifically, we adapt the dynamic spatial equilibrium framework developed by Kleinman, Liu, and Redding (2023) to our setting of considering quarterly variation in productivity that drives business cycles and hence employment. The model features multiple states, multiple sectors, input-output linkages, an Armington model of trade, and discrete choices by workers that lead to migration between states and transitions between industries. Crucially, the model incorporates capital accumulation, a necessary feature for capturing business cycle dynamics.

We then apply the methods developed in Foerster, Sarte, and Watson (2011) to this context of multiple

sectors and states. Specifically, we apply the model to our data series of employment at the state-industry level. The model is calibrated using observed data on state-level employment, industry composition, bilateral trade flows, and migration patterns. The use of the model helps disentangle movements in the equilibrium outcome of employment growth driven by internal propagation from economic mechanisms versus those driven by external shocks. As a result, this approach allows us to make inference about the underlying productivity shocks, as well as their correlation structure. We decompose the shocks hitting the various state-industries into those stemming from aggregate or common factors, state-specific factors, industry-specific factors, or simply idiosyncrasies in that state-industry. This factor structure allows us to analyze the interplay between localized shocks and broader economic adjustments, and enables us to quantitatively analyze how shocks propagate across different regions and sectors, influencing both local and aggregate employment outcomes.

The estimated shocks reveal a clear distinction between the national cycle and differences across states. The aggregate factor generates 88 percent of the observed standard deviation of national employment growth, but only 35 percent of the observed cross-state dispersion. State factors display the opposite pattern: they generate little national movement but substantial differences across states. Industry and state-industry-specific shocks also matter much more for regional differences than their national effects alone would suggest. Employment gains in some states can offset losses in others, so shocks that are important locally may be difficult to detect in the national aggregate. The sources of the national cycle therefore provide an incomplete account of the employment experiences of individual states.

We next hold the estimated shocks fixed and examine how changes in the economy's structure affect employment fluctuations. Reducing the role of capital or intermediate inputs, restricting interstate trade, and reducing interstate worker mobility have relatively small effects on national employment volatility but much larger effects on differences across states. Reducing capital and intermediate-input use or removing interstate trade increases cross-state dispersion. Reducing interstate worker movements instead lowers dispersion by 44 percent, while leaving national volatility essentially unchanged. Taken together, the results show that economic mechanisms with similar implications for national employment can have very different implications for its distribution across states.

This paper contributes to the literature on the sources and propagation of business-cycle fluctuations. Foerster et al. (2011) distinguish aggregate from sector-specific shocks while accounting for propagation through input-output linkages, while Vom Lehn and Winberry (2022) emphasize investment networks as a source of sectoral comovement. Our analysis also relates to research on the timing and propagation of regional recessions (Owyang et al., 2005; Hamilton and Owyang, 2012) and on using regional evidence to discipline aggregate business-cycle models (Beraja et al., 2019). We bring these perspectives together by examining shocks shared across the country, within states, and within industries, and tracing their employment effects through economic connections across locations and sectors. Our focus is on whether the shock sources and transmission mechanisms important for national employment fluctuations also account for differences across states.

Second, our framework builds on quantitative models of regional and sectoral adjustment through trade, production linkages, and labor mobility (Caliendo et al., 2018, 2019; Dix-Carneiro et al., 2023). Specifically, our structural model adapts the dynamic spatial equilibrium approach of Kleinman et al. (2023) and the spatial mobility dynamics framework in the spirit of Caliendo, Parro, Rossi-Hansberg, and Sarte (2018) and Rodríguez-Clare, Ulate, and Vasquez (2025). By incorporating forward-looking investment behavior and endogenous labor mobility, we enrich the standard spatial general equilibrium models, enabling us to capture the intricate dynamics of employment volatility across regions and sectors. We combine these dynamics with

an estimated factor structure for productivity shocks to examine the sources and propagation of employment fluctuations.

Third, our analysis relates to the literature on regional labor-market adjustment and migration. Blanchard and Katz (1992) study how U.S. regional labor markets adjust to changes in employment opportunities. Foschi et al. (2025) study changes in migration in the United States, including relating changing migration patterns to industry-level shocks across states. Beyer and Smets (2015) compare labor-market adjustment and migration in Europe and the United States, while Russ et al. (2024) examine regional cyclical sensitivity within currency areas. We contribute by studying how worker mobility interacts with trade, production linkages, and capital accumulation to shape employment fluctuations across states. Our analysis distinguishes changes in the geographic distribution of employment from changes in national employment, a distinction that matters when interpreting the role of mobility in regional adjustment.

Finally, this study relates closely to literature exploring the interplay between sectoral reallocation and business cycle dynamics. Chodorow-Reich and Wieland (2020) highlight the role of sectoral reallocation shocks in influencing aggregate employment outcomes, particularly during economic downturns. By investigating intra-industry regional variations, our analysis complements their findings.

The remainder of the paper is as follows. Section 2 introduces the data on employment growth across states and sectors, and establishes our key facts about this data. Section 3 lays out a quantitative dynamic spatial equilibrium model with states and sectors. Section 4 describes the model's calibration and solution and illustrates its responses to aggregate, state, and industry productivity shocks. Section 5 combines the employment data with the model to estimate the shocks and examine their implications for national and state employment fluctuations. Then, in Section 6, we unpack the baseline results using decompositions of the data and model-based counterfactuals, which helps us better understand the drivers of business cycles. Finally, Section 7 concludes. We put additional background results, derivations, and details in the Technical Appendix.

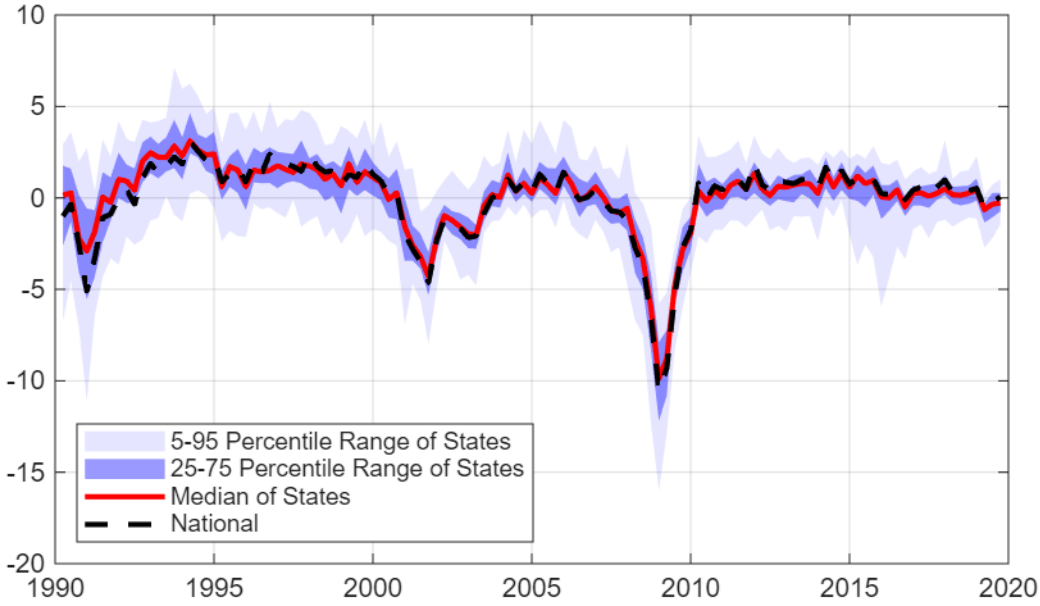
2 Employment Growth Across States and Sectors

In this Section, we introduce the data used in our analysis, and document that while there is a national business cycle, states experience substantial differences in their employment outcomes. We quantify the size of these differences and ask how far they can be understood from states' industrial composition. In particular, we find: (i) dispersion across states is large relative to national fluctuations; (ii) even within industry classifications, there is substantial heterogeneity across states; and (iii) differences in state-level outcomes depends on differences in underlying industry dynamics rather than differences in states' industry shares.

2.1 State-Level Employment Growth Data

Our dataset comprises quarterly employment levels for U.S. private-sector industries, disaggregated by state and NAICS classification, spanning 1990 to 2019. The data are sourced from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages (QCEW). Monthly observations are aggregated to the quarterly frequency by averaging across months. The dataset includes all 50 U.S. states and the District of Columbia. We exclude all government sectors and unclassified private industries retaining the following ten private sectors or industries: Natural Resources and Mining; Construction; Manufacturing; Trade, Transportation,

Figure 1: Aggregate and State Employment Growth



and Utilities; Information; Financial Activities; Professional and Business Services; Education and Health Services; Leisure and Hospitality; and Other Services.[1]

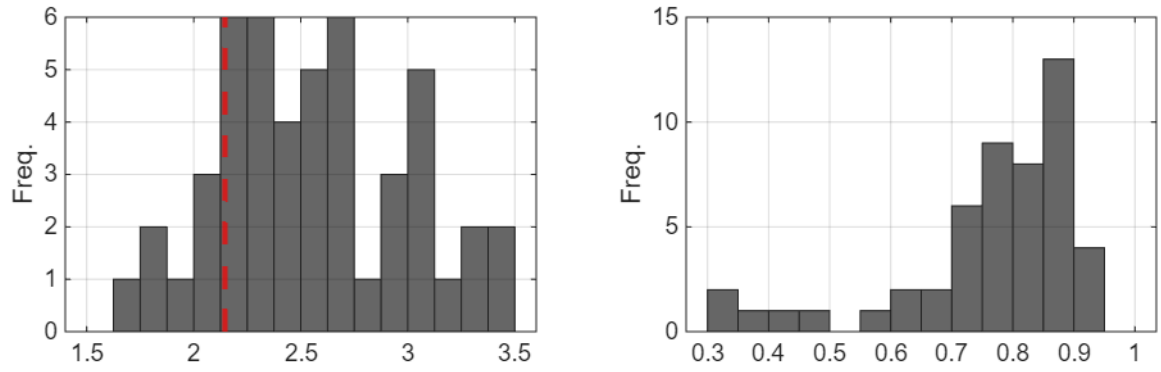
Let $\ell_{i,n,t}$ denote employment in industry i, state n, and time t. We compute annualized quarterly log growth, $400\Delta \log \ell_{i,n,t}$, and subtract each state-industry's sample mean. Denote the resulting series by $g_{i,n,t}$. Demeaning focuses attention on fluctuations around each cell's own average growth rather than persistent differences in trend growth across places.[3]

We compute state-, national-, and industry-level employment growth by aggregating the state-industry series using fixed employment shares.[4] Define $\bar{\omega}_{i,n}$ as the time average of state-industry (i,n)'s share of national private employment, and let $s_{i,n} = \bar{\omega}_{i,n} / \sum_j \bar{\omega}_{j,n}$ be its normalized share within state n. State, national, and industry growth are thus

$$g_{n,t} = \sum_i s_{i,n} g_{i,n,t}, \quad g_t = \sum_{i,n} \bar{\omega}_{i,n} g_{i,n,t}, \text{ and } \quad g_{i,t} = \sum_n (\bar{\omega}_{i,n} / \sum_m \bar{\omega}_{i,m}) g_{i,n,t} \tag{1}$$

Figure 2: Moments of State Employment Growth

(a) Standard Deviations **(b) Correlation with National**



2.2 Heterogeneity in Employment Growth Across States

Figure 1 highlights our first empirical fact: in any given quarter, there is substantial variation in the performance of states, as measured by their demeaned employment growth rates. The figure shows the growth rate in national employment g_t (black dashed line), along with shaded regions representing the 25-75 percentile range (dark blue region), and the 5-95 percentile range (light blue region). Employment growth at the national level is known to fluctuate with the business cycle–the recessions in 1990-1991, 2001, and 2007-2009 are clearly visible.

The figure shows that while the median of states closely tracks national employment growth, there is a wide range of state outcomes in any given quarter. A comparison between the variation in the black line and the width of the state ranges illustrates this point. National volatility – the time-series standard deviation of g_t – is 2.15 annualized percentage points. We assess spatial dispersion by focusing on the cross-sectional standard deviation of $g_{n,t}$; the sample median of this statistic is 1.51 percentage points. Thus, the typical amount of cross-state variation in a quarter is about 70 percent of the time-series variation in national employment growth. In other words, the aggregate series masks substantial heterogeneity at the state level.

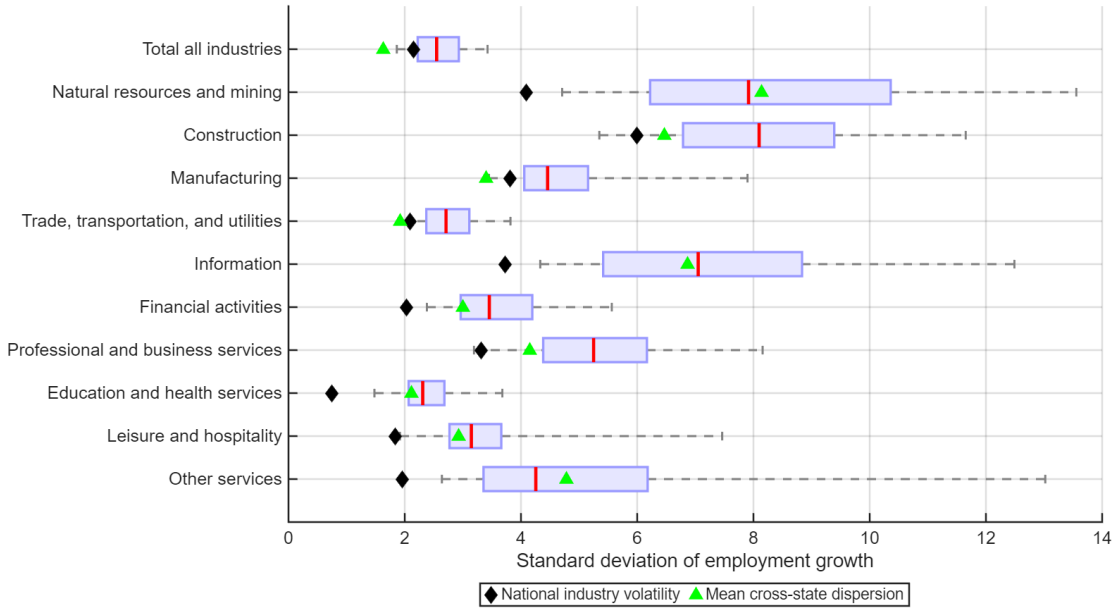
States also exhibit heterogeneity in both the volatility of their employment growths, as well as the degree of comovement with the nation as a whole. Figure 2 investigates this state-level heterogeneity by showing two key moments for state employment growth: the standard deviations and correlations with national employment growth. Figure 2a shows a histogram of each state's time-series standard deviation of employment growth, along with the national standard deviation of 2.15 (red dashed line). Here there is a wide range in the variability of states' growth rates from about 1.5 percent to around 3.5 percent. Most states are more volatile than the nation overall, but there are some that are less volatile. And the histogram is relatively flat in the sense that there is a wide dispersion in volatilities. Figure 2b shows a histogram of the correlation coefficient

[1]Each quarterly state-industry series is seasonally adjusted using the U.S. Census Bureau's X-13ARIMA-SEATS seasonal adjustment software, accessed via the R package "seasonal". Seasonal adjustment is conducted using default settings, which include automatic ARIMA model selection and detection of outliers. From the resulting decomposition, we extract the trend component, excluding seasonal and irregular components.[2] Following seasonal adjustment, employment levels for the District of Columbia were combined with those of Virginia, and the two were treated as a single geographical unit in all subsequent analysis.

[3]Carlino and Sill (2001) find both common trends and cycles across regions in the US; we ignore the former to focus on the latter.

[4]In the Technical Appendix, we document that aggregation using time-varying weights (e.g., the time-varying share of employment in a given state-industry relative to national employment) yields nearly identical series if we aggregate using mean weights (averaging across quarters). This result implies that shifts in employment levels across states and industries plays little role in explaining changes in aggregated growth rates. We thus use mean weights in all our aggregation.

Figure 3: Distribution of State-Level Employment Volatility



between each state's growth rate and the national. The histogram shows a few states have low coefficients indicating weak comovement with the national series, most states have a moderately high co-movement, but no state has a coefficient exceeding 0.95 that would indicate very close co-movement with the nation. Put together, these two panels indicate that states often differ both in how volatile their growth is, but also how closely they adhere to national employment growth.

2.3 Heterogeneity within Industries

Why do states experience different volatilities and comovement with the nation? A natural candidate explanation is that these outcomes are due to different industry compositions. In our framework, states' overall growth is due to different growth rates of employment across 10 different industries. If states vary in the shares of how many people are employed by the various sectors, then this differing industrial structure can lead to different volatilities and co-movements across states.

Sectoral composition, or industrial structure, could possibly play a role in two ways. First, states that have higher shares of their own employment in industries that are either more pro-cyclical or more volatile then inherit the properties of these sectors. For example, if Manufacturing and Construction tend to have more volatile and pro-cyclical employment than Education and Health, then states with higher shares in those industries will exhibit higher volatility and correlation with the national business cycle. Second, states that are more concentrated in one sector or another, regardless of the type of sector might have more volatility due to higher exposure to sector-specific shocks.

Figures 3 and 4 show why the explanation must also allow for differences within industries. They expand the view of states' volatilities and comovements, originally shown in Figure 2, to the industry level and across states. The within-industry distributions are wide on both dimensions.

Figure 4: Distribution of State-Level Comovement with National Employment

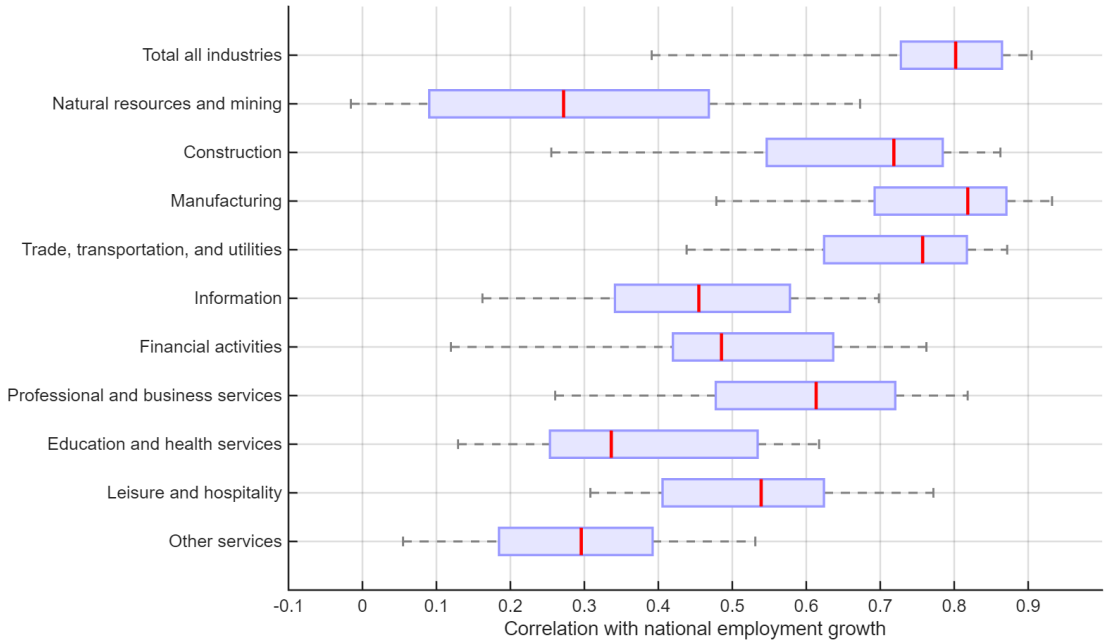
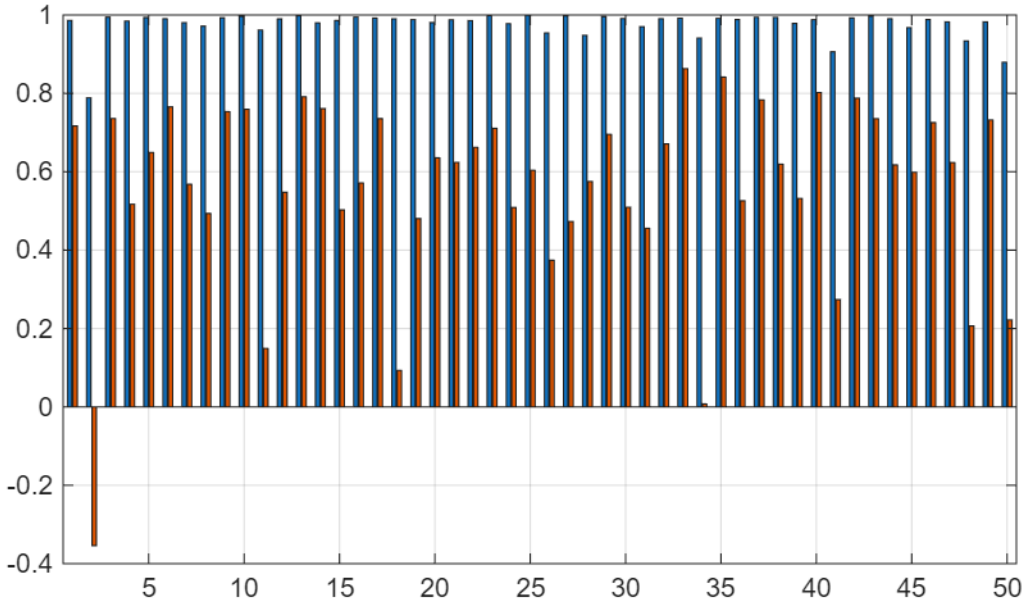


Figure 3 compares the cross-state distribution of each industry's volatility with the volatility of that industry nationally. Each row of the chart shows the nation-level volatility, the mean over time of the cross-state standard deviation, and a box and whiskers plot showing the variation in state-level standard deviations. For example, the first line shows the total of all industries. These elements have been previously discussed: from Figure 1 the national volatility is 2.15 percent (black diamond) and the mean cross-state standard deviation is 1.51 percentage points (green triangle), and from Figure 2a the standard deviations vary somewhat and tend to be above the national volatility. The remaining rows show this calculation repeated for the 10 industries we consider. The clear takeaway is that for all industries, there is substantial cross-sectional variability compared to the time-series of the industry at the national level, and the state-level industry volatility tends to have a large range.

Figure 4 reports how strongly each state-industry moves with its corresponding national industry, repeating the calculation of Figure 2b but at the industry level. Each row of the chart shows a box and whiskers plot for the correlation between each state's industry growth and the national industry growth ($\{\rho(g_{i,n,t}, g_{i,t})\}$). The top row show the comovement for all industries seen in Figure 2b, where the median correlation is 0.8, but with a longer right tail and no mass near unity. Comovement clearly differs within and across industries. Each industry has different degrees of comovement between states and the national, but all have nontrivial amounts of heterogeneity.

These results indicate that, across states but within an industry, there is substantial variation in employment growth at a given time. States vary both in how volatile they are in a given industry, and the degree to which their industries co-move with that industries' growth at a national level. These results would seemingly contradict the assertion that industrial composition plays a primary role in generating differences in employment growth at the state level. States do not simply differ in their relative exposure to different

Figure 5: State Growth Under Alternative Industry Aggregations



Notes: Each pair of bars reports the fit statistic in equation (3) for one state. Blue bars show the national-weight counterfactual (NW) in equation (2); orange bars show the national-growth counterfactual (NG) in equation (2). States are ordered as in the underlying QCEW panel.

industries. Instead, these results indicate that, even at the industry level, there are fundamental differences in outcomes across states.

2.4 Industrial Composition as an Accounting Explanation

We use two accounting exercises to separate differences in broad industry shares from differences in employment growth within the same broad industry. Let $s_i = \sum_n \bar{\omega}_{i,n}$ be industry i's national employment share. The first exercise gives every state these national industrial shares while retaining its own observed local industry growth. The second is a Bartik-style exercise, which retains each state's industry shares but replaces local industry growth with the corresponding national industry growth:

$$g_{n,t}^{\mathrm{NW}} = \sum_i s_i g_{i,n,t}, \quad g_{n,t}^{\mathrm{NG}} = \sum_i s_{i,n} g_{i,t}. \tag{2}$$

These exercises change different elements of the state-growth identity. The first asks how much the observed path changes when its aggregation weights change. The second asks how closely national industry movements, weighted by local employment shares, track that path.

Figure 5 measures the fit of each constructed series to observed state growth using

$$R_n^2(x) = 1 - \frac{\sum_t (g_{n,t} - g_{n,t}^x)^2}{\sum_t g_{n,t}^2}, \qquad x \in \{\mathrm{NW}, \mathrm{NG}\}. \tag{3}$$

The national-weight exercise (blue bars) changes the measured state outcomes very little. The median fit of the national-weight series is 0.99. If we redo our previous calculations with the new national weight

series the key moments remain unchanged: spatial dispersion falls only from 1.51 to 1.47, the median state volatility changes from 2.56 to 2.58, and the median correlation with national growth remains 0.80. In contrast, the Bartik-style series that replaces local industry growth with national industry growth (red bars) yields a median fit of 0.62. In other words, thes counterfactual series summarize an important part of state employment movements, but leave substantially more unexplained than the national-weight benchmark.

2.5 Implications and Questions

To conclude this section, we summarize our findings and turn to questions that these arise. First, we find substantial heterogeneity in employment growth across states, and that states have differing volatilities and co-movement with national growth. Second, we find that these facts extend to the industry level: even within an industry there is substantial heterogeneity across states, with different volatilities and co-movements. Lastly, we argue that industrial composition is not directly a key reason why these differences occur.

The empirical evidence here establishes a national employment cycle with unequal local incidence. Broad industry composition is visible in the data, and national industry growth has meaningful predictive content, but neither is sufficient to explain why states differ. The remaining variation could reflect state or state-industry shocks, heterogeneous exposure to aggregate and industry shocks, or equilibrium propagation through economic forces such as trade, production linkages, capital dynamics, or worker mobility.

Observed employment growth cannot distinguish these explanations on its own: a local employment decline need not originate from a local disturbance, and common employment growth need not originate from a common shock. These results lead to a number of questions that we will seek to address in the remainder of the paper. Do states experience different outcomes from one another due to differential exposure to either aggregate or industry-level shocks, and do state-level shocks play a key role? Similarly, looking within an industry, and across states, do outcomes vary so much because of different exposure to either aggregate, state, or industry level shocks?

Answering these questions requires a structural model that helps differentiate types of shocks, and disentangles comovement and its drivers. We want to know to what extent there is propagation of shocks across states and across industries. And since employment is inherently about labor, which can move across both states and sectors, we want to know to what extent migration and labor mobility plays a role in either exacerbating or dampening volatility and comovement. We now turn to a model that has these features.

3 A Quantitative Dynamic Spatial Model with States and Sectors

In this Section, we lay out a workhorse quantitative dynamic spatial model with states and industries. The framework, developed by Kleinman et al. (2023) and adapted to our setting of considering quarterly variation in productivity that drives business cycles and hence employment, features multiple states, multiple sectors, input-output linkages, an Armington model of trade, and discrete choices by workers that lead to migration between states and transitions between industries. Crucially, the model features capital accumulation, which differentiates it from other dynamic spatial models, but is a necessary feature for capturing business cycle dynamics.[5]

[5] Kleinman et al. (2023) highlight how capital and labor dynamics interact to explain declining income convergence across U.S. states. In addition, capital accumulation plays a key role in affecting how shocks propagate across industries that are linked via the production network (Foerster et al., 2011; Vom Lehn and Winberry, 2022).

3.1 Environment and Timing

There are locations indexed by $n \in \{1, ..., N\}$, and industries index by $i \in \{1, ..., I\}$. If workers are employed in a production sector, they supply one unit of labor inelastically; they can also not be employed by any of the sectors and engaged in home production (Caliendo et al., 2019), which is denoted industry $i = 0$.

Labor and capital are predetermined at the beginning of quarter t. Productivity innovations arrive, firms produce and trade, and goods and factor prices clear current markets. Workers then choose their state, industry, and employment status for quarter $t + 1$ subject to moving costs, while capital owners choose next quarter's installed capital. Thus wages, prices, and choices respond to current innovations even though current employment and installed capital do not. Workers do not have access to saving or investment, and instead live hand-to-mouth. Capital owners are geographically immobile and cannot switch sectors; they face a consumption-investment decision and hence are forward-looking.

Production in a given state and industry uses labor and capital, plus intermediate inputs sourced from other locations. States differ in their employment and production shares, input-output structure, expenditure patterns, bilateral trade exposure, and worker transition rates. This structure gives way to a rich production network that spans both locations and industries. A specific state-industry is subject to technology shocks, which we use as drivers of the business cycle. Due to both trade and migration links, shocks can propagate across and within states and industries.

3.2 Production

A competitive firm in industry i and state n produces gross output $y_{i,n,t}$ from value added $v_{i,n,t}$ and a composite material input $m_{i,n,t}$:

$$y_{i,n,t} = \left(\frac{v_{i,n,t}}{\gamma_{i,n}}\right)^{\gamma_{i,n}} \left(\frac{m_{i,n,t}}{1-\gamma_{i,n}}\right)^{1-\gamma_{i,n}}. \tag{4}$$

The value-added share $\gamma_{i,n}$ varies across state-industries. Value added combines labor $\ell_{i,n,t}$ and installed capital $k_{i,n,t}$,

$$v_{i,n,t} = z_{i,n,t} \left(\frac{\ell_{i,n,t}}{\mu_{i,n}}\right)^{\mu_{i,n}} \left(\frac{k_{i,n,t}}{1-\mu_{i,n}}\right)^{1-\mu_{i,n}}, \tag{5}$$

where $\mu_{i,n}$ is labor's share in value added, $z_{i,n,t}$ denotes the state-industry specific total factor productivity that follows a random walk process

$$\log z_{i,n,t} = \log z_{i,n,t-1} + \varepsilon_{i,n,t}, \tag{6}$$

where $\mathbb{E}_{t-1}\varepsilon_{i,n,t} = \mathbf{0}$. Thus $\varepsilon_{i,n,t}$ is an innovation to productivity growth, and a one-time innovation permanently shifts the level of productivity. The covariance structure of these latent innovations–that is, how productivity evolves across regions and industries–is crucial to our analysis and discussed later.

Firms assemble a composite material input by bundling inputs $m_{ji,n,t}$, which are intermediates flowing from industry j to industry i in location n. They do so according to the relevant share $\phi_{ji,n}$:

$$m_{i,n,t} = \prod_j \left(\frac{m_{ji,nt}}{\phi_{ji,n}}\right)^{\phi_{ji,n}}, \qquad \sum_j \phi_{ji,n} = 1. \tag{7}$$

where the share $\phi_{ji,n}$ measures industry i's use of inputs from industry j in state n. Equations (4)–(7) create

propagation across industries: a productivity disturbance changes the cost of the affected good and therefore the production costs of downstream users.

Trade of gross output between locations is subject to iceberg variable trade costs, where $\tau_{i,nm} \geq 1$ units of gross output from industry i must be shipped from location n in order for one unit to arrive in location m, where $\tau_{i,nm} > 1$ if $n \neq m$ and $\tau_{i,nm} = 1$ if $n = m$. Trade transmits a disturbance across states; input-output demand transmits it across industries; and their interaction determines which state-industries gain or lose expenditure.

3.3 Goods Importers

Goods importers in a given location n and for an industry i purchase goods from the various locations $m \in \{1, ..., N\}$ and combine them using a constant elasticity of substitution technology:

$$q_{i,n,t} = \left(\sum_m q_{i,mn,t}^{\frac{\theta}{\theta+1}} \right)^{\frac{\theta+1}{\theta}}. \tag{8}$$

This Armington model of trade has goods of type i that flow from m to n, and get combined with a trade elasticity of θ. Goods importers then sell their products $q_{i,n,t}$ to be used in materials inputs in n ($\{m_{ij,n,t}\}_{j=1}^I$, for workers' consumption ($\{c_{ij,n,t}^w\}_{j=1}^I$), and for capital owners' consumption and investment ($\{c_{ij,n,t}^k, x_{ij,n,t}^k\}_{j=1}^I$).

3.4 Workers' Consumption and Migration

As noted, workers are either employed in an industry $i \in \{1, ..., I\}$ or are unemployed and have home production ($i = 0$).

An employed worker in location n in industry i supplies a unit of labor inelastically, and earns a wage $w_{i,n,t}$. They have flow utility

$$\log u_{i,n,t}^w = \log \left(b_n c_{i,n,t}^w \right) \tag{9}$$

where b_n is a location-specific utility fixed component such as amenities, and the consumption level is given by

$$c_{i,n,t}^w = \prod_j \left(\frac{c_{ji,n,t}}{\psi_{j,n}} \right)^{\psi_j}, \qquad \sum_j \psi_{j,n} = 1. \tag{10}$$

where $\psi_{j,n}$ denote the consumption shares of different industry products, which potentially differ across locations.

An unemployed worker receives utility from home production $u_{0,n,t} = b_n h_n$.

At the beginning of time t, workers' location is fixed, with a mass $\ell_{i,n,t}$ at each location n in industry i. Total population is normalized to unity, so

$$\sum_{n=1}^N \sum_{i=0}^I \ell_{i,n,t} = 1. \tag{11}$$

At the end of each period t, workers make forward-looking decisions on location and industry. They have

value function given by

$$v_{i,n,t} = \log u_{i,n,t} + \max_{\substack{m\in\{1,\ldots,N\}\\ j\in\{0,\ldots,I\}}} \left\{\beta \mathbb{E}_t v_{j,m,t+1} + \kappa_{ij,nm} + \rho\varepsilon_{j,m,t}\right\}. \tag{12}$$

Migration between location-industry (i,n) and (j,m) is governed by a fixed migration cost $\kappa_{ij,nm}$ and a shock $\varepsilon_{j,m,t}$ that follows an extreme value distribution with CDF $F(\varepsilon) = e^{-e^{(-\varepsilon-\bar{\gamma})}}$, where $\bar{\gamma}$ is the Euler-Mascheroni constant.

This discrete location and industry choice framework implies that the probability of migrating from (i,n) to (j,m), the out-migration probability, is given by

$$D_{ij,nm,t} = \frac{\left(\frac{\exp(\beta \mathbb{E}_t v_{j,m,t+1})}{\kappa_{ij,nm}}\right)^{1/\rho}}{\sum_{o=1}^{N}\sum_{k=0}^{I}\left(\frac{\exp(\beta \mathbb{E}_t v_{k,o,t+1})}{\kappa_{ik,no}}\right)^{1/\rho}}. \tag{13}$$

As a result, these probabilities implies an evolution of labor relating labor at (m,j) to all previous (i,n):

$$\ell_{j,m,t+1} = \sum_{n=1}^{N}\sum_{i=0}^{I} D_{ij,nm,t}\ell_{i,n,t}. \tag{14}$$

Mobility affects regional cycles through two margins. Workers can switch between employment and home production, which changes aggregate employment. Conditional on employment, they can switch industries or spatial units, which redistributes employment across the 500 producing state-industries. Because labor and capital are stocks, a one-time productivity innovation can generate employment adjustment over many subsequent quarters.

3.5 Capital Owners

Capital owners in location n in industry i have utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log\left(c^k_{i,n,t}\right) \tag{15}$$

where the consumption level $c^k_{i,n,t}$ follows the same functional form as in the workers' problem (10). They purchase consumption and investment goods across industries, own capital $k_{i,n,t}$ with return $r_{i,n,t}$. The have budget constraint

$$\sum_{j=1}^{I} p_{j,n,t}\left(c^k_{ji,n,t} + x^k_{ji,n,t}\right) = r_{i,n,t}k_{i,n,t} \tag{16}$$

and face a law of motion of capital given by

$$k_{i,n,t+1} = (1-\delta)\,k_{i,n,t} + x_{i,n,t}. \tag{17}$$

where $x_{i,n,t}$ is assembled using the same form as consumption, and δ is the rate of depreciation.

Investment uses the same local final-goods bundle as consumption. Define the gross real return

$$R_{i,n,t} = 1 - \delta + \frac{r_{i,n,t}}{p_{n,t}}. \tag{18}$$

where $p_{n,t}$ is the local consumption price index. Log utility implies a constant saving rate:

$$c^k_{i,n,t} = (1-\beta)R_{i,n,t}k_{i,n,t}, \qquad k_{in,t+1} = \beta R_{i,n,t}k_{i,n,t}. \tag{19}$$

Capital owners therefore carry a fraction β of current wealth into the next quarter, so the current real rental return determines the evolution of installed capital. A productivity disturbance changes current investment expenditure and the productive capacity inherited next quarter. The resulting capital stock can affect the size and persistence of later employment responses.

3.6 Market Clearing

Market clearing for goods implies that the value of gross output across industries and locations satisfies

$$p^y_{i,n,t}y_{i,n,t} = \sum_m \sum_j S_{ij,nm,t}p^y_{j,m,t}y_{j,m,t}, \tag{20}$$

where $S_{ij,nm,t}$ denotes the share of expenditures on goods produced from (i,n) as a fraction of total expenditures within (j,m).

In addition, market clearing, along with the cost-minimizing bundle for (10), means that consumer prices in each state are related to gross-output prices across all (j,m) and iceberg trade costs according to

$$p_{n,t} = \prod_j \left(\left(\sum_m \left(\tau_{j,mn}p^y_{j,m,t} \right)^{-\theta} \right)^{-\frac{1}{\theta}} \right)^{\psi_{j,n}}, \tag{21}$$

Similarly, industry-state materials prices are determined by

$$p^m_{i,n,t} = \prod_j \left(p^q_{j,n,t} \right)^{\phi_{ji}} \tag{22}$$

3.7 Equilibrium

A competitive equilibrium is a sequence of quantities, prices, worker values, migration shares, and trade shares such that firms and importers minimize costs, workers and capital owners solve their dynamic problems, goods and factor markets clear, and the laws of motion for labor and capital hold. The equilibrium conditions of the model collapse into a system of equations encompassing both price and quantity variables, as shown in the Technical Appendix.

4 Model Solution and Dynamics

In this section, we describe our model solution and give and overview of the dynamics, discuss how we use the model solution as a filter for the data described in Section 2, and discuss our baseline empirical results.

4.1 Solving the Model

To solve the model, we begin by assembling the equilibrium conditions as described in Section 3.7. We then log-linearize these equations around a deterministic steady state in the aggregate. This steady state is

associated with no changes in technology and a stationary distribution of workers across states and industries. The linearized equations take the form

$$\mathbf{A}\mathbb{E}_t\mathbf{y}_{t+1} = \mathbf{B}\mathbf{y}_t + \mathbf{C}\mathbf{z}_t \tag{23}$$

where $\mathbf{y}_t = \begin{bmatrix} \tilde{\mathbf{w}}_t' & \mathbf{p}_t' & \tilde{\mathbf{p}}_t^{m\prime} & \tilde{\mathbf{p}}_t^{y\prime} & \tilde{\mathbf{p}}_t^{v\prime} & \tilde{\mathbf{v}}_t' & \tilde{\mathbf{l}}_t' & \tilde{\mathbf{k}}_t' \end{bmatrix}'$ is the set of endogenous variables.

As shown in the Technical Appendix, the matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ are functions of the underlying parameters of the economy. We set the discount factor $\beta = 0.99$, the depreciation rate $\delta = 0.026$, the trade elasticity $\theta = 5$, and the migration elasticity $\rho = 5.34$. The solution also depends on a number of matrices: the value-added shares $\mathbf{\Gamma_d} = \text{diag}(\gamma_{i,n})$, the labor shares $\mu_\mathbf{d} = \text{diag}(\mu_{i,n})$, the input-output shares $\mathbf{\Phi} = [\phi_{ij,n}]$, the consumption shares $\mathbf{\Psi_d} = \text{diag}(\psi_{i,n})$, the steady-state out-migration matrix $\mathbf{D} = [D_{ij,nm,ss}]$, the steady-state in-migration matrix $\mathbf{E} = [E_{ij,nm,ss}]$, the steady-state expenditure share matrix $\mathbf{S} = [S_{ij,nm,ss}]$, and the steady-state income share matrix $\mathbf{T} = [T_{ij,nm,ss}]$. We set these equal to observed shares based upon state-industry level data; see the Technical Appendix for details.

We solve the resulting linear system using the algorithm in King and Watson (1998, 2002), which has the advantage of being able to solve a large linear system such as ours. The solution of the system (23) has the form

$$\begin{aligned} \mathbf{y}_t &= \mathbf{\Pi}\mathbf{s}_t \\ \mathbf{s}_t &= \mathbf{M}\mathbf{s}_{t-1} + \mathbf{G}\varepsilon_t \end{aligned} \tag{24}$$

where $\mathbf{s}_t = \begin{bmatrix} \tilde{\mathbf{l}}_t' & \tilde{\mathbf{k}}_t' & \tilde{\mathbf{z}}_t' \end{bmatrix}'$. For our model filtering, we assume a factor structure of ε_t, to be discussed in Section 5.1.

4.2 Model Dynamics

Before moving to our empirical results, we present three experiments that illustrate the dynamics of the model solution. The first raises productivity in all states and industries, the second in all industries of one state at a time, and the third in one industry across all states. In each experiment, productivity in the affected state-industries rises by approximately 1%, with no subsequent innovations. These experiments show how the same model translates disturbances of different geographic and industrial scope into national employment growth and differences across states.

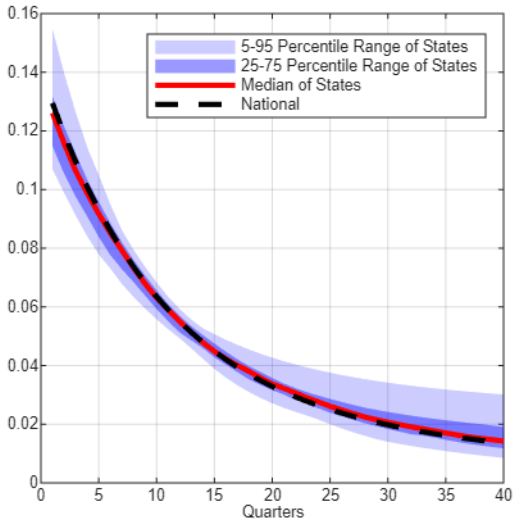
4.2.1 A Common Productivity Shock

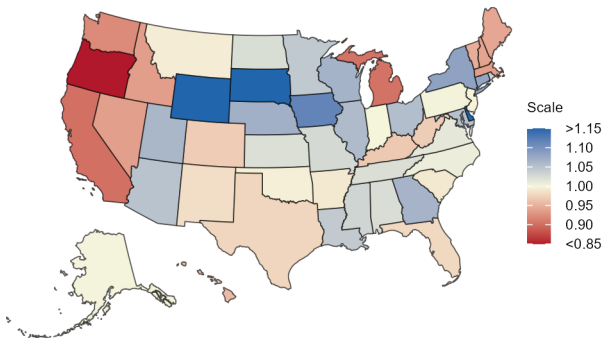
The first experiment examines the model's response to an equal-sized productivity shock across all states and industries. Since our model has non-employment (or home production), an increase in productivity across all sectors raises the real wage, thereby raising the value of employment in any state or sector relative to non-employment and draw workers into employment. National employment can therefore expand.

Figure 6 shows the effects of this common shock. Panel (a) plots the impulse response of national employment growth together with the median and percentile ranges of state employment growth. Employment begins to adjust in the quarter after the shock, with employment taking roughly 20 quarters to converge to its new long-term level. Not all states have equal employment increases, as some experience relatively higher growth; second, the dynamics following an aggregate shock varies across states. In other words, the shock induces a strong initial rise in employment that later subsides, suggesting that workers enter employment in certain states before relocating to others offering better employment opportunities.

Figure 6: Effects of a 1% Shock to all State-Industries

(a) National and State Growth Rates **(b) Cumulative Change, Relative to Aggregate**





Panel (b) shows the geographic distribution of cumulative employment changes over 40 quarters, expressed relative to the national cumulative change. Blue states have larger proportional employment gains than the nation, while red states have smaller gains. A value below one indicates a smaller gain relative to the national response, not necessarily an employment decline. These results highlight that the model can generate regional dispersion even when the underlying productivity innovation is common to all locations.

4.2.2 Shocks Originating in Individual States

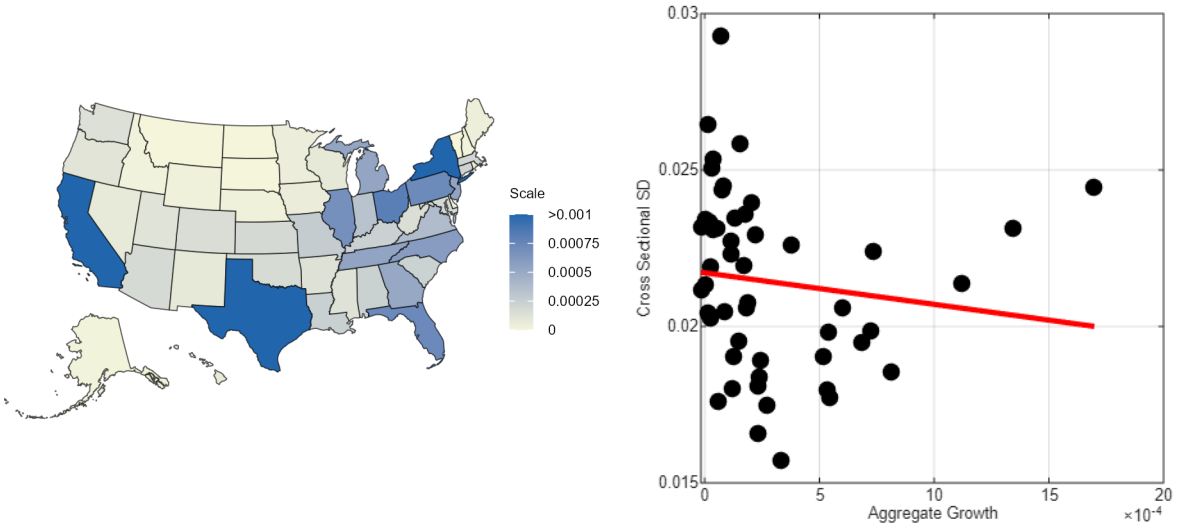
The second experiment considers shocks originating in individual states. We raise productivity by the same amount in all producing industries of one state, leaving productivity elsewhere unchanged, and repeat the exercise for each of the 50 states. Trade, input-output linkages, and worker mobility allow each disturbance to affect employment beyond the state in which it originates.

Figure 7(a) maps the national employment-growth response 40 quarters after each state-level shock. Each state's color therefore records the national effect of a shock originating in that state, not that state's response to a single common shock. Shocks to large states such as California, Texas, and New York generate relatively large national responses. These comparisons reflect differences in state size as well as industrial structure and economic connections: an equal percentage productivity increase in each state is not an equal-sized disturbance to the national economy.

Panel (b) compares the national growth response with the cross-state standard deviation of employment-growth responses at the same horizon. Each point corresponds to a different state receiving the shock. The relationship is not strongly positive: shocks producing larger national growth responses need not produce greater dispersion across states. The national effect of a disturbance and the regional differences it creates are therefore distinct features of its transmission.

Figure 7: Effects of 1% Shock to Each State

(a) Aggregate Growth from Each States' Shock **(b) Dispersion vs. National Growth**



4.2.3 Industry-wide shocks

The third experiment raises productivity in one industry across all states, leaving productivity in other industries unchanged. We repeat the exercise for each of the ten industries. Figure 8 reports the distribution across states of cumulative changes in total employment over 40 quarters. The black diamonds show the corresponding change in total national employment.

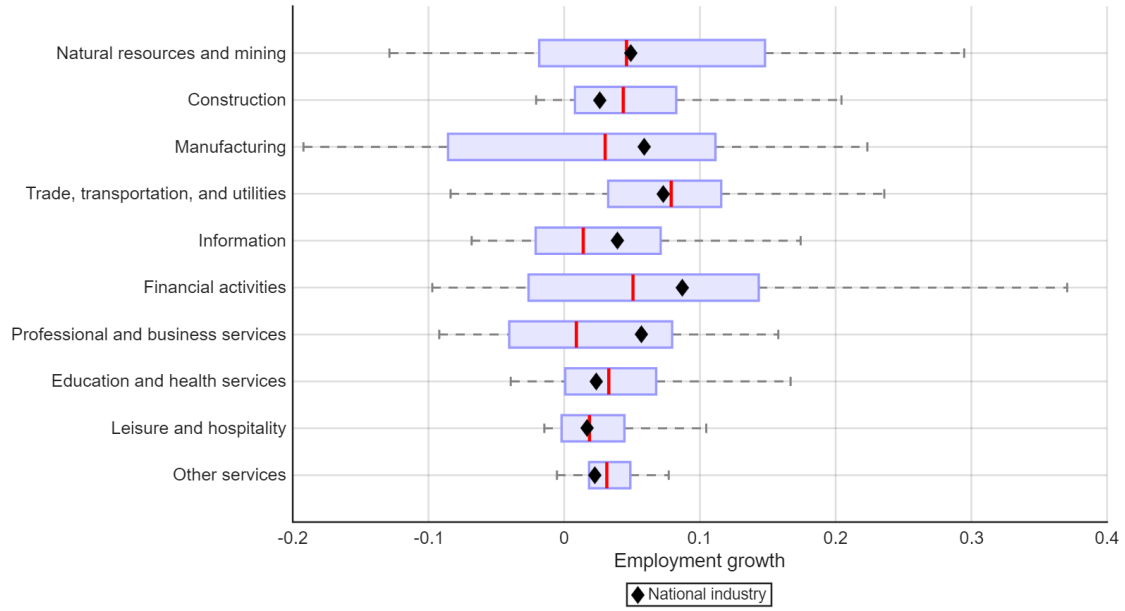
Industry shocks also produce heterogeneous state outcomes. A manufacturing productivity improvement, for example, raises total employment nationally but produces employment gains in some states and declines in others. Other industries generate more tightly clustered state responses. These differences reflect the interaction of industrial specialization with the model's production and adjustment mechanisms. A shock shared nationally within an industry can therefore generate substantial geographic variation, and a positive national employment response need not imply employment gains in every state.

Taken together, these experiments illustrate a key feature of the model: the location and industry in which a disturbance originates need not coincide with the locations and industries in which its employment effects are felt. Common shocks can generate regional differences, and geographically or industrially concentrated shocks can have effects throughout the economy. The experiments describe this transmission but do not establish which shocks account for observed employment fluctuations. We next combine the model solution with a factor structure for productivity innovations to study the sources of the national and regional variation documented in Section 2.

5 Baseline Empirical Results

Having described the model solution and its dynamics, we turn to merging the model with our data on employment growth at the state-industry level.

Figure 8: Effects of 1% Shock to Each Industry: Distribution of State Responses



5.1 Estimation and Filtering

To do so, we combine the model solution described in Section 4 with a factor structure for state-industry productivity innovations:

$$\varepsilon_{i,n,t} = \lambda_{i,n} f_t + \lambda_i^n f_{n,t} + \lambda_n^i f_{i,t} + u_{i,n,t}. \tag{25}$$

This factor structure implies that a shock to a given state-industry in quarter t depends on an aggregate factor f_t with loading $\lambda_{i,n}$, a state factor $f_{n,t}$ with loading λ_i^n, an industry factor $f_{i,t}$ with loading λ_n^i, and an idiosyncratic portion $u_{i,n,t}$ specific to the state-industry. For example, a productivity innovation in California's construction industry may reflect aggregate conditions, conditions specific to California, conditions shared by construction across states, or forces unique to construction in California. Employment responds to these innovations through the model's dynamics and economic connections.

Let Λ^A, Λ^S, and Λ^I collect the coefficients for the aggregate, state, and industry factors, respectively. We assume that the factors are mutually independent and normalize each factor's variance to one. Shocks specific to individual state-industries are independent of the factors and of one another, with variances collected in the diagonal matrix Σ_u. The covariance matrix of productivity shocks is therefore

$$\Sigma = \Lambda^A(\Lambda^A)' + \Lambda^S(\Lambda^S)' + \Lambda^I(\Lambda^I)' + \Sigma_u.$$

Each term captures a different source of shared movements in productivity. These assumptions imply that a shock originating in one state or industry can affect employment elsewhere through the model's economic connections.

Observations and latent states. The empirical panel contains quarterly employment growth for 50 states and 10 industries. Each of the 500 series is expressed as annualized log growth and demeaned over the estimation sample. We hold the calibrated model parameters and equilibrium response matrices fixed and estimate the disturbances and their factor structure.

Combining the model solution with equation (25) yields a state-space representation for employment growth. The model is not invertible, as we have $N(I+1) = 550$ labor, and $NI = 500$ capital and productivity series to back out from only 500 observations on labor growth. In other words, the observations on employment do not uniquely pin down paths for productivity growth. Instead, we develop a Gibbs sampler tailored to our specific application, which allows us to jointly estimate the shocks and the underlying factor structure, which we can then use for inference.

Two-block Gibbs procedure. Our estimation framework alternates between two steps; here we give a brief overview, see the Technical Appendix for details.

First, given an estimate of the shock covariance matrix, we use observed employment growth and the model's dynamics to draw a possible history of productivity shocks and the economy's underlying states (Durbin and Koopman, 2002). In this setup, employment is predetermined when a current innovation arrives, so the productivity innovation first affects employment one period after it enters productivity. Inference therefore uses the dynamic relationship between shocks and employment to infer a possible sequence of shocks.

Second, given a drawn productivity-shock history, we draw a factor decomposition of aggregate, state, industry, and state-industry-specific components. We update both the factors and the coefficients measuring each state-industry's exposure to them, together with the variances of the remaining state-industry-specific shocks. These updates imply a new covariance matrix, which is then used in the next draw of the shock history.

Iterating between these steps – first drawing a sequence of shocks given the factor structure and hence covariance matrix, and second drawing the factor structure given a sequence of shocks – represents a Gibbs sampler that characterizes the joint distribution. We perform 1000 draws, and use the medians of the MCMC chain for the analysis. We next turn to an empirical investigation into the shock structure, and then turn to employment growth.

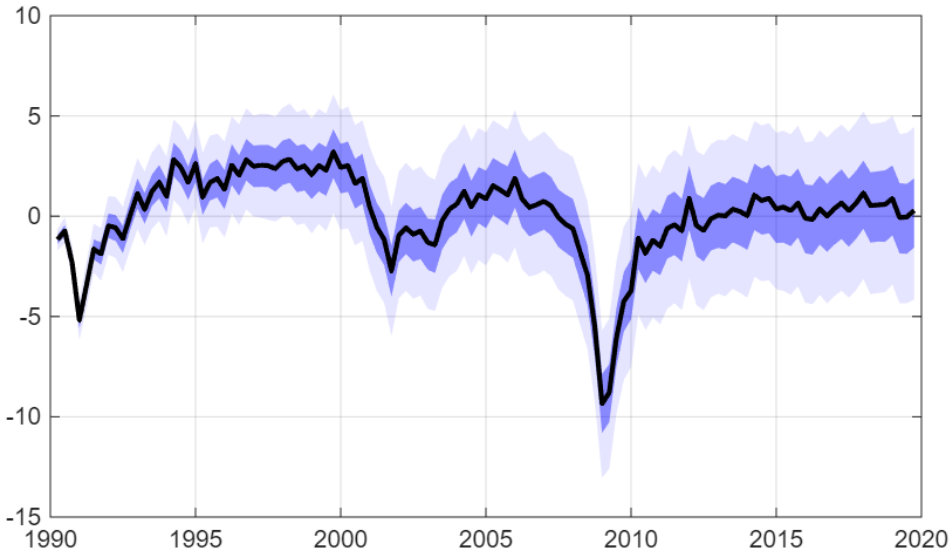
5.2 Aggregate, State, and Industry Components of Productivity Shocks

Our first set of empirical results investigates the sources of the productivity shocks $\{\varepsilon_{i,n,t}\}$. In particular, we examine the evolution of the aggregate factor f_t and the extent to which different state-industries are exposed to it; similarly, we examine heterogeneity in how state-industries are affected by both state- and industry-level factors.

First, examining the aggregate factor, Figure 9 plots the cumulated sum of aggregate factor f_t which summarizes common movements in productivity levels. Because productivity follows a random walk in the model, cumulating the factor summarizes the level component associated with successive common innovations. Different industries and locations have different exposures to this common component, as captured by their estimated factor loadings $\lambda_{i,n}$.

The dynamics exhibit a clear business-cycle pattern. The estimated common component declines around the recessions in the early 1990s and 2000s, followed by a precipitous drop in 2008–09 associated with the Global Financial Crisis. Outside those recessions, the expansions take on varying intensities, with the

Figure 9: Cumulated Aggregate Productivity Factor



Notes: The black line is the pointwise median of the cumulated aggregate factor across retained draws. Dark and light bands are the 25th–75th and 5th–95th draw percentiles. Cumulation is performed within each draw before computing percentiles. The vertical scale reflects the factor normalization.

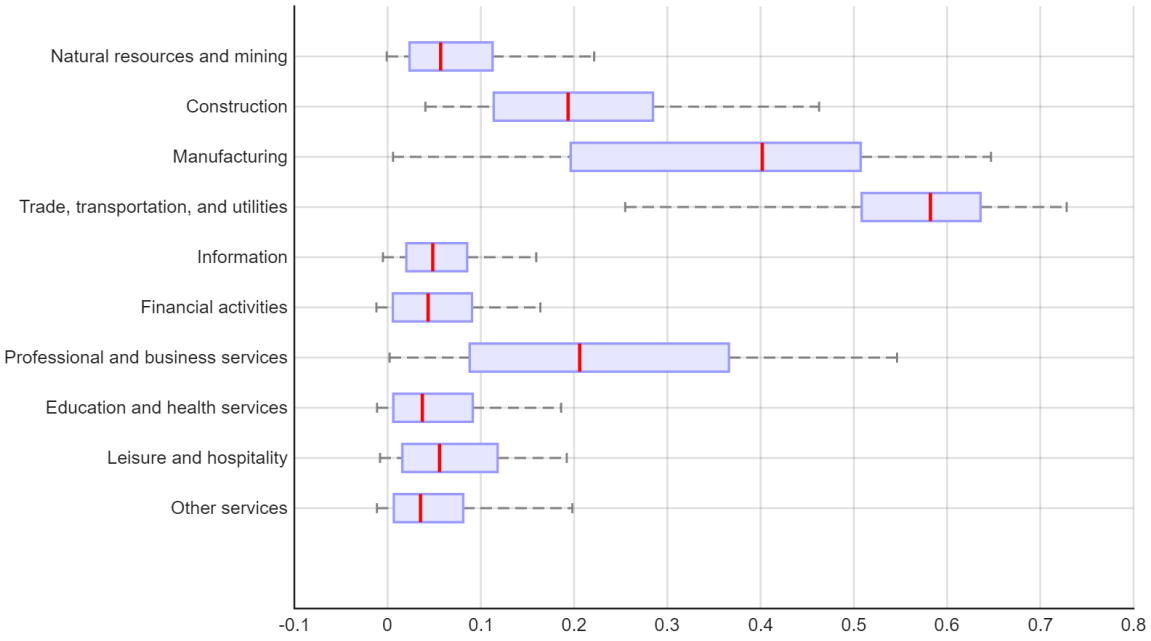
prolonged expansion of the 1990s and the relatively subdued movements of the 2010s. This correspondence establishes that the recovered common component has a recognizable cyclical pattern.

Despite the fact that the aggregate factor loads on all states and industries, state-industries vary both in their exposure to it and in how closely its contribution tracks their inferred productivity innovations. For component c, let $\varepsilon^c_{i,n,t}$ denote its contribution to the innovation in equation (25). The fit is given by $R^2_{i,n}(c) = 1 - \frac{\sum_t \left(\varepsilon_{i,n,t} - \varepsilon^c_{i,n,t}\right)^2}{\sum_t \varepsilon^2_{i,n,t}}$, where $c \in \{A, S, I, u\}$. For the aggregate component, $\varepsilon^A_{i,n,t} = \lambda_{i,n} f_t$. A value of one means that this component is responsible for the entire innovation history for that unit.

Figure 10 summarizes the cross-state variation in $R^2_{i,n}(A)$, grouped by industry. The figure shows substantial variation both across industries and across states within an industry. Each row summarizes the distribution across states within an industry. The aggregate component tracks productivity shocks relatively closely in Trade, Transportation, and Utilities, where the median R^2 is approximately 0.58, and in Manufacturing, where it is approximately 0.40. Construction and Professional and Business Services occupy an intermediate position. In the other industries, the median R^2 is much smaller. Thus, even though all state-industries can be exposed to the same aggregate factor, that factor accounts for their productivity shocks to differing degrees.

Figure 11 reports the corresponding statistics for state and industry factors. In many industries, the state factor has a low R^2 for the typical state. In some states, however, it provides a much closer fit to the estimated productivity shocks. This variation is particularly pronounced in financial activities, where the importance of the state component differs sharply across states. Industry factors also have relatively low median R^2 values, although the medians are higher in leisure and hospitality and manufacturing than in most other industries. Their importance also varies across states within the same industry. Taken together, these results show that the common sources of productivity shocks differ across states and industries. The model also allows for shocks specific to each state-industry, captured by the remaining component $u_{i,n,t}$.

Figure 10: Aggregate-Factor Fit for State-Industry Productivity Innovations



Notes: Each row summarizes the cross-state distribution of the R^2 statistic for the named industry. Boxes span the 25th–75th percentiles, red lines mark medians, and whiskers span the 5th–95th percentiles in the retained exhibit. Cross-state percentiles are computed within each draw and then summarized by their medians across draws. Boxes and whiskers describe cross-state heterogeneity.

There are thus three main takeaways from this analysis of the shocks. First, the aggregate factor captures common movements that resemble an aggregate business cycle. Second, the aggregate factor captures varying degrees of state-industry productivity shocks, and there is substantial heterogeneity both across states and industries. And third, both state and industry factors explain wide ranges of state-industry dynamics, with there significant heterogeneity.

These figures describe the sources of productivity shocks, but do not yet tell us how much each source matters for employment. Employment depends both on the shocks themselves and on how their effects spread through the economy as workers and capital adjust. We next use the model to examine the national and state employment movements generated by each shock component.

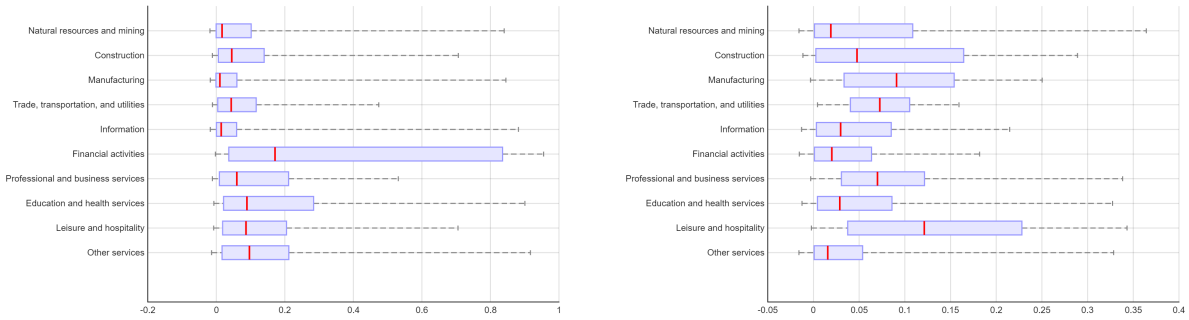
5.3 Sources of National and State Employment Fluctuations

We now turn to the core of our empirical analysis, which is using the estimated shock processes and factor structure to understand the sources of national and state employment fluctuations. As noted in Section 2, and in particular Figure 1, there is substantial variability in both national employment growth, and the cross-state dispersion in any given quarter.

To assess how each shock component affects employment, we conduct four decompositions. We calculate employment growth when productivity innovations come only from the aggregate factor, only from state factors, only from industry factors, or only from the cell-specific component. In each exercise, we set the remaining shock components to zero while holding the model's parameters and initial conditions fixed. These

Figure 11: State- and Industry-Factor Fit for Productivity Innovations

(a) State Factors **(b) Industry Factors**



Notes: Panel (a) uses $\varepsilon^S_{i,n,t} = \lambda^n_i f_{n,t}$; panel (b) uses $\varepsilon^I_{i,n,t} = \lambda^i_n f_{i,t}$. Both calculate R^2 fit and apply the conventions of Figure 10.

exercises allow each component to affect employment through all of the model's economic linkages.

Figure 12 compares observed national employment growth (shown originally in Figure 1) with the employment paths generated by these exercises. The aggregate factor produces much of the movement in national employment. It captures the declines associated with the recessions of the early 1990s and early 2000s and the much larger contraction in 2008–09. It also tracks much of the movement during the intervening expansions. Industry factors generate additional national fluctuations but at smaller scale, most notably contributing to the declines in the early 2000s and during the Global Financial Crisis. State factors and cell-specific shocks generate much smaller movements in national employment.

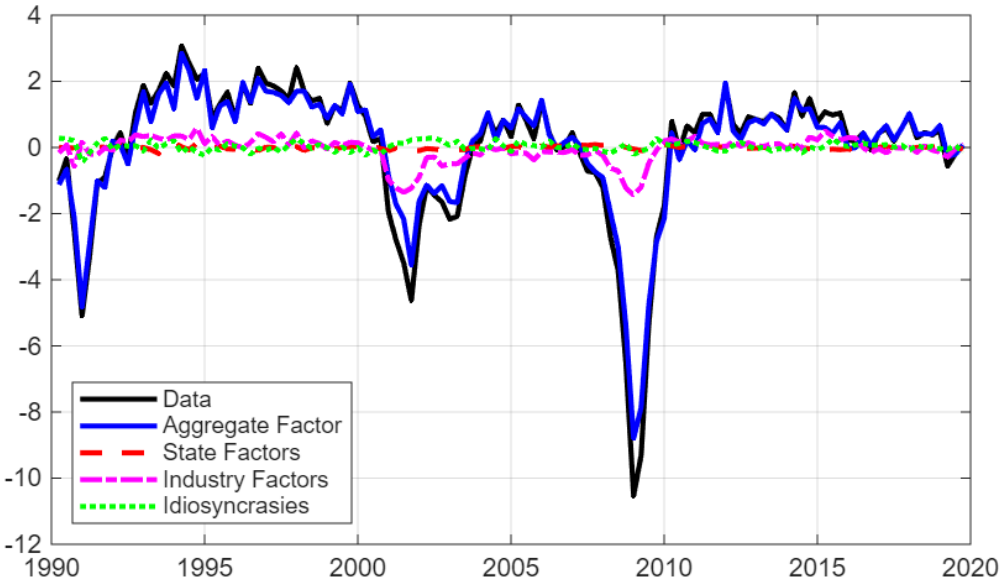
The importance of the aggregate factor for national employment does not imply that the other components are unimportant for individual states. Employment gains in some states can offset losses in others, leaving a small national movement despite substantial differences across states. We therefore next examine how much variation across states each component generates.

Figure 13 plots the cross-sectional standard deviation in each quarter; in other words, the variability across states corresponding to the shaded regions in Figure 1. The lines follow the convention in Figure 12. If the aggregate factor was the only driver, then the heterogeneity across states would typically be much lower than is present in the data. State and Industry factors by themselves generate variability across states that is slightly lower than in the data. And the idiosyncrasies specific to state-industries tend to generate much more dispersion across states.

These general trends mask some interesting differences over time however. In the recessions of the early 1990s, early 2000s, and the Global Financial Crisis, the cross-sectional standard deviation increases and is driven by notable changes in the aggregate factor. Other movements in the cross-section have interpretable movements generated by either state or industry-factors. State factors – specifically in Louisiana during Hurricane Katrina – move dispersion in 2006. Industry factors matter for the Global Financial Crisis and and an increase witnessed in the mid-2010s.

Table 1 summarizes the national and state employment movements generated by each shock component. We measure national volatility by the standard deviation of national employment growth over time (the variation in the data line in Figure 12). To measure differences across states, we calculate the cross-state standard deviation of employment growth in each quarter and average it over time (the average of the data line in Figure 13). We express both measures relative to their counterparts in the data, so that a value of one

Figure 12: National Employment Growth by Shock Component



Notes: The black line shows observed national employment growth. The colored lines show employment growth when only the indicated shock component is retained, holding the model's parameters and initial conditions fixed. Simulated paths are summarized by their pointwise medians across retained draws. Employment growth is demeaned and expressed at an annualized rate.

means that the exercise generates the same amount of volatility or dispersion as observed.

The aggregate factor generates 88 percent of the observed standard deviation of national employment growth, but only 35 percent of the observed dispersion across states. Thus, the same component that closely tracks the national cycle generates a more limited amount of cross-state variation. Its effects are nevertheless uneven: states differ both in their exposure to the aggregate factor and in how their economies respond to productivity shocks.

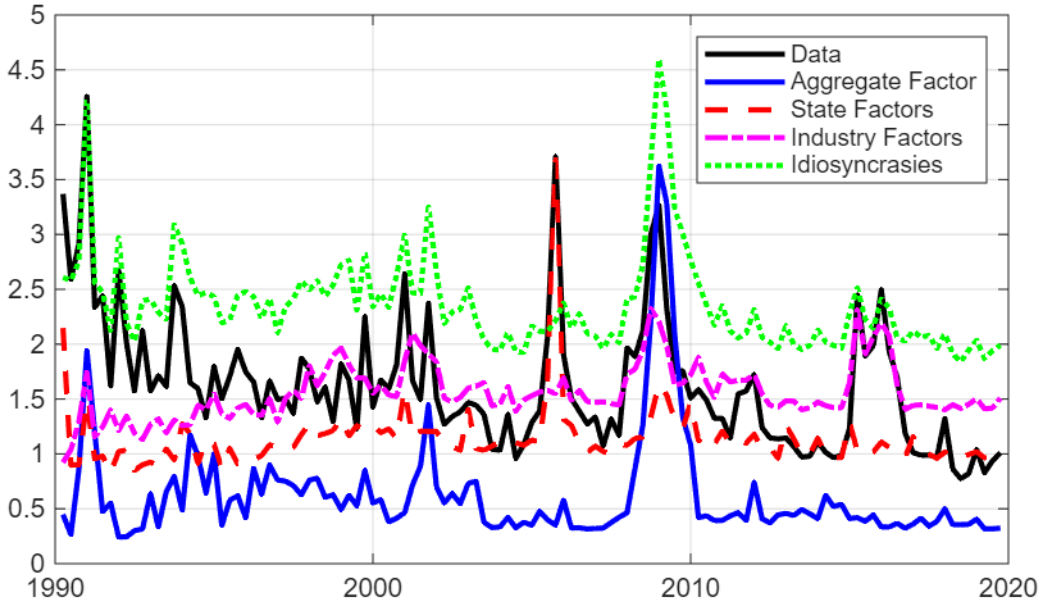
State factors display the opposite pattern. They generate only 3 percent of the observed national standard deviation, but 82 percent of observed cross-state dispersion. Their importance for individual states is therefore largely obscured in national employment. Industry factors also generate substantial differences across states: the industry-factor exercise produces 90 percent of observed cross-state dispersion, compared with 26 percent of national employment volatility.

The state-industry-factor exercise produces relatively little national volatility but more cross-state dispersion than observed in the full model: the corresponding ratios are 0.11 and 1.44.[6]

Taken together, these results show why understanding the national cycle and understanding differences across states are distinct tasks. The aggregate factor is central to national employment fluctuations, while state, industry, and cell-specific shocks are important for the uneven employment experiences of individual states.

[6]The latter does not mean that explains 144 percent of cross-state variation. It means that retaining these shocks alone generates greater dispersion than retaining all components together. The employment movements associated with different components can offset one another, so the individual ratios need not sum to one.

Figure 13: Cross-State Standard Deviations by Shock Component



Notes: The black line shows the observed cross-sectional standard deviation of state employment growth. The colored lines show the cross-sectional standard deviations when only the indicated shock component is retained, holding the model's parameters and initial conditions fixed. Simulated paths are summarized by their pointwise medians across retained draws. Employment growth is demeaned and expressed at an annualized rate.

Differences across states. The national results conceal substantial differences in how closely the aggregate factor tracks employment in individual states. Figure 14 compares each state's employment growth in the baseline model with its employment growth when only the aggregate shock component is retained. We use the same R^2 measure introduced above, but apply it to state employment growth. Higher values indicate that the aggregate-factor exercise more closely matches the state's baseline employment path.

The figure shows considerable geographic variation. The aggregate factor tracks employment relatively closely in North Carolina and South Carolina, for example, while the fit is much weaker in North Dakota and Louisiana. In Alaska and Hawaii, the fit is negative: retaining the aggregate component alone produces an employment path that is further from the baseline than a zero-growth benchmark.

These differences can arise for several reasons. States have different exposures to the aggregate factor, their economies respond differently to productivity shocks, and other shock components play different roles in their employment movements. Thus, the aggregate factor's importance for the national cycle does not translate into an equally close description of employment in every state.

The role of individual industry factors. We next examine the ten industry factors separately. In each exercise, we retain the estimated factor for one industry and set all other productivity shock components to zero. We hold model parameters and initial conditions fixed and allow the shocks to affect employment throughout the economy. Each exercise therefore measures movements in total employment, including effects outside the industry where the shocks originate. These exercises use the estimated historical shocks.

Table 2 shows that the industries whose factors generate the largest national fluctuations are not necessarily those that generate the largest differences across states. Among the industry factors, trade, transportation,

Table 1: National Employment Volatility and Cross-State Dispersion by Shock Component

Source	Relative national employment SD	Relative cross-state dispersion
Data / Full model	1.00	1.00
Aggregate factor only	0.88	0.35
State factors only	0.03	0.82
Industry factors only	0.26	0.90
Idiosyncrasies only	0.11	1.44

Notes: The first column reports the standard deviation of national employment growth relative to the data. The second reports the time average of the cross-state standard deviation of employment growth, also relative to the data. Each exercise retains only the indicated shock component, holding model parameters and initial conditions fixed. Entries are medians across retained draws. These ratios are not additive shares of employment variation.

and utilities generates the greatest national employment volatility, at 12 percent of the observed national standard deviation. Manufacturing and professional and business services follow, at 10 and 9 percent, respectively.

The ranking differs for cross-state dispersion. Manufacturing generates the largest relative dispersion in the table, followed by natural resources and mining and by education and health services. The latter two industry factors generate relatively small national fluctuations but substantial differences across states. Conversely, trade, transportation, and utilities generates less cross-state dispersion despite its greater importance for national volatility.

These results extend the distinction between national and regional fluctuations to individual industry factors. An industry's importance for national employment does not provide a reliable guide to its importance for differences across states. Both depend on the estimated shocks, states' exposures to them, and how their effects spread through the economy. We next examine the economic mechanisms that shape these responses.

6 Economic Mechanisms and Employment Fluctuations

The previous section identified the shock components associated with national employment fluctuations and differences across states. We now examine how the economy's structure shapes these outcomes. Holding the estimated productivity shocks fixed, we change individual features of the model and calculate the resulting employment paths. We do not re-estimate the shocks to fit employment under each alternative model.

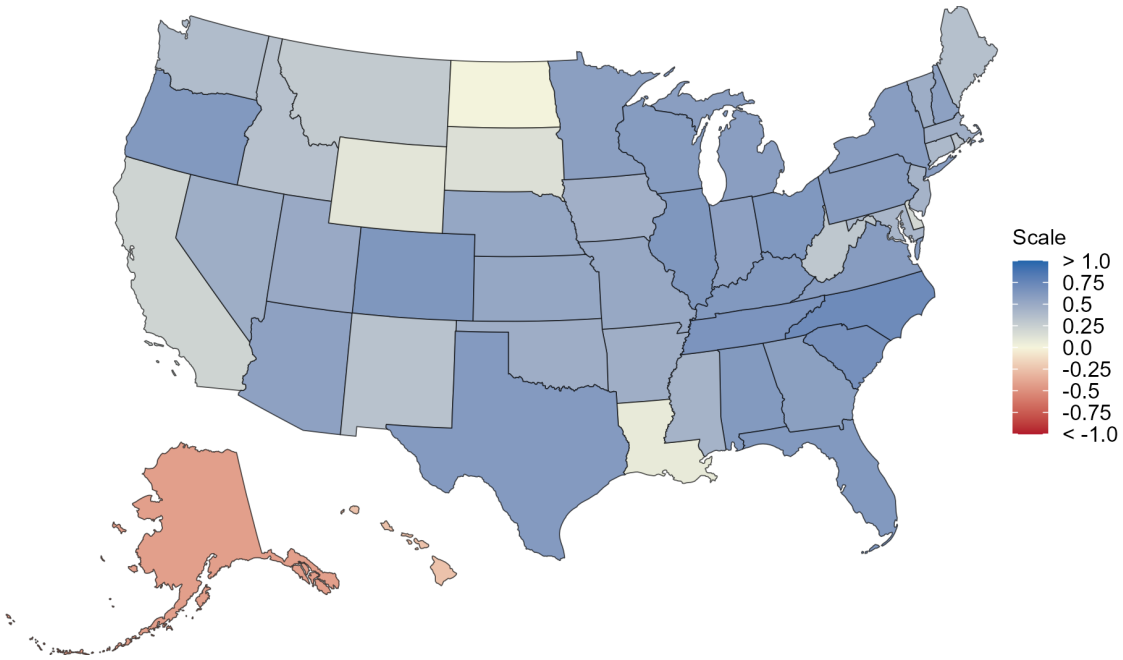
We consider five changes: reducing the role of capital in production, reducing the use of intermediate inputs, limiting interstate worker mobility, removing interstate trade, and replacing state-specific production and expenditure shares with common industry shares. Each exercise starts from the same estimated initial conditions.

Table 3 shows that these changes have relatively small effects on national employment volatility but much larger effects on differences across states. National volatility remains between 0.97 and 1.04 times its baseline value. Cross-state dispersion, by contrast, ranges from 0.56 to 1.75 times the baseline.

Reducing the role of capital raises cross-state dispersion by 47 percent. Reducing intermediate-input use raises it by 75 percent, while removing interstate trade raises it by 31 percent. Thus, in these exercises, capital, intermediate inputs, and trade each help moderate differences in employment growth across states, even though their effects on national volatility are comparatively small.

The mobility experiment produces a different result. Reducing interstate worker movements lowers

Figure 14: Aggregate-Factor Fit for State Employment Growth



Notes: Each state's shading reports the R^2 fit of employment growth in the aggregate-factor-only exercise relative to employment growth in the full model. Values are summarized by their medians across retained draws. Higher values indicate a closer fit. Negative values indicate a larger squared error than a benchmark of zero demeaned employment growth. Model parameters and initial conditions are held fixed.

cross-state dispersion by 44 percent, while leaving national employment volatility essentially unchanged. In this exercise, greater mobility therefore increases the dispersion of employment growth across states.

Replacing state-specific production and expenditure shares with common industry shares also increases cross-state dispersion, to 1.56 times the baseline.

Together, these exercises show that relatively similar national employment fluctuations can coexist with very different regional outcomes. Understanding those regional outcomes requires attention not only to the sources of shocks, but also to the economic mechanisms through which they affect employment.

7 Conclusion

In this paper, we have examined the sectoral origins of regional and aggregate business cycles. While national variation in employment growth follows a well-known business cycle dynamic, there is considerable variation in state-level outcomes in any given quarter. States have heterogeneity in both the volatility of their employment growth, and the extent to which they correlate with the nation as a whole. These findings also hold for employment growth at the industry level across states.

Using a quantitative spatial dynamic equilibrium model that captures economic mechanisms such as trade, migration, production networks, and capital accumulation, we inferred the underlying productivity shocks driving employment growth and uncovered their factor structure. An aggregate factor drives most of the movements in employment growth at the national level, yet only about a third of the state-level variation. State- and industry-specific factors account for less of the movement in the national series, but more substantial amounts of the cross-sectional variation. Trade, production networks, and capital accumulation

Table 2: Employment Effects of Individual Industry Factors

Industry factor retained	Relative national employment SD	Relative cross-state dispersion
Natural resources and mining	0.04	0.35
Construction	0.01	0.17
Manufacturing	0.10	0.41
Trade, transportation, and utilities	0.12	0.15
Information	0.03	0.14
Financial activities	0.03	0.27
Professional and business services	0.09	0.25
Education and health services	0.04	0.31
Leisure and hospitality	0.03	0.17
Other services	0.02	0.13

Notes: Each row retains only the named industry's factor, holding model parameters and initial conditions fixed. Outcomes refer to total employment across all industries. The entries compare national employment volatility and cross-state dispersion with their baseline counterparts; they are not additive shares.

all serve to mute the disparity of state outcomes.

Our quantitative framework, by matching observed data using both exogenous shocks and a rich set of mechanisms, is thus well-suited for further analysis of stabilization policy at both the national and state levels.

References

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute.

Beraja, M., E. Hurst, and J. Ospina (2019). The aggregate implications of regional business cycles. *Econometrica 87*(6), 1789–1833.

Beyer, R. C. M. and F. Smets (2015). Labour market adjustments and migration in europe and the united states: How different? *Economic Policy 30*(84), 643–682.

Blanchard, O. J. and L. F. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity 1992*(1), 1–75.

Caliendo, L., M. Dvorkin, and F. Parro (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica 87*(3), 741–835.

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2018). The impact of regional and sectoral productivity changes on the u.s. economy. *The Review of Economic Studies 85*(4), 2042–2096.

Carlino, G. and K. Sill (2001). Regional income fluctuations: Common trends and common cycles. *The Review of Economics and Statistics 83*(3), 446–456.

Chodorow-Reich, G. and J. Wieland (2020). Secular labor reallocation and business cycles. *Journal of Political Economy 128*(6), 2245–2287.

Dix-Carneiro, R., J. P. Pessoa, R. Reyes-Heroles, and S. Traiberman (2023). Globalization, trade imbalances, and labor market adjustment. *The Quarterly Journal of Economics 138*(2), 1109–1171.

Table 3: Economic Mechanisms and Employment Fluctuations

	Baseline		No Capital		No Materials		No Migration		No Trade		No Het	
Source	N.E.	X-S	N.E.	X-S	N.E.	X-S	N.E.	X-S	N.E.	X-S	N.E.	X-S
Full Model	1.00	1.00	0.98	1.47	0.97	1.75	1.00	0.56	1.04	1.31	0.97	1.56
Agg Factor	0.88	0.35	0.85	0.61	0.86	0.59	0.90	0.27	0.92	0.49	0.88	0.54
State Factors	0.03	0.82	0.04	1.02	0.04	1.44	0.03	0.34	0.02	0.99	0.03	1.12
Ind Factors	0.26	0.90	0.32	1.51	0.33	1.13	0.31	0.57	0.29	1.28	0.30	0.88
Idiosyncrasies	0.11	1.44	0.13	2.18	0.16	2.23	0.11	0.84	0.11	1.89	0.11	1.84

Notes: All exercises retain the full estimated productivity shock history. National employment volatility is measured by its time-series standard deviation. Cross-state dispersion is the time average of the quarterly cross-state standard deviation of employment growth. Both measures are expressed relative to the data. The capital and intermediate-input exercises leave small positive input shares; the mobility exercise also reduces interstate worker movements.

Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika 89*(3), 603–615.

Foerster, A. T., P.-D. G. Sarte, and M. W. Watson (2011). Sectoral versus aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy 119*(1), 1–38.

Foschi, A., C. L. House, C. Proebsting, and L. L. Tesar (2025). Inter-state labor mobility and the u.s. economy. Working Paper.

Hamilton, J. D. and M. T. Owyang (2012). The propagation of regional recessions. *The Review of Economics and Statistics 94*(4), 935–947.

King, R. G. and M. W. Watson (1998). The solution of singular linear difference systems under rational expectations. *International Economic Review 39*(4), 1015–1026.

King, R. G. and M. W. Watson (2002). System reduction and solution algorithms for singular linear difference systems under rational expectations. *Computational Economics 20*(1-2), 57–86.

Kleinman, B., E. Liu, and S. J. Redding (2023). Dynamic spatial general equilibrium. *Econometrica 91*(2), 385–424.

Owyang, M., J. Piger, and H. Wall (2005). Business cycle phases in u.s. states. *The Review of Economics and Statistics 87*(4), 604–616.

Rodríguez-Clare, A., M. Ulate, and J. P. Vasquez (2025). Trade with nominal rigidities: Understanding the unemployment and welfare effects of the china shock. Journal of Political Economy, Forthcoming.

Russ, K. N., J. C. Shambaugh, and S. R. Singh (2024). Currency Areas, Labor Markets, and Regional Cyclical Sensitivity. *IMF Economic Review 72*(1), 152–195.

Vom Lehn, C. and T. Winberry (2022). The investment network, sectoral comovement, and the changing us business cycle. *The Quarterly Journal of Economics 137*(1), 387–433.