Secular Labor Reallocation and Business Cycles*

Gabriel Chodorow-Reich
Harvard University and NBER

Johannes Wieland
UC San Diego and NBER

December 2015

Abstract

We study the effect of mean-preserving labor reallocation on business cycle outcomes. We develop an empirical methodology using a local area's exposure to industry reallocation based on the area's initial industry composition and employment trends in the rest of the country over a full employment cycle. Using confidential employment data by local area and industry over the period 1980-2014, we find sharp evidence of reallocation contributing to worse employment outcomes during national recessions but not during national expansions. We repeat our empirical exercise in a multi-area, multi-sector search and matching model of the labor market. The model reproduces the empirical results subject to inclusion of two key, empirically plausible frictions: imperfect mobility across industries, and downward nominal wage rigidity. Combining the empirical and model results, we conclude that reallocation can generate substantial amplification and persistence of business cycles at both the local and the aggregate level.

*Chodorow-Reich: Harvard University, Department of Economics, Littauer Center, Cambridge, MA 02138 (e-mail: chodorowreich@fas.harvard.edu); Wieland: University of California, San Diego, Department of Economics, 9500 Gilman Dr. #0508, La Jolla, CA 92093 (email: jfwieland@ucsd.edu). We thank for their comments David Berger, Gordon Hansen, Larry Katz, Pat Kline, Giuseppe Moscarini, Valerie Ramey, Lawrence Uren, Gianluca Violante, Ivan Werning, and numerous seminar and conference participants. This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. The views expressed here do not necessarily reflect the views of the BLS or the U.S. government. We are grateful to Jessica Helfand and Michael LoBue of the BLS for their help with the Longitudinal Database. An appendix to the paper is available on the authors' webpages.

1. Introduction

Industries can experience idiosyncratic shocks, generating changes in the distribution of employment. Whether such industry labor reallocation matters quantitatively in causing, amplifying, or propagating the business cycle has important implications for our understanding of business cycles, labor markets, and the scope for policy in affecting business cycle outcomes. Yet, the issue remains one of great debate.

In this paper we study the employment effects of secular labor reallocation, which we define as the change in an economy's allocation of labor in response to mean-preserving idiosyncratic industry shocks. We make two main contributions. First, we propose a novel method to estimate the effect of secular labor reallocation on business cycles and implement it using a confidential employment data set from the Bureau of Labor Statistics. We find large negative effects of reallocation on aggregate employment when the reallocation coincides with a national recession, but roughly neutral effects when it occurs coincident with a national expansion. Second, we build a multi-area, multi-sector search-and-matching model to interpret and extend the empirical results. Together, the model and data indicate that reallocation can be an important contributor to the amplification and persistence of business cycles at both the local and the aggregate level.

A number of empirical challenges have hampered analysis of the effect of secular labor reallocation on aggregate outcomes. First, the small number of national business cycles in periods with high frequency, high quality industry level data limit inference based only on national variation. Second, reallocation within a business cycle may reflect cyclical sensitivities that vary across industries (Abraham and Katz, 1986), and business cycles can cause permanent reallocation of inputs (Schumpeter, 1942). Finally, we generally do not observe pure reallocation shocks, i.e. mean-preserving demand or productivity shocks across industries.

We design and implement a methodology that addresses these challenges. To circumvent the small number of national business cycles, we use variation in reallocation and business cycle outcomes across broadly defined local labor markets in the United States. To address the problem of cyclical sensitivity, we divide the national business cycle into a recession-recovery and an expansion period, where by definition the national economy experiences the same level of cyclical "tightness" at the beginning and end of each subcomponent. Our metric of reallocation consists of summing the absolute value of industry employment share changes between the start and end of the recession-recovery or expansion cycle. We address the endogeneity of reallocation to local conditions by developing a measure of reallocation exposure based on a local area's initial industry composition and the pattern of industry reallocation in the rest of the country. Finally, we transform this measure of exposure to total reallocation into the part based on mean-preserving shocks by orthogonalizing it with respect to the growth rate

predicted by an area's industry composition.

We implement our exercise using confidential employment data by local area and industry from the Bureau of Labor Statistics Longitudinal Database, merged with the public use counterpart of these data, the QCEW. We use the public use version to extend the analysis back to 1979. The resulting data set tracks industry reallocation and aggregate employment in more than 200 urban local labor markets.

An example illustrates the identification. The share of workers employed in the wholesale trade sector has declined over the past thirty years, due in part to better inventory management technology, while the share in health care has expanded in response to an aging population's rising demand for health services. By contrast, the share of employment at establishments engaged in management of companies has changed relatively little. As a result, a local labor market with employment concentrated in wholesale trade and health services would have experienced greater secular reallocation than an area with low employment in these sectors but a high concentration of corporate headquarters. At the same time, the average industry demand shock across these areas would be very similar. Our exercise amounts to comparing aggregate employment outcomes in these two areas during national recession-recoveries and expansions.

We have two main findings. First, reallocation has economically important effects on an area's aggregate employment growth during a recession-recovery cycle. On average, an area with reallocation one standard deviation above the mean has employment roughly 2% lower at the end of a national recession-recovery cycle. Second, reallocation does not have a meaningful effect on aggregate employment if it occurs during an expansion. These results are statistically strong and robust to inclusion of a variety of local area time-varying control variables, inclusion of local area fixed effects, changing the metric used to partition time into recessions, recoveries, and expansions, and excluding the areas with the largest shares of employment in manufacturing, construction, or health care.

The paper's second contribution is to build a multi-sector, multi-area model of reallocation unemployment and to discipline the model's features using our empirical results. Each area in the model consists of a number of industries, each of which contains firms and workers who interact according to a search and matching framework. The shares of workers and firms in each industry depends on industry-specific productivity and consumer preferences. In line with the data, the model features two-way gross flows of workers across industries each period. The model counterpart to secular industry reallocation is an increase in the cross-sectional variance of industry-level productivities.

We conduct the same empirical exercise in the model as in the data. Specifically, we define predicted reallocation in the model based on an area's initial industry distribution and industry employment trends in the rest of the model economy, and compare outcomes across areas with industry distributions which give rise to different amounts of reallocation but experience the

same mean shock. We repeat this comparison during an "expansion" in which the increase in the cross-sectional variance of industry productivity constitutes the only set of shocks, and when we combine the industry reallocation shocks with an increase in an interest rate wedge to simulate a demand-induced recession.

With only within-industry search and matching frictions, the mean-preserving spread in industry productivities generates a small decline in employment regardless of whether it occurs coincident with a demand-induced recession or not. Incorporation of reallocation frictions and empirically plausible downward wage rigidity breaks this symmetry and produces a model response of similar magnitude to that observed in the data. Intuitively, during expansions higher wages draw job seekers into the expanding sectors, while wage compression during recessions pushes the adjustment into a larger difference in job finding rates. The dispersion in job finding rates produces reallocation unemployment, similar to the mismatch unemployment in Sahin, Song, Topa, and Violante (2014). We provide evidence of such asymmetric wage compression using hourly wage data from the CPS.

Having disciplined the model to match the cross-sectional facts from the data, we next use it to clarify the relevance of other proposed empirical signatures of reallocation. A popular account suggests that a reallocation shock must engender high wages in the growing sector and falling wages in the declining sector (see e.g. DeLong, 2010; Krugman, 2014). While strictly true in our model, the magnitude of this wage differential can be quite small. Moreover, it is precisely when this wage differential is small that the unemployment response to the reallocation shock is magnified. The model also shows that reallocation shocks can produce a decline in vacancies as unemployment increases, *contra* Abraham and Katz (1986). The increase in wages in expanding sectors limits vacancy creation there, while at the same time wage rigidity reduces the employer's share of match surplus in contracting sectors, sharply reducing vacancies in those sectors. The net result is that total vacancies fall following a reallocation shock.

We also use the model to highlight the interaction of industry reallocation shocks and labor market fluidity. Davis and Haltiwanger (2014) define worker fluidity as the sum of worker separation and hiring rates, and find in a U.S. state-year panel data set that an exogenous increase in fluidity raises aggregate employment. The same result obtains in our model. Indeed, the positive correlation holds in our benchmark calibration even conditional on an industry reallocation shock, as both fluidity and aggregate employment decline in response. More important, higher exogenous fluidity mitigates the adverse employment consequences of a reallocation shock by making industry transitions easier, suggesting more fluid labor markets may reduce the volatility of unemployment in addition to affecting the mean level.

Finally, we use the model to link our local area empirical results to national effects. We vary the size of the mean-preserving productivity shock to generate variation in reallocation at the national level. Reallocation has substantial effects on both the amplification and persistence of

national recessions in the model. Thus, our results suggest that secular labor reallocation can be an important contributor to the nature of local and aggregate business cycles.

The paper relates to literatures on the causes and consequences of input reallocation and business cycles. In an early and influential contribution, Lilien (1982) argued that sectoral shifts were responsible for much of the fluctuation in unemployment in the 1970s, a point subsequently disputed by Abraham and Katz (1986) and Murply and Topel (1987). The critiques of Lilien inform our methodological approach. Debate over the importance of sectoral reallocation has renewed in the context of the slow recoveries from the most recent two national recessions.[1]

Methodologically, our paper follows Autor, Dorn, and Hanson (2013) and Charles, Hurst, and Notowidigdo (2014) in using industry shocks to local labor markets. Autor et al. study the effects of China's export growth on U.S. commuting zones which had previously produced goods exported by China, while Charles et al. examine outcomes in MSAs experiencing large manufacturing declines. Our paper differs in its focus on business cycles outcomes. As such, we construct a measure that does not rely on a specific source of variation in sectoral reallocation, and also construct a model to interpret our findings. Our findings complement recent work on the consequences of reallocation at the worker level (Jaimovich and Siu, 2014; Fujita and Moscarini, 2013; Davis and Haltiwanger, 2014).

Our general equilibrium search-and-matching model with nominal frictions builds on Christiano, Eichenbaum, and Trabandt (2015), and earlier work by Walsh (2005). We incorporate an industry structure and labor reallocation frictions following Kline (2008), Pilossoph (2014), and Dvorkin (2014). Downward nominal wage rigidity has recently been emphasized by Schmitt-Grohé and Uribe (2013), Baqaee (2014) and Daly and Hobijn (2014), and following Hall (2005), our implementation does not violate bilateral bargaining efficiency conditions (Barro, 1977).

Section 2 explains our methodology and empirical identification strategy. Section 3 describes the employment data and our concept of local labor markets. Section 4 presents summary statistics of reallocation and local business cycles, as well as a series of specification checks of our framework. Section 5 contains the paper's core results on the effects of reallocation on business cycle outcomes. In section 6 we describe the labor market model and demonstrate its ability to match the patterns found in section 5. Section 7 concludes.

2. Measurement and Identification

Our first contribution is to design and implement a methodology to measure the effects of secular industry reallocation. Our strategy rests on three innovations. First, we measure

[1] See e.g. Groshen and Potter (2003); Koenders and Rogerson (2005); Berger (2014); Garin, Pries, and Sims (2013); Mehrotra and Sergeyev (2012) for papers which highlight the importance of input reallocation, and Aaronson, Rissman, and Sullivan (2004); Pilossoph (2014); Dvorkin (2014); Hall and Schulhofer-Wohl (2015) for an opposing view. Sahin et al. (2014) stake a middle ground using an empirical decomposition.

reallocation over a full national employment recession-recovery or expansion cycle, rather than period by period. Second, we define a local area's exposure to reallocation based on its initial industry composition and national trends. Third, we orthogonalize local reallocation exposure with respect to other labor demand shocks associated with an area's industry distribution.

2.1. Measure of Reallocation

We start by defining an index of reallocation based on the change in industry employment shares. The economy consists of A distinct areas, each with I industries. Let $e_{a,i,t}$ be employment in area a and industry i at time t, $e_{a,t} = \sum_{i=1}^{I} e_{a,i,t}$ the total employment in the area, and $s_{a,i,t} = e_{a,i,t}/e_{a,t}$ industry i's employment share. Define reallocation $R_{a,t,t+j}$ in area a between months t and $t+j$ as the scaled sum of absolute sectoral employment share changes:

$$R_{a,t,t+j} = \left(\frac{12}{j}\right)\left(\frac{1}{2}\sum_{i=1}^{I}|s_{a,i,t+j} - s_{a,i,t}|\right). \tag{1}$$

The measure $R_{a,t,t+j}$ is easily interpreted. The term in the second parentheses of $R_{a,t,t+j}$, $\frac{1}{2}\sum_{i=1}^{I}|s_{a,i,t+j} - s_{a,i,t}|$, equals zero if employment grows at an identical rate in every industry between t and $t+j$. The term is equal to one if all industries with positive employment in t disappear by $t+j$. In general, this term is between zero and one, with higher realizations indicating more reallocation. The first term, $12/j$, translates the reallocation between t and $t+j$ into a monthly flow expressed at an annual rate, such that $R_{a,t,t+j} \subseteq [0, 12/j]$. For the same area and time period, $R_{a,t,t+j}$ is additive in the level of industry aggregation. For example, $R_{a,t,t+j}$ constructed over NAICS 4 digit industries equals the reallocation across 3 digit industries plus the reallocation that occurs across 4 but within 3 digit industries. As an immediately corollary, $R_{a,t,t+j}$ is weakly increasing in the level of industry disaggregation. Finally, the measure has a natural economic interpretation when total employment does not change between the two periods, $e_{a,t} = e_{a,t+j}$, a case of particular interest in what follows. In that case,

$$R_{a,t,t+j}|_{e_{a,t}=e_{a,t+j}} = \frac{12}{j}\frac{1}{2}\sum_{i=1}^{I}\left|\frac{e_{a,i,t+j} - e_{a,i,t}}{e_{a,t}}\right| = \frac{12}{j}\frac{1}{2}\sum_{i=1}^{I}\frac{|e_{a,i,t+j} - e_{a,i,t}|}{e_{a,t}}. \tag{2}$$

Equation (2) rewrites $R_{a,t,t+j}$ as the minimum fraction of total period t employment that changes industries between t and $t+j$, expressed as a monthly flow at an annual rate.

The reallocation measure defined in equation (1) also has a close relationship to two measures of period-by-period reallocation used in previous studies. First, $R_{a,t,t+1}$ is equivalent to the sectoral dispersion measure defined in Lilien (1982) up to a first order approximation and substitution of our L1 for his L2 norm. We prefer the absolute value to the squared metric because it has less sensitivity to outliers. Second, $R_{a,t,t+1}$ is equal to the job reallocation rate

defined in Davis and Haltiwanger (1992) if total employment remains unchanged between the two periods. Appendix A contains details of these comparisons.

2.2. Empirical Challenges

We use a parsimonious model of production and demand across areas and industries to illustrate the identification issues which arise with using $R_{a,t,t+j}$ to understand how reallocation affects business cycle outcomes. The model has a similar structure to that of the model in section 6, but does not incorporate any of the labor market frictions which give rise to a causal role for reallocation on area employment. It thus provides a valid benchmark under the null of no causal effect and allows us to make precise the empirical challenges.

2.2.1. A Simple Model Each industry $i \in I$ in area $a \in A$ has a constant returns to scale production function $Q_{a,i,t} = \eta_{i,t} e_{a,i,t}$ and sells its output to a local wholesaler at price $P^Q_{a,i,t}$. The wholesaler combines the local industry output into an area-specific good $Q_{a,t}$ using the CES technology $Q_{a,t} = [\sum_i^I \tau_{a,i,t}^{\frac{1}{\zeta}} Q_{a,i,t}^{\frac{\zeta-1}{\zeta}}]^{\frac{\zeta}{\zeta-1}}$, which it sells to a national retailer at price $P_{a,t}$. All agents consume the same retail good, $Q_t = [\sum_a^A \tilde{\tau}_{a,t}^{\frac{1}{\xi}} Q_{a,t}^{\frac{\xi-1}{\xi}}]^{\frac{\xi}{\xi-1}}$. This setup allows for area-specific shocks to $\tilde{\tau}_{a,t}$, industry-specific shocks to $\eta_{i,t}$, and area-industry shocks to $\tau_{a,i,t}$. Up to a redefinition of variables, this division is almost without loss of generality; for example, all of what follows would go through with area-industry specific productivities $\eta_{a,i,t}$. Likewise, we need not specify for now how labor markets or goods market clear.

Industry and wholesaler prices freely adjust and labor can move frictionlessly across areas and industries.[2] These two assumptions ensure a common marginal revenue product $\bar{p}_t$,

$$\bar{p}_t = \left[\sum_i^I \tilde{\tau}_{i,t} \eta_{i,t}^{\zeta-1} \right]^{\frac{1}{\xi-1}}, \tag{3}$$

where $\tilde{\tau}_{i,t} \equiv \sum_b^A \tilde{\tau}_{b,t} \tau_{b,i,t} \Phi_{b,t}^{\frac{\xi-\zeta}{\zeta-1}}$ is an effective aggregate industry demand share, and $\Phi_{a,t} \equiv \sum_i^I \tau_{a,i,t} \eta_{i,t}^{\zeta-1}$ is an area-specific index of industry labor demand. Details of this and other calculations are provided in appendix B.

We present closed form solutions for $e_{a,i,t}, s_{a,i,t}$, and $e_{a,t}$:

$$e_{a,i,t} = \tau_{a,i,t} \eta_{i,t}^{\zeta-1} \tilde{\tau}_{a,t} \Phi_{a,t}^{\frac{\xi-\zeta}{\zeta-1}} \bar{p}_t^{-(\xi-1)} e_t, \tag{4}$$

$$s_{a,i,t} = \frac{\tau_{a,i,t} \eta_{i,t}^{\zeta-1}}{\Phi_{a,t}}, \tag{5}$$

$$e_{a,t} = \tilde{\tau}_{a,t} \Phi_{a,t}^{\frac{\xi-1}{\zeta-1}} \bar{p}_t^{\xi} Q_t, \tag{6}$$

[2] We provide a generalization to sticky-prices in appendix B.

from which we derive an expression for the local area employment change:

$$\ln e_{a,t+j} - \ln e_{a,t} = \frac{\xi - 1}{\zeta - 1} \ln \frac{\Phi_{a,t+j}}{\Phi_{a,t}} + \varepsilon_{t,t+j} + \varepsilon^a_{a,t,t+j}, \tag{7}$$

where $\varepsilon_{t,t+j} \equiv -\xi \left(\ln \bar{p}_{t+j} - \ln \bar{p}_t\right) + \ln Q_{t+j} - \ln Q_t$ is an aggregate shock which may be absorbed by a time fixed effect and $\varepsilon^a_{a,t,t+j} \equiv \ln \tilde{\tau}_{a,t+j} - \ln \tilde{\tau}_{a,t}$ is a local area shock.

We interpret $\{\eta_{i,t}\}$ as physical productivities, and $\{\tau_{a,i,t}\}$ as consumer taste preferences. From equation (5), changes in the joint distribution of these variables will generate within-area reallocation across industries. This setup does not, however, give a direct role to reallocation in determining local employment outcomes $e_{a,t}$, a subject we turn to in section 6. We proceed for now by using equation (7) to illustrate the issues that may arise in a regression of local area employment on reallocation.

2.2.2. Three Identification Challenges The first identification challenge arises because not all reallocation results from a mean-preserving change in the underlying fundamentals. In our simple model, a change in the distribution of $\{\eta_{i,t}, \tau_{a,i,t}\}$ which changes the index $\Phi_{a,t}$ will have a direct effect on local area employment through the $\ln \frac{\Phi_{a,t+j}}{\Phi_{a,t}}$ term in equation (7). Intuitively, some areas have more growth industries and others more declining industries, and such differences affect local employment directly even absent effects from the need to reallocate inputs. In our example, the special case of a mean-preserving change occurs if and only if $\Phi_{a,t+j} = \Phi_{a,t}$; more generally, the change in $\Phi_{a,t}$ constitutes a potential omitted variable.[3]

The second problem mirrors the critique of Lilien's period-by-period reallocation measure by Abraham and Katz (1986). Abraham and Katz point out that industries differ in their cyclical sensitivities. For example, demand for durable goods falls by more than demand for education during a recession. As a result, the employment share in durable goods production falls during recessions and the share in education increases, generating reallocation at the recession frequency. In this case, however, the business cycle causes a temporary reallocation across industries, rather than industry reallocation affecting the business cycle. Specializing our example to two industries, a decline in local demand, $\varepsilon^a_{a,t,t+1} < 0$, might differentially affect the consumer preferences $\tau_{a,1,t+1} = \tau_{a,1,t} + \beta\varepsilon^a_{a,t,t+1}$ and $\tau_{a,2,t+1} = \tau_{a,2,t} - \beta\varepsilon^a_{a,t,t+1}$. Given initial conditions $\tau_{a,1,t} = \tau_{a,2,t}$ and $\eta_{1,t} = \eta_{1,t+1} = \eta_{2,t} = \eta_{2,t+1} = 1$, we would measure positive reallocation, $R_{a,t,t+1} = 12(\frac{1}{2})\beta|\varepsilon^a_{a,t,t+1}|$, generating a negative correlation between reallocation and local employment, but with causality running from $\varepsilon^a_{a,t,t+j}$ to reallocation $R_{a,t,t+j}$.

The third challenge also concerns the direction of causality between reallocation and aggregate outcomes. A literature stretching from Schumpeter (1942) to Berger (2014) has explored

[3]Other authors have focused exactly on idiosyncratic industry or even firm shocks which do not net out in the aggregate (Long and Plosser, 1983; Gabaix, 2011). One can view such shocks as combinations of the mean-preserving shocks studied here and non-zero mean shocks.

how the state of the business cycle might affect incentives for firms to restructure and hence the reallocation of inputs. For example, a low opportunity cost of restructuring during periods of weak demand might induce reallocation to concentrate during cyclical downturns, generating a correlation between local reallocation and local employment without a causal effect of the former on the latter.

2.3. Mean-preserving Changes and Predicted Growth

A measure of predicted employment growth based on an area's industry mix addresses the concern of mean-altering changes in industry productivities and tastes. Let $e_{-a,i,t}$ denote employment in industry i at time t summing over all areas *other than* area a, $s_{-a,i,t} = \frac{e_{-a,i,t}}{e_{-a,t}}$ the employment share excluding area a, and $g_{-a,i,t,t+j} = \frac{e_{-a,i,t+j}}{e_{-a,i,t}} - 1$ the employment growth rate. We refer to these objects as the elsewhere employment share and the elsewhere employment growth rate. Following a literature beginning with Bartik (1991), we define an area's predicted monthly growth rate (expressed at an annual rate), $g^b_{a,t,t+j}$, where the superscript b refers to "Bartik", as

$$g^b_{a,t,t+j} = \frac{12}{j} \sum_{i=1}^{I} s_{a,i,t} g_{-a,i,t,t+j}. \tag{8}$$

Intuitively, $g^b_{a,t,t+j}$ predicts an area's employment growth rate between t and $t+j$ based on the area's industry employment shares at t and the common component of industry employment growth across areas. To show formally why this variable may solve the first identification challenge, it helps to write out $g^b_{a,t,t+j}$ in the case of our simple example:

$$\ln\left(1 + g^b_{a,t,t+j}\right) = \ln\left[\frac{\bar{p}^{-(\xi-1)}_{t+j} e_{t+j}}{\bar{p}^{-(\xi-1)}_{t} e_{t}}\right] + \ln \Phi^b_{a,t+j} - \ln \Phi_{a,t}, \tag{9}$$

where $\Phi^b_{a,t+j} \equiv \sum_i^I \tau_{a,i,t} \eta_{i,t}^{\zeta-1} \left(\frac{\tilde{\tau}_{-a,i,t+j} \eta^{\zeta-1}_{i,t+j}}{\tilde{\tau}_{-a,i,t} \eta^{\zeta-1}_{i,t}}\right)$ is the Bartik predicted local industry index and $\tilde{\tau}_{-a,i,t} \equiv \sum_{b \neq a}^{A} \tilde{\tau}_{b,t} \tau_{b,i,t} \Phi_{b,t}^{\frac{\xi-\zeta}{\zeta-1}}$. If each taste parameter $\tau_{a,i}$ grows at the same rate locally as nationally, then $\Phi^b_{a,t+j} = \Phi_{a,t+j}$. Using the log approximation $\ln(1 + g^b_{a,t,t+j}) \approx g^b_{a,t,t+j}$, the Bartik employment growth variable would then control exactly for the effect on local area employment of a mean-altering change in industry productivities and tastes.

2.4. Reverse Causality and Predicted Reallocation

We address the reverse causality concerns with a multipronged approach. The first and most important involves defining the reallocation analog of Bartik predicted growth as the reallocation which would obtain in the local area if employment growth in each local industry exactly followed its elsewhere counterpart. To derive such an expression, we start from an

equivalent representation of equation (1) in terms of dispersion of employment growth rates,

$$R_{a,t,t+j} = \frac{12}{j}\frac{1}{2}\sum_{i}^{I} s_{a,i,t}\left|\frac{1+g_{a,i,t,t+j}}{1+g_{a,t,t+j}} - 1\right|,$$

substitute in the numerator using $g_{-a,i,t,t+j}$ and in the denominator using $g^b_{a,t,t+j}$, and rearrange terms to yield:

$$R^b_{a,t,t+j} = \left(\frac{12}{j}\right)\left(\frac{1}{2}\sum_{i=1}^{I}\frac{s_{a,i,t}}{s_{-a,i,t}}\left|\left(\frac{1+g_{-a,t,t+j}}{1+g^b_{a,t,t+j}}\right)s_{-a,i,t+j} - s_{-a,i,t}\right|\right). \tag{10}$$

The ratio $s_{a,i,t}/s_{-a,i,t}$ in equation (10) characterizes the relative importance of industry i in area a. The first term inside the absolute value, $\left(\frac{1+g_{-a,t,t+j}}{1+g^b_{a,t,t+j}}\right)s_{-a,i,t+j}$, is the elsewhere final employment share rescaled by the ratio of local area predicted growth to elsewhere total employment growth. The rescaling forces the term to sum to one across all industries. Intuitively, an area with only one industry and where that industry is predicted to decline should have zero predicted reallocation.[4] Thus, the full term inside the absolute value is the appropriately scaled contribution to reallocation for each industry predicted by elsewhere employment. When $e_{-a,t} = e_{-a,t+j}$, a derivation similar to equation (2) shows that $R^b_{a,t,t+j}$ has the interpretation of the predicted net quantity of industry employment reshuffling between t and $t+j$ as a share of total employment at t, expressed at an annual rate.

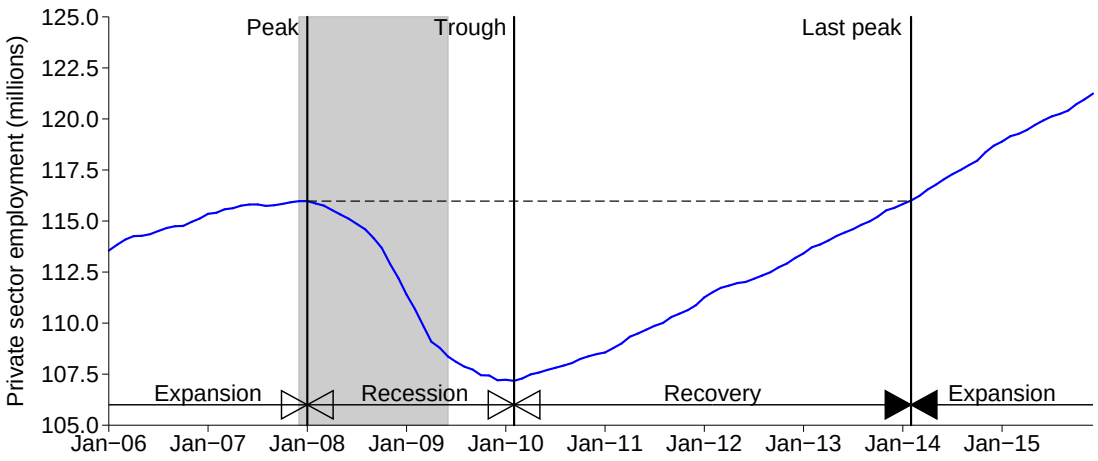
2.5. Timing

Our second innovation in response to the possibility of endogenous reallocation within a business cycle is to measure reallocation separately over full national recession-recovery cycles and expansion cycles. A full recession-recovery cycle starts at an employment peak and lasts for T periods until the economy regains the previous peak's employment level, which we call the last-peak. Figure 1 illustrates the timing for total private sector employment in the U.S. economy between 2006 and 2014. For comparison, the shaded area shows the NBER recession. Using our timing convention, the peak occurs in January 2008 at 116 million employees. The private sector regains its last-peak level of employment in February 2014, $T = 73$ months. The period from June 2005 (the last-peak from the previous recession cycle) to January 2008, and the period beginning in February 2014, are expansions.

Assumption 1 *The cyclical sensitivity of national industry employment is the same at a peak and a last-peak.*

[4]The scaling factor matters only when an area's initial industry mix predicts employment growth very different from the actual national growth rate. Such instances are rare, and all of our quantitative results remain essentially unchanged if we omit the scaling factor. We thank Ivan Werning for suggesting it.

Figure 1 – Business Cycle Timing: 2006-2015



Notes: The figure shows our business cycle timing procedure as applied to national private sector employment over 2006-2015. The shaded area indicates the NBER defined recession for comparison.

Under this assumption, our timing convention addresses the Abraham and Katz (1986) critique. Implicitly, a recession-recovery cycle consists of a temporary decline below trend, and an expansion cycle a movement along the trend, similar to the "gaps" view of business cycles advocated by DeLong and Summers (1988).[5] We present evidence supporting this timing assumption in section 4, and report sensitivity analysis to our dating procedure in section 5.

Measuring reallocation based on the snapshots of industry employment composition at the peak and last-peak effectively filters out any cyclical reallocation which occurs during a recession but reverses during the recovery. The approach thus isolates secular reallocation trends without requiring parametric time series models for either the cyclical component or the trend component of employment shares (see Brainard and Cutler, 1993; Aaronson et al., 2004; Mehrotra and Sergeyev, 2012, for articles that take the parametric time series approach). Together, national timing of peaks and last-peaks and predicted reallocation make $R^b_{a,t,t+T}$ pre-determined with respect to local outcomes following the start of the national cycle, addressing the Schumpetarian reverse causality concern.

[5]The underlying employment trend excludes population growth for three reasons. First, the predicted reallocation measure $R^b_{a,t,t+j}$ has the natural interpretation of the predicted minimum fraction of total peak employment that changes industries between the peak and last-peak when $e_t = e_{t+j}$, as it does (up to discrete time error) for $R^b_{a,t,t+T}$. Second, there exist much better high frequency measures of employment than of population at the local level. Third, we do not see an obvious alternative for how to adjust for demographic trends. For example, not only has the national employment-population ratio not recovered its pre-recession level as of the end of 2015, the peak of the series predates the 2000 recession as well. Similarly, the employment-population ratio for prime age males did not regain its previous peak following *any* downturn since 1975.

2.6. Empirical Specification and Comments

Our strategy consists of exploiting cross-sectional variation in the part of predicted reallocation $R^b_{a,t,t+T}$ orthogonal to predicted growth $g^b_{a,t,t+T}$, and where t corresponds to a peak (or last-peak) and T the elapsed time until the subsequent last-peak (or peak). We formally state two further assumptions needed for this approach to identify the causal effect of reallocation.

Assumption 2 *Growth rates of the same industry in different local areas have a common component, and the loading of the local area-industry on the common component does not depend on the local industry share.*

The first part of Assumption 2 is a relevance condition for the Bartik predicted variables. It will hold if the local area industry response to national shocks is positively correlated with the national industry response. In our example, the relevance conditions holds trivially if the $\{\eta_{i,t}\}$ are not constant. The second part of the assumption states that when industries grow they do not systematically grow more or less in areas where they already have large or small shares. Through the lens of equation (9), Bartik growth $g^b_{a,t,t+j}$ will prove an inadequate control if changes to η_i systematically prompt differential changes to $\tau_{a,i}$ depending on the initial level of $\tau_{a,i,t}$. The same issue arises with $R^b_{a,t,t+j}$ and would affect the interpretation of the effect of a unit increase in predicted reallocation. Fortunately, both parts of Assumption 2 have testable implications: Bartik growth should predict actual growth and Bartik reallocation should predict actual reallocation (part 1), and the pass through coefficient for predicted reallocation should equal 1 (part 2). We verify these implications in the data and find strong evidence in support.

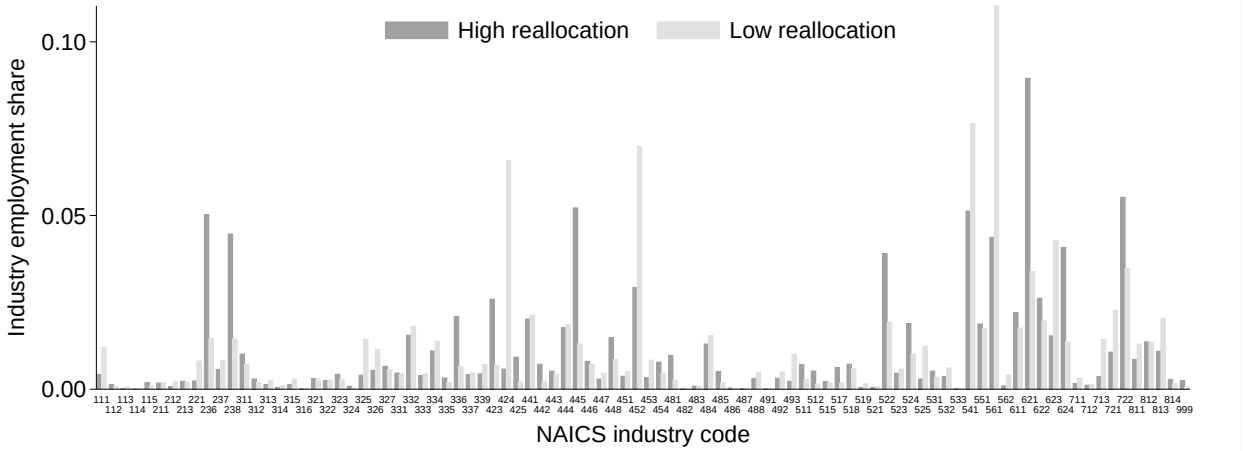
Assumption 3 *Conditional on Bartik predicted growth, area-specific shocks are uncorrelated with Bartik predicted reallocation.*

In our example, shocks to the local demand determinant $\varepsilon^a_{a,t,t+j}$ must be uncorrelated with residualized Bartik predicted reallocation.

Less formally, the two variables Bartik growth $g^b_{a,t,t+j}$ and Bartik reallocation $R^b_{a,t,t+j}$ depend only on the industry distribution in a base period t and industry employment growth in other areas between t and $t+j$. Fixing j, both variables are thus pre-determined as of time t with respect to local outcomes. Under mild conditions, $g^b_{a,t,t+j}$ has the interpretation of the average labor demand shock implied by an area's industry distribution. Therefore, the part of $R^b_{a,t,t+j}$ orthogonal to $g^b_{a,t,t+j}$ is the reallocation predicted by the area's initial industry mix but orthogonal to the direct labor demand consequences of that initial mix.

The Bartik research design has the advantage of not requiring the researcher to take a stand on the deep shocks determining reallocation in any given period, such as changes in technology, consumer tastes, exchange rates, or trade policy. We do not often directly observe $\{\eta_{i,t}\}$ and

Figure 2 – Identification Example



Notes: The bar graph displays two hypothetical industry distributions at the onset of the 2008-14 cycle. The high reallocation area has predicted reallocation during the 2008-14 cycle of .89% per year, and predicted employment growth of .001% per year. The low reallocation area has predicted reallocation during the 2008-14 cycle of .62% per year, and predicted employment growth of .001% per year.

$\{\tau_{a,i,t}\}$. Rather, the evolution of employment shares nationally summarizes the consequences of the combination of these deep shocks for reallocation, similar to the sufficient statistic approach described by Chetty (2009) for welfare economics. The Bartik approach simply requires that the deep shocks produce a common component of industry employment growth across areas, and that, after residualizing with respect to the consequences for average labor demand, these shocks affect local areas only through their effect on reallocation.

Figure 2 provides a heuristic example of the identification strategy. The figure shows two hypothetical industry distributions at the start of the 2008-14 cycle. The dark bars correspond to an industry distribution which generates R_a^b of 0.89% per year during the cycle, while the light bars give rise to R_a^b of 0.62%. That is, over the full cycle, an area with the distribution given by the dark bars would be predicted to have an additional 1.6% of employment changing industry between the national peak and last-peak. However, both the dark bars and the light bars correspond to predicted employment growth during the cycle of essentially zero. Our methodology amounts to comparing business cycle outcomes in these two areas.

3. Data

We implement our exercise in broadly defined local labor markets in the United States.

Data on employment by county and industry come from the Bureau of Labor Statistics Longitudinal Database (LDB) and Quarterly Census of Employment and Wages (QCEW). The LDB reports employment by establishment and month and covers the period 1990-2014. The source data come from quarterly reports filed by employers with state employment security agencies as part of the unemployment insurance system; as a result, the LDB contains essentially

universal coverage of private sector employment. Each establishment in the LDB has a 6 digit NAICS code associated with its primary activity. Our LDB sample contains 42 states which allow access to their data through the BLS visiting researcher confidential data access program.

The QCEW is the public use version of the LDB. It reports employment at the industry-county level for all 50 states from 1975-2014, subject to disclosure limitations to prevent the release of identifying information regarding single establishments. Even at the NAICS 2 digit level and with counties already aggregated into metropolitan statistical areas (MSAs), roughly one-fifth of potential cells get suppressed for disclosure reasons; the suppressed share rises to 35% for MSA-industry cells at the NAICS 3 digit level. Thus, analysis of reallocation requires the use of the confidential data.

Two details of the data collection procedure merit mention as they affect the periods included in our analysis. First, the Federal Unemployment Compensation Amendments of 1976 expanded the number of industries and establishments covered by unemployment insurance laws, with the result that the QCEW expanded its coverage of employment between 1976 and 1980.[6] We exclude data prior to 1978 because the staggered implementation of the coverage expansion across states produces substantial measurement difficulties during that period. In effect, we exclude the 1976-1980 expansion from the analysis. Second, in 1990 and 1991 the BLS lowered the threshold requirements for multi-establishment employers to report employment by single establishment (Farmer and Searson, 1995). As a result, an unusually high number of establishments change industry code during those years. While predicted reallocation between the 1990 peak and 1993 last-peak should remain mostly unaffected by the reclassifications as long as the changes roughly net out at the national level, actual reallocation at the local level has sufficient measurement error to render it unusable.[7]

We combine our LDB sample with NAICS 2 and 3 digit employment from the QCEW for counties in states not in the LDB, and with 2 digit SIC data for 1975-2000.[8] We seasonally adjust all series at the industry-county level using the multi-step moving average approach contained in the Census Bureau's X-11 algorithm. We also define a new SIC classification, "SIC 1.5," which groups 2 digit SIC industries into 2 digit NAICS industries using the modal employment for 2 digit SIC industries which split into multiple 2 digit NAICS industries.[9] Relative to other data sets with employment by geography and industry, such as the Census

[6]See http://www.bls.gov/cew/cewbultncur.htm#Coverage.

[7]We are grateful to Jessica Helfand and David Hiles of the BLS for helping to clarify the issues related to the 1990 and 1991 reporting change.

[8]The QCEW reports employment by county and SIC 2 digit industry beginning in 1975, and by 3 and 4 digit industry for 1984-2000. We date the 1980s expansion as beginning in October 1983, making the introduction in 1984 of the SIC 3 and 4 digit industry detail redundant for our analysis. The 1987 revision of the SIC made large changes to a handful of industry definitions which if uncorrected would result in spurious reallocation. We adjust for the classification changes by combining each of SIC 36 and 38, SIC 60 and 61, and SIC 73, 87, and 89 into a single composite industry.

[9]See appendix C for details of the construction of the SIC 1.5 classification.

Bureau's County Business Patterns or Longitudinal Business Database (LBD), the BLS data have the important advantage of providing monthly rather than annual frequency, a requirement for the timing procedure described in section 2. In what follows, we make SIC 2/NAICS 3 our baseline level of industry detail.

We aggregate county-level data into Core Based Statistical Areas (CBSAs). The Office of Management and Budget (OMB) defines CBSAs as areas "containing a large population nucleus and adjacent communities that have a high degree of integration with that nucleus," and distinguishes between Metropolitan (MSA) and Micropolitan (MiSA) areas depending on whether the urban core contains at least 50,000 inhabitants. We further aggregate CBSAs into Combined Statistical Areas (CSAs).[10] CSAs consist of adjacent CBSAs that have "substantial employment interchange," and thus better capture the local labor market. Not all CBSAs belong to a CSA. For example, the San Diego MSA is not part of a CSA, but the Boston-Cambridge-Newton MSA is one of five MSAs in the Boston-Worcester-Providence CSA.

Our final sample includes all MSAs and CSAs containing at least one MSA, with employment of at least 50,000 in one month, an agricultural share of employment of less than 20%, and where we observe at least 95% of employment at the industry level.[11] The final sample contains 1,314 of the 3,144 counties in the United States and covers 86% of 2013 employment.

4. Summary Statistics

4.1. Trends in National Reallocation

We begin with an overview of reallocation at the national level. Table 1 reports national full cycle reallocation for each recession-recovery and expansion, and at six levels of industry aggregation. The shaded rows indicate the recession-recovery episodes. Here and elsewhere, we group together the 1980 recession-recovery, 1981 expansion, and 1981-1983 recession-recovery cycles into a single episode.[12] We measure reallocation using SIC definitions for the episodes between 1975 and 2000, and using NAICS definitions for the episodes beginning after 1990. It helps to group SIC "1.5" with NAICS 2, SIC 2 with NAICS 3, and SIC 4 with NAICS 6, based on similarity in the number of industries. Indeed, reallocation measures for the overlapping episodes of the March 1990-April 1993 recession-recovery cycle and the April 1993-December

[10] We use the 2013 OMB county classifications of CBSAs and CSAs for our entire sample to avoid discontinuities from counties switching CBSAs.

[11] We exclude areas with a large agricultural share because of the particular difficulty of seasonally adjusting agricultural employment. The 95% coverage restriction binds because of disclosure limits in CSAs/MSAs located at least partly in states not in our LDB sample.

[12] We do so for a number of reasons. Measured separately, each of these episodes is much shorter than the recession-recovery and expansion cycles that follow. Grouping them together generates a cycle 6 months longer than the 1990-1993 recession-recovery cycle, and 11 months shorter than the 2000-2005 recession-recovery cycle. At the local level, many areas do not have a last-peak between the two recessions. Results are qualitatively robust to this choice.

Table 1 – Reallocation by Episode and Industry Detail

Epsiode	Months	Expansion	Industry definition						
			SIC 1.5	NAICS 2	SIC 2	NAICS 3	NAICS 4	SIC 4	NAICS 6
Mar80-Oct83	43	No	1.14		1.29			1.68	
Oct83-Mar90	77	Yes	0.71		0.93				
Mar90-Apr93	37	No	0.82	1.02	0.97	1.14	1.30	1.34	1.54
Apr93-Dec00	92	Yes	0.42	0.60	0.85	0.77	0.95	1.14	1.13
Dec00-Jun05	54	No		0.80		0.96	1.23		1.40
Jun05-Jan08	31	Yes		0.61		0.74	1.02		1.22
Jan08-Feb14	73	No		0.67		0.74	0.91		1.06
$R^2: \frac{1}{j}\lvert\Delta s_{i,t,t+j}\rvert = \alpha_i + \varepsilon_{i,t,t+j}$			0.72	0.72	0.73	0.59	0.63	0.75	0.64
Industry count			18	20	73	92	305	963	1028

Notes: The table reports values of $R_{us,t,t+j}$ for all complete national recession-recovery and expansion cycles between 1975 and 2014, and at varying levels of industry detail. The table omits the entry for SIC 4 between 1983 and 1990 because of the SIC classification revision in 1987.

2000 expansion cycle appear roughly comparable across these definitions, validating the groupings and facilitating comparison across time and classification.

A number of interesting patterns emerge. First, the rate of reallocation during episodes containing recessions has trended down. We find 4.6% (1.29*43/12) of employment changed SIC 2 digit industry between the March 1980 private sector employment peak and the October 1983 last-peak. The same fraction changed NAICS 3 digit industry between the January 2008 peak and the February 2014 last-peak, despite the latter episode lasting 30 months longer. As a result, monthly reallocation fell from 1.29% (at an annual rate) during the 1980-83 episode to 0.74% during the 2008-14 episode. The decline in between is monotonic. Despite the widespread attention to industry reallocation during the 2008-2014 episode, our measure of secular reallocation suggests a decline in reallocation intensity during the Great Recession (see also Foster, Grim, and Haltiwanger, 2014).

The second interesting pattern comes from contrasting reallocation during the recession-recovery and expansion cycles. Beginning with Schumpeter (1942), a long tradition has viewed recessions as periods with low opportunity cost of reallocation, with the implication that reallocation will concentrate during downturns. More recently, Jaimovich and Siu (2014) find that the disappearance of employment in occupations in the middle of the skill distribution has occurred almost entirely during recessions. In contrast, Caballero and Hammour (2005) argue that recessions reduce cumulative reallocation across establishments within manufacturing. What does our approach say? At the national level, substantial industry reallocation has occurred during both expansion and recession-recovery cycles. However, slightly more secular

reallocation concentrates during episodes containing recessions.

Third, comparing reallocation measures using the same industry classification and for the same recession reveals the monotonicity property in aggregation level discussed in section 2.1. For example, of the 6.5% (1.06% per year multiplied by 6.08 years) of employment changing 6 digit NAICS industry between the January 2008 peak and the February 2014 last-peak, 4.1 p.p. constituted movement across 2 digit industries, 0.4 p.p. movement within 2 digit but across 3 digit industries, 1.0 p.p. movement within 4 digit but across 3 digit industries, and 0.9 p.p. movement within 4 digit but across 6 digit industries.

Fourth, while individual industries exhibit persistence in their contribution to national reallocation, the explanatory power of this relationship lies well below one. We establish this fact by reporting in the penultimate line of the table the R^2 from the regression:

$$\frac{1}{j}|\Delta s_{i,t,t+j}| = \alpha_i + \varepsilon_{i,t}. \tag{11}$$

For example, at the NAICS 3 level, the R^2 of this regression equals 0.59. Thus, individual industry trends leave unexplained 40% of the variation in the contribution of industries to national reallocation. This time variation in industry employment trends will in turn contribute to substantial variation over time in the predicted reallocation in individual areas.

We next evaluate our timing convention using national share changes. As explained above, the division of calendar time into peak-to-last-peak, which we call recession-recovery, and last-peak-to-peak, which we call expansion, addresses the Abraham and Katz (1986) critique if industries exhibit the same cyclical sensitivity at a peak and last-peak. Otherwise, industries will systematically gain or lose share during a recession-recovery depending on their cyclicality. Table 2 presents formal diagnostic statistics based on the following specification test:

$$\frac{1}{j}\Delta s_{i,c(t),x(t),t,t+j} = \alpha_{i,c(t)} + \beta_i \mathbb{I}\{x = \text{recession-recovery}\} + \varepsilon_{i,tc(t),x(t),t,t+j}. \tag{12}$$

The dependent variable in equation (12) is the average monthly change in industry i's employment share during either the recession-recovery or the expansion part of the national last-peak-to-last-peak cycle c. We allow non-parametrically for changes in trend share growth over time by saturating the regression with a full set of industry-by-last-peak-to-last-peak cycle fixed effects $\{\alpha_{i,c(t)}\}$. The variable $\mathbb{I}\{x = \text{recession-recovery}\}$ is an indicator for whether the period t to $t+j$ corresponds to a recession-recovery or an expansion. The coefficients $\{\beta_i\}$ estimate the average recession-recovery share change premium for each industry. These coefficients should all equal zero if the peak and last-peak occur at the same cyclical point for each industry.

Column (1) of table 2 reports summary results for NAICS 3 digit industries, grouping the 1993-2000 expansion with the 2000-2005 recession-recovery, and the 2005-2008 expansion with the 2008-2014 recession-recovery. The p-value for the joint F-test of $\beta_i = 0\ \forall i$ does not reject at

Table 2 – Equality Tests of Loading: Recession, Recovery and Expansion

	Reces.-Recov. v. Expans.			Reces. v. Recov.		
	NAICS 3	SIC 2	SIC/ NAICS	NAICS 3	SIC 2	SIC/ NAICS
	(1)	(2)	(3)	(4)	(5)	(6)
Equality test of loading:						
Joint test p-value	0.21	0.00	0.02	0.00	0.00	0.00
P-value ex. construction, health	0.71	0.44	0.95	0.00	0.00	0.00
Last-peak-to-last-peak X industry FE	Yes	Yes	Yes	No	No	No
Reces.-recov. X industry indicator	Yes	Yes	Yes	No	No	No
Peak-to-last-peak X industry FE	No	No	No	Yes	Yes	Yes
Recession X industry indicator	No	No	No	Yes	Yes	Yes
R^2	0.86	0.94	0.90	0.91	0.94	0.89
Number of industries	92	73	18	92	73	18
Observations	366	288	144	549	288	144

Notes: In columns (1)-(3), the dependent variable is the monthly share change during a recession-recovery or expansion episode. The sample includes the 1976-80 expansion grouped with the 1980-83 recession-recovery, the 1983-90 expansion grouped with the 1990-93 recession-recovery, the 1993-2000 expansion grouped with the 2000-05 recession-recovery, and the 2005-08 expansion grouped with the 2008-14 recession-recovery. Equation (12) provides the estimating equation. In columns (4)-(6), the dependent variable is the monthly share change during a recession or recovery episode. The sample includes the 1980-82 recession grouped with the 1982-83 recovery, the 1990-92 recession grouped with the 1992-93 recovery, the 2000-03 recession grouped with the 2003-05 recovery, and the 2008-10 recession grouped with the 2010-14 recovery. Equation (13) provides the estimating equation. The p-value reports the p-value from a joint hypothesis test that the coefficients interacting the industry categorical variables with the recession-recovery (columns 1-3) or recession (columns 4-6) indicator all equal zero. Construction is SIC 0C17 and NAICS 238. Health is SIC 0I80 and NAICS 62.

the 10% level. Column (2) reports results for SIC 2 digit industries. Here the p-value rejects at the 5% level. However, only two sectors account for the rejection: specialty trade contractors (NAICS 238), and health care (NAICS 621, 622, and 623). The test fails to reject for all other sectors, with a p-value of 0.44. Column (3) combines NAICS 2 digit industries with our SIC 1.5 industries. The sample period for column 3 spans 1976 to 2014 and includes 4 full last-peak-to-last-peak episodes. Despite the concomitant increase in power, excluding construction and health the data do not reject equality with zero, with a p-value of 0.95.

What about the construction and health sectors? In appendix D, we plot the time series of employment shares in these two sectors. Whether the result of having a finite sample or particular structural loadings, these sectors exhibit marked differences between recession-recoveries and expansions irrespective of exactly where one defines the last-peak. We assess the sensitivity of our results to the behavior of these two sectors in robustness exercises.

Columns (4)-(6) of table 2 repeat the exercise in equation (12), but after removing all expansion periods from the sample and defining $\alpha_{i,c(t)}$ as an industry-by-recession-recovery

fixed effect:

$$\frac{1}{j}\Delta s_{i,c(t),x(t),t,t+j} = \alpha_{i,c(t)} + \beta_i \mathbb{I}\{x = \text{recession}\} + \varepsilon_{i,tc(t),x(t),t,t+j}. \tag{13}$$

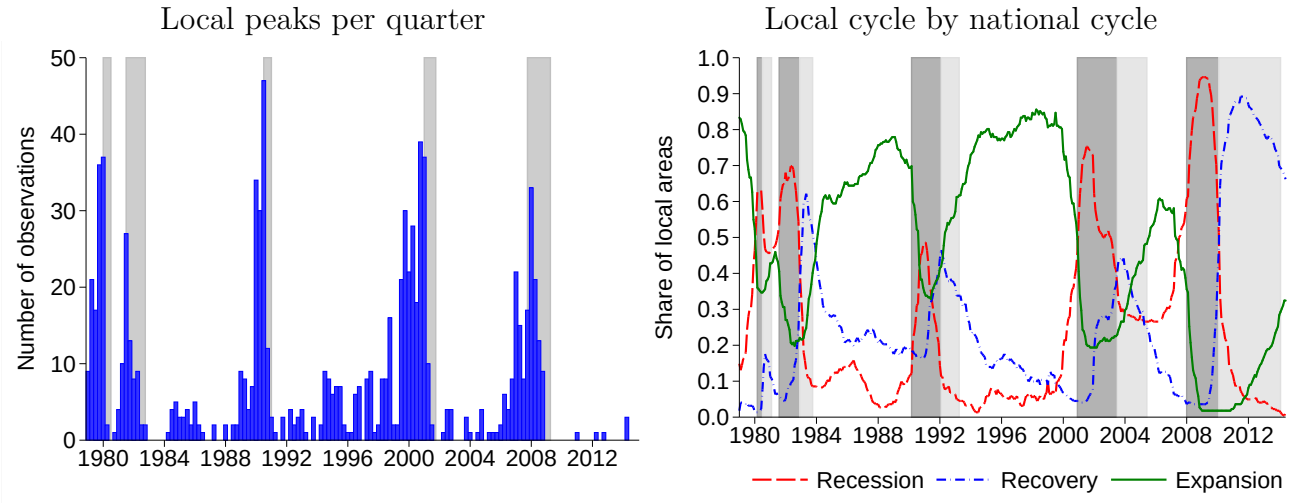
The $\{\beta_i\}$ now assess whether cyclical differences exist between recessions and recoveries, as suggested by Abraham and Katz (1986). We find that they do. The data reject $\beta_i = 0\ \forall i$ at all conventional confidence levels, even excluding the construction and health sectors. This suggests that the failure to reject in columns (1)-(3) does not reflect low power, but instead provides support for the timing convention we adopt as our baseline.

4.2. Local Business Cycles

We now turn to a comparison of the timing of local and national business cycles. This comparison matters to the interpretation of our empirical results contrasting national recession-recovery and expansion episodes insofar as the national cycle may not coincide with the local cycles. The pattern of local cycles and their relation to national cycles also has independent interest for a number of questions in regional economics.

We identify 915 local employment peaks in CSA/MSAs between 1979 and 2014.[13] The left panel of figure 3 displays their calendar frequency. Local peaks cluster around national business cycle peaks; more than three-quarters of the local peaks occur in the six quarters preceding or during an NBER recession.

Figure 3 – Correlation of Local and National Business Cycles



Notes: In the left panel, the shaded areas indicate NBER recessions. The panel shows the number of local recessions beginning in each quarter. In the right panel, the dark shaded area indicates a national recession, and the light shaded area indicates a national recovery. The panel shows the fraction of local areas in the indicated business cycle state.

[13] To avoid misattributing statistical noise or temporary disruptions such as strikes to cyclical causes, for local cycles we require that the period of the recession last at least 6 months and that employment at the trough has fallen by at least 1% and 2000 employees from employment at the peak.

The right panel of figure 3 reports time series of the fraction of local areas in a local recession, recovery, or expansion. The dark shaded areas indicate national recessions, and the light shaded areas indicate national recoveries. For this plot only, we show separately the 1980-81 and 1981-83 recession-recovery cycles. Perhaps not surprisingly, the fraction of local areas in recession spikes during national recessions, and conversely for expansions. The fraction in recession remains around 0.5 even during the brief 1981 expansion, providing support for our decision to group the episodes together. The relatively high share of areas in recession during the 2005-08 expansion reflects areas such as the Detroit-Warren-Ann Arbor CSA which never recovered from the sharp manufacturing losses imposed by the 2000 recession. Strikingly, fewer than 10% of areas escaped recession during the 2008 downturn.

Taken together, the evidence in figure 3 points to a strong correlation between national and local cycles, echoing earlier findings of comovement at the state level (Blanchard and Katz, 1992; Hamilton and Owyang, 2012).

4.3. Variation in Predicted Reallocation

Figure 4 shows a map of the variation in predicted reallocation during the 2008-2014 recession-recovery cycle.[14] We split the MSA/CSA observations into quintiles based on their Bartik reallocation, and mark higher reallocation levels with darker shades of red. Note that all CBSAs belonging to an observation have the same color. The map shows that predicted reallocation is not easily explained by geographic factors. Indeed, many of the areas with high predicted reallocation border areas with low predicted reallocation.

Table 3 reports the pairwise correlations in predicted local reallocation across each national recession-recovery and expansion. A consistent pattern does not emerge. Bartik predicted reallocation has a positive correlation across some national recession-recoveries and a negative correlation across others. The absence of strong serial correlation again helps in expanding the available variation. Nonetheless, in what follows we cluster all standard errors by CSA/MSA to account for arbitrary correlation within a CSA/MSA over time.

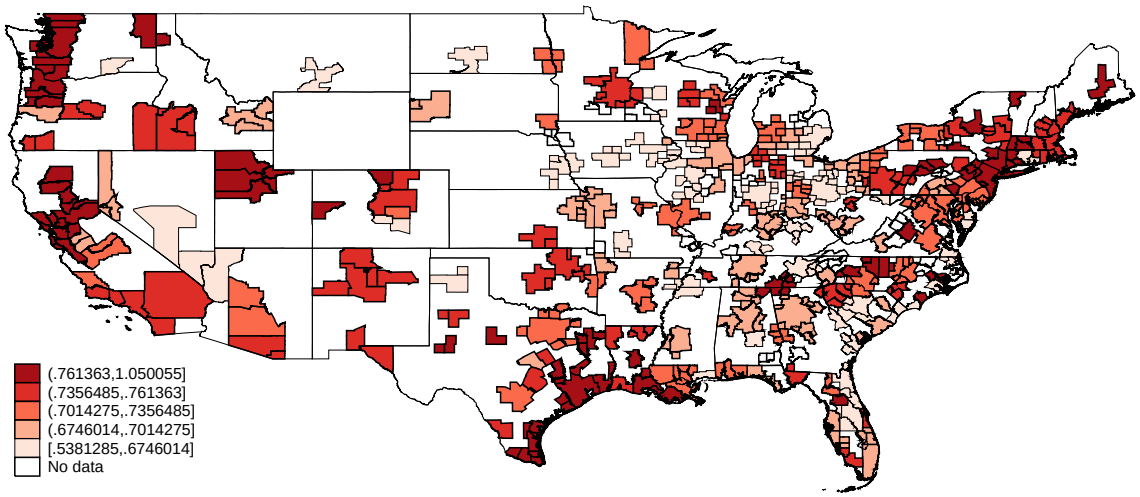
4.4. Predicted and Actual Reallocation

For predicted reallocation to matter to local outcomes, it must actually predict realized reallocation, as formalized in assumption 2. Here we assess this first stage relationship.

Table 4 reports regressions of actual on predicted reallocation, separately for recession-recovery and expansion cycles. For this table only, the recession-recovery cycles exclude the 1990-93 period. As discussed in section 3, the substantial industry reclassification of establishments during that period generates sufficient measurement error to render it unusable. Columns (1) and (4) contain our minimum set of control variables: national cycle fixed effects, predicted

[14]For data confidentiality reasons, the map uses only the public-use QCEW data. Greater disclosure limitations prior to 2008 make it impossible to report maps at the same level of industry detail for earlier cycles.

Figure 4 – Map of Predicted Reallocation Per Year, 2008-2014 Cycle



Notes: the figure shows the geographic distribution predicted reallocation per year for the national employment peak in January 2008. Due to disclosure limitations, for this figure only we use only data from the public-use QCEW and require a minimum industry employment coverage of 80%.

Table 3 – Correlation of Predicted Reallocation Across Episodes

	1980-1983	1983-1990	1990-1993	1993-2000	2000-2005	2005-2008	2008-2014
1980-1983	1.00						
1983-1990	0.24	1.00					
1990-1993	0.21	0.09	1.00				
1993-2000	−0.02	0.52	−0.16	1.00			
2000-2005	0.30	0.26	0.07	0.43	1.00		
2005-2008	0.17	0.35	0.04	0.43	0.52	1.00	
2008-2014	0.10	0.15	0.05	0.26	0.49	0.71	1.00

employment growth over the cycle, and a lag of population growth. Columns (2) and (5) add a number of time-varying area-specific controls.[15] Columns (3) and (6) replace the additional control variables with location fixed effects.

Across specifications, predicted reallocation has strong explanatory power. Treating these as first stage regressions, in most specifications the instrument would clear conventional thresholds for absence of a weak first stage.

The point estimates relating predicted to actual reallocation also cluster around one, with implications for the validity of using local variation to infer effects at a national level. If the

[15]These include: lagged house price growth; lagged employment growth; area size, measured by the log of sample mean employment; the Herfindahl of industry concentration at the cycle start; and for the recession-recovery cycles, the predicted employment decline at the national recession trough, based on peak employment shares. This last variable gives a measure of the cyclicality of an area's industry mix. Appendix E provides further description of the variables and reports partial correlations with predicted reallocation.

Table 4 – First Stage Regressions

	Dependent variable: actual reallocation					
	Recession-recovery cycles			Expansion cycles		
	(1)	(2)	(3)	(4)	(5)	(6)
Right hand side variables:						
Predicted reallocation	1.03**	1.06**	1.04**	0.64**	1.28**	0.99**
	(0.24)	(0.24)	(0.34)	(0.24)	(0.16)	(0.26)
National cycle FE	Yes	Yes	Yes	Yes	Yes	Yes
Predicted growth, lagged population	Yes	Yes	Yes	Yes	Yes	Yes
Other time-varying controls	No	Yes	No	No	Yes	No
Geographic area FE	No	No	Yes	No	No	Yes
Predicted reallocation partial F stat.	19.1	19.9	9.7	7.2	67.0	14.9
R^2	0.43	0.63	0.81	0.24	0.53	0.76
CSA-MSA clusters	218	218	218	218	218	218
Observations	534	534	534	557	557	557

Notes: The dependent variable is actual reallocation, R_a. The variable predicted reallocation is the reallocation measure R_a^b. The recession-recovery sample excludes the 1990-1993 recession cycle. The other time-varying controls include lagged house price growth; lagged employment growth; area size, measured by the log of sample mean employment; the Herfindahl of industry concentration at the cycle start; and for the recession-recovery cycles, the predicted employment decline at the national recession trough. Standard errors in parentheses and clustered by CSA-MSA. $+, *, **$ indicate significance at the 0.1, 0.05, and 0.01 levels, respectively.

expansion or shrinkage of an industry occurs disproportionately in areas that have a high initial employment share, then predicted reallocation would translate into much higher local reallocation. Conversely, if industry trends disproportionately concentrate in areas with low initial employment share in those industries, then predicted reallocation would overestimate local reallocation. The first stage coefficient of one indicates that on average national reallocation maps one-for-one into local reallocation, irrespective of the local employment share. Thus, local areas provide appropriate microcosms of the national economy in line with assumption 2.

In what follows, we report reduced form rather than second stage results for two reasons. First, otherwise we would have to exclude the 1990-93 cycle from the analysis. Second, a first stage coefficient of one with a large partial F statistic suggests that the reduced form and second stage coefficients would have a similar magnitude in any case. Indeed, we find very similar coefficients in unreported second stage results excluding the 1990-93 cycle. We therefore proceed with reduced form regressions of business cycle outcomes on predicted reallocation.

5. Empirical Results

5.1. Baseline Recession-Recovery Results

We begin by establishing a negative effect of predicted reallocation on aggregate employment during a national recession-recovery cycle. Table 5 reports OLS regressions of the form:

$$y_{a,t,t+j} = \alpha_t + \beta_1 R^b_{a,t,t+T} + \beta_2 g^b_{a,t,t+T} + \gamma' X_{a,t} + \varepsilon_{a,t,t+j}, \tag{14}$$

where $y_{a,t,t+j}$ is an outcome in area a at between the time t and $t+j$. The right hand side variable of interest, $R^b_{a,t,t+T}$, measures the predicted monthly flow of secular reallocation over the course of the national recession-recovery, based on the area's initial industry distribution and the evolution of industry employment in the rest of the country. All specifications control for at a minimum the Bartik predicted growth in the area over the course of the national recession-recovery, the area's population trend $\Delta \ln l_{t-60,t-12}$, and national recession-recovery cycle fixed effects. The table reports two outcomes: an area's annualized employment growth from the beginning to the end of the national cycle; and an area's employment growth in the four years following a national peak in employment. While the average national recession-recovery cycle lasts approximately four years, these measures will differ because of the heterogeneity across national cycles in the cycle length. In the regressions where the dependent variable is actual growth in the four years following a national peak, we always also control for predicted growth over the four year horizon, $g^b_{a,t,t+48}$. We report standard errors clustered by CSA/MSA.

Table 5 shows a negative effect of reallocation on employment growth during a recession-recovery cycle. Each column in each panel restricts in a different way the type of variation in reallocation exposure exploited. Columns (1) and (5) present the most parsimonious specifications, including only the control variables just described. These variables in fact absorb substantial variation; in particular, the predicted growth rate in column (1) has a coefficient of 1.19 and a standard error of 0.21, such that the data do not reject a coefficient of one for predicted growth but strongly reject a coefficient of zero. Columns (2) and (6) add the control variables described in appendix E: lagged house price growth, lagged employment growth, area size, the industry Herfindahl at the employment peak, and the predicted employment decline at the national recession trough. Here the point estimate falls but only slightly, while the R^2 rises by 16 p.p. reflecting the explanatory power of the control variables for the outcome variable.[16]

Columns (3) and (7) present regressions which control non-parametrically for the Bartik predicted growth rate by including episode-specific indicator variables for belonging to each of twenty quantiles of predicted growth. Similar to the identification illustrated by figure 2, these columns compare employment growth across areas with different predicted reallocation but in the same vigintile of predicted growth. The absolute value of the point estimate for reallocation

[16] We provide the corresponding table showing the coefficients and standard errors on the control variables in appendix E.

Table 5 – Effects of Reallocation on Employment During Recession-Recovery Cycles

	Dependent variable:							
	$\frac{12}{T}\ \Delta \ln e_{t,t+T}$				$\Delta \ln e_{t,t+48}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Right hand side variables:								
Predicted reallocation	−3.13**	−2.65**	−3.23**	−2.85**	−11.7**	−10.3**	−13.2**	−11.0**
	(0.66)	(0.65)	(0.84)	(0.79)	(2.63)	(2.60)	(3.44)	(3.04)
Predicted growth	1.19**	0.95*	1.04	1.64**	−1.30	−0.67	−1.35	−7.31*
	(0.21)	(0.43)	(0.78)	(0.34)	(2.68)	(2.52)	(4.26)	(3.29)
$\Delta \ln l_{t-60,t-12}$	0.068**	0.023	0.068**	−0.14**	0.20**	0.022	0.21**	−0.56**
	(0.017)	(0.026)	(0.016)	(0.035)	(0.065)	(0.10)	(0.061)	(0.14)
$g^b_{a,t,t+48}$					1.57*	1.43+	1.28+	3.81**
					(0.63)	(0.78)	(0.68)	(0.82)
National cycle FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other controls	No	Yes	No	No	No	Yes	No	No
Bartik growth quantiles	No	No	Yes	No	No	No	Yes	No
Geographic area FE	No	No	No	Yes	No	No	No	Yes
Pred. reallocation mean	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Pred. reallocation s.d.	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22
Dep. var. mean	0.46	0.46	0.46	0.46	1.66	1.66	1.66	1.66
Dep. var. s.d.	1.80	1.80	1.80	1.80	8.10	8.10	8.10	8.10
R^2	0.33	0.49	0.42	0.61	0.48	0.60	0.55	0.69
CSA-MSA clusters	219	219	219	219	219	219	219	219
Observations	748	748	748	748	748	748	748	748

Notes: The dependent variable is indicated in the table header. $\frac{12}{T}\ \Delta \ln e_{t,t+T}$ is the monthly employment growth rate during the national recession-recovery cycle, expressed at an annual rate. $\Delta \ln e_{t,t+48}$ is the log change in employment over the four years following the national peak. The other time-varying controls include lagged house price growth; lagged employment growth; area size, measured by the log of sample mean employment; the Herfindahl of industry concentration at the cycle start; and the predicted employment decline at the national recession trough. Standard errors in parentheses and clustered by CSA-MSA.

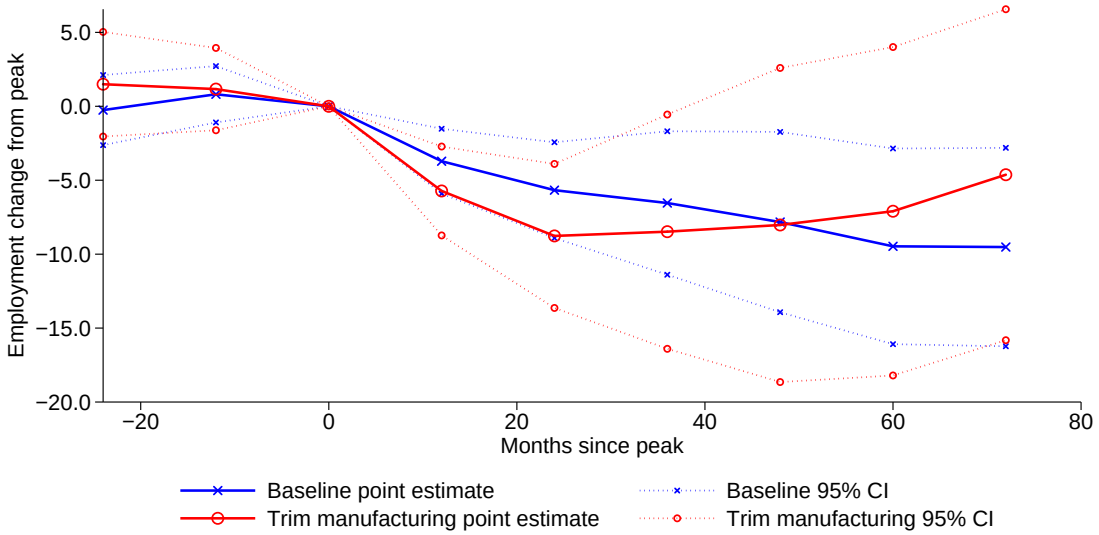
rises slightly under this specification.

Columns (4) and (8) substitute area fixed effects for the area-specific control variables.[17] Inclusion of both area and time fixed effects restricts the variation in predicted reallocation to coming from within a CSA/MSA and relative to the national mean. The point estimates for the effect of reallocation again remain similar. In all specifications, the data reject no effect of reallocation on outcomes at the 1% level.

The point estimates translate easily into economic magnitudes. In the first panel, a marginal 1 p.p. increase in predicted reallocation per year results in employment growth 3 p.p. slower per year during the course of the recession-recovery cycle. Predicted reallocation has a sample

[17]We substitute rather than augment because the control variables include a lag of the dependent variable.

Figure 5 – Local Projection



Notes: The figure plots the coefficients on reallocation from a regression of the change in employment at different horizons on reallocation.

standard deviation during recession-recovery cycles of about 0.2. Thus, a one standard deviation increase in reallocation results in employment growth about 0.6 p.p. slower per year, culminating in an employment level about 2% lower at the end of the national cycle.

Figure 5 further explores the trajectory of the employment response to reallocation. The solid line with X hashmarks plots the coefficients $\beta_{1,j}$ from a local projection of employment growth on reallocation:

$$\Delta \ln e_{a,t,t+j} = \alpha_{t,j} + \beta_{1,j} R^b_{a,t,t+T} + \beta_{2,j} g^b_{a,t,t+T} + \gamma'_j X_{a,t} + \sum_{h \in \{j\}} \delta_{j,h} g^b_{a,t,t+h}. \tag{15}$$

These coefficients trace out an impulse response function. Areas undergoing reallocation during a national recession-recovery cycle experience a sharp relative fall in employment immediately following the national peak. The decline appears persistent, with the trough after five years.The coefficients for one and two years before the national peak indicate little evidence of areas with large predicted reallocation during the recession-recovery experiencing differential employment trends immediately prior to the national peak. For comparison, we also plot in the solid line with O hashmarks the impulse response for a robustness sample excluding areas in the top quartile of manufacturing share and described further below. The impulse response function generated by the robustness sample lies within the 95% confidence interval bands of the baseline impulse response function but has a more U-shaped pattern and a trough after 24 months.

To what extent does the fall in employment coincide with out-migration of the working-age population? Table 6 repeats the regressions reported in table 5, but replacing the log of

Table 6 – Effects of Reallocation on Migration During Recession-Recovery Cycles

	Dependent variable:							
	$\frac{12}{T}\,\Delta \ln l_{t,t+T}$				$\Delta \ln l_{t,t+48}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Right hand side variables:								
Predicted reallocation	-1.02^{**}	-0.72^{+}	-1.12^{*}	-0.73^{+}	-4.18^{**}	-3.08^{*}	-4.83^{*}	-2.95^{+}
	(0.35)	(0.37)	(0.48)	(0.40)	(1.43)	(1.49)	(2.02)	(1.54)
National cycle FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other controls	No	Yes	No	No	No	Yes	No	No
Bartik growth quantiles	No	No	Yes	No	No	No	Yes	No
Geographic area FE	No	No	No	Yes	No	No	No	Yes
CSA-MSA clusters	219	219	219	219	219	219	219	219
Observations	748	748	748	748	748	748	748	748

Notes: The dependent variable is indicated in the table header. $\frac{12}{T}\,\Delta \ln l_{t,t+T}$ is the monthly working-age population growth rate during the national recession-recovery cycle, expressed at an annual rate. $\Delta \ln l_{t,t+48}$ is the log change in the working-age population over the four years following the national peak. The other time-varying controls include all of the covariates reported in columns (2) and (6) of table 5. Standard errors in parentheses and clustered by CSA-MSA.

employment with the log of the working-age population in the dependent variable. Depending on the specification, roughly one-third to one-half of the decline in employment constitutes workers leaving the geographic area. This magnitude lies below but reasonably close to the estimate in Blanchard and Katz (1992) for arbitrary labor demand shocks.[18] The persistence of population movements may partly explain the persistence of the impulse response function.

5.2. Baseline Expansion Results and Comparison

Table 7 reports a parallel set of results for reallocation occurring during national expansions. Reallocation has a quantitatively smaller effect on employment outcomes during an expansion than during a recession-recovery. In almost all cases, the data do not reject zero effect.

Table 7 also reports the p-values from equality tests of the effect of reallocation during a recession-recovery and an expansion. We construct the p-values by estimating pooled regressions and interacting each covariate and predicted reallocation with an indicator for recession-recovery or expansion cycle. As apparent from comparing the predicted reallocation coefficients in tables 5 and 7, the interacted coefficient Reallocation X Recession-recovery has a negative

[18]Two differences in sample may explain the differences in magnitudes. First, Blanchard and Katz (1992) use states as their unit of analysis, whereas we use combined and metropolitan statistical areas. While states are larger, from the perspective of migration they have arbitrary borders. Second, migration has trended down in the United States over the past 30 years (Molloy, Smith, and Wozniak, 2011; Kapan and Schulhofer-Wohl, Forthcoming; Ganong and Shoag, 2015), implying much less unconditional migration during our sample period than during the period analyzed in Blanchard and Katz.

Table 7 – Effects of Reallocation on Employment During Expansion Cycles

	Dependent variable:							
	$\frac{12}{T}\Delta \ln e_{t,t+T}$				$\Delta \ln e_{t,t+48}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Right hand side variables:								
Predicted reallocation	−0.28	−0.66	−1.86*	0.23	1.40	−1.54	−4.82	5.24
	(0.79)	(0.87)	(0.75)	(1.04)	(3.48)	(3.79)	(4.24)	(5.68)
Predicted growth	0.93**	0.91**	2.33**	1.66**	−3.44+	−3.75+	0.54	−6.38*
	(0.21)	(0.22)	(0.77)	(0.33)	(2.07)	(2.03)	(3.85)	(3.06)
$\Delta \ln l_{t-60,t-12}$	0.094**	0.098**	0.11**	−0.15**	0.21*	0.15	0.24*	−0.87**
	(0.017)	(0.021)	(0.017)	(0.049)	(0.095)	(0.12)	(0.098)	(0.31)
$g^b_{a,t,t+48}$					2.18**	2.12**	2.69**	3.85**
					(0.51)	(0.46)	(0.54)	(0.69)
National cycle FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other controls	No	Yes	No	No	No	Yes	No	No
Bartik growth quantiles	No	No	Yes	No	No	No	Yes	No
Geographic area FE	No	No	No	Yes	No	No	No	Yes
Pred. reallocation mean	0.76	0.76	0.76	0.76	0.76	0.76	0.76	0.76
Pred. reallocation s.d.	0.12	0.12	0.12	0.12	0.12	0.12	0.12	0.12
Dep. var. mean	2.37	2.37	2.37	2.37	6.87	6.87	6.87	6.87
Dep. var. s.d.	1.58	1.58	1.58	1.58	10.02	10.02	10.02	10.02
R^2	0.41	0.47	0.51	0.73	0.63	0.67	0.69	0.80
CSA-MSA clusters	218	218	218	218	218	218	218	218
Observations	557	557	557	557	557	557	557	557
P-value: $\beta^{\text{expansion}} = \beta^{\text{recession-recovery}}$	0.002	0.050	0.202	0.009	0.001	0.055	0.112	0.010

Notes: The dependent variable is indicated in the table header. $\frac{12}{T}\Delta \ln e_{t,t+T}$ is the monthly employment growth rate during the national expansion, expressed at an annual rate. $\Delta \ln e_{t,t+48}$ is the log change in employment over the four years following the national last-peak. The other time-varying controls include lagged house price growth; lagged employment growth; area size, measured by the log of sample mean employment; and the Herfindahl of industry concentration at the cycle start. Standard errors in parentheses and clustered by CSA-MSA. The line $\beta^{\text{expansion}} = \beta^{\text{recession-recovery}}$ reports the p-value from a t-test on predicted reallocation interacted with recession-recovery in a pooled regression including both recession-recovery and expansion episodes and interacting each covariate as well as predicted reallocation with an indicator for recession-recovery or expansion.

sign and large economic magnitude in all specifications. In the most and least parsimonious specifications, corresponding to columns (1), (4), (5), and (8), we reject equality of coefficients during recession-recovery and expansion cycles at the 1% level. We reject at the 5% level in column (2), and at the 10% level in column (6).

Table 8 – Robustness

	Recession-recovery sample				Expansion sample			
Specification	β	S.E.	CSAs	Obs.	β	S.E.	CSAs	Obs.
1. Baseline	-3.13^{**}	0.66	219	748	-0.66	0.87	218	557
2. Drop 1980-83	-1.98^{**}	0.71	217	610				
3. Drop 1990-93	-2.20^{*}	0.88	218	534				
4. Drop 2000-05	-4.20^{**}	0.78	219	556				
5. Drop 2008-14	-3.68^{**}	0.69	216	544				
6. Trim construction	-2.74^{**}	0.69	196	526	-0.05	1.01	198	416
7. Trim health care	-2.82^{**}	0.84	188	522	-0.42	0.89	186	415
8. Trim manufacturing	-3.50^{**}	0.99	185	520	-0.76	0.95	177	415
9. Drop if large industry	-2.97^{**}	0.80	202	640	0.01	0.92	198	467
10. Extend recession/recovery	-2.59^{**}	0.78	217	685	0.09	0.81	216	440
11. HP filter dating	-2.96^{**}	0.82	217	741	-0.32	0.72	219	550
12. NAIRU dating	-2.76^{**}	0.70	217	637	0.60	0.58	218	556
13. Peak-to-peak reallocation	-3.07^{**}	1.00	219	748	-0.23	0.84	218	557

Notes: Each panel of each row of the table reports the coefficient and standard error of predicted reallocation from a separate regression, with the dependent variable monthly employment growth during the cycle. Each regression also includes predicted growth, a lag of population growth, and cycle fixed effects. The first row, labeled "Baseline", reproduces column (1) of table 5 (left panel) and the corresponding result for expansions (right panel). Rows 6-8 exclude observations in the cycle's top quartile of employment share in the industry indicated. Row 9 excludes observations where the area contains at least one industry with employment of 5% or more of the national total in that industry at the cycle start. Row 10 extends the national recession-recovery cycle by 3 months on each side. In row 11, a recovery ends when the cyclical component of an HP filter of national employment (smoothing parameter 129,600) equals zero. In row 12, a recovery ends when the national unemployment rate first equals the CBO's estimate of the NAIRU. Row 13 constructs predicted reallocation on a peak-to-peak basis.

5.3. Robustness

The negative effect of labor reallocation on aggregate employment during recession-recovery cycles is a robust result. Table 8 groups together a number of sensitivity exercises. Each panel of each row of the table reports the coefficient and standard error from a separate regression, with the dependent variable monthly employment growth during the cycle. For brevity, we report results only for the parsimonious specification controlling for predicted growth, a lag of population growth, and cycle fixed effects. Thus, the first row, labeled "Baseline", reproduces column (1) of table 5 (left panel) and the corresponding result for expansions (right panel).

Rows 2-5 provide subsample analysis by dropping one recession-recovery at a time. No single episode dominates the results.[19]

[19]The small increase in magnitude moving from row 2 to row 5 suggests slightly larger effects of reallocation earlier in the sample. The decline in the marginal effect of reallocation mirrors the decline in total reallocation reported in table 1. The model described in section 6 predicts this form of convexity; in the model economy, a given difference in reallocation between two areas results in a larger difference in employment outcomes if the average reallocation rate is higher.

Rows 6-8 explore the importance of particular industries. In section 4, we identified two sectors, specialty trade contractors and health care, which in sample have systematically different share changes during recession-recovery cycles and expansions. Already, the columns of table 5 with area fixed effects suggest that permanent exposure to these industries cannot fully explain the negative relationship during recession-recovery cycles. Row 6 provides an alternative sensitivity test by removing from the sample observations above the 75th percentile in beginning-of-cycle employment share in specialty trade contractors. Row 7 does the same, but for the share of employment in health care. Neither sample restriction generates a meaningful difference in the estimated effect of reallocation. Row 8 assesses sensitivity to removing observations with the highest exposure to the manufacturing sector, a sector which has contracted in employment share almost continuously through our sample and as a result has received substantial attention elsewhere (Autor et al., 2013; Charles et al., 2014). We cannot reject equality of the recession-recovery coefficients in rows 6-8 with the baseline, and all remain statistically significant at the 1% level. Likewise, the expansion coefficients remain much smaller in absolute magnitude and statistically indistinguishable from zero.

Row 9 addresses one possible concern with the Bartik predicted reallocation and employment variables. If employment in an industry concentrates in a few areas (for example, auto manufacturing employment in Detroit and Birmingham), and if the firms in different areas engage in strategic interaction (for example, negative shocks to plants in Detroit induce expansion at plants in Birmingham), then local industry employment in an area may correlate with elsewhere employment in that industry because of a strategic response to shocks specific to another local area's industry, rather than because of the common response to a common set of shocks. The coefficient of roughly one in the regressions in table 4 of actual on predicted reallocation and the positive coefficient on predicted growth in table 5 suggest absence of such concerns in our sample. Nonetheless, a more direct robustness exercise is to remove from the sample areas which contain an industry with employment relative to national employment in that industry above some threshold. Row 9 shows the result for a threshold of 5%. This restriction removes 17 CSA/MSAs from the sample (primarily the largest ones), but has almost no effect on the estimated coefficients.

Rows 10 to 12 explore robustness to the precise timing definition. Row 10 expands the recession-recovery symmetrically by 3 months on either side. For example, row 10 treats the peak of the Great Recession as occurring in October 2007 instead of January 2008, and the last-peak as May 2014 instead of February 2014.[20] The point estimate falls but only slightly. In row 11, we define the peak as in the baseline sample, but define the last peak as the first

[20]Because the 3 month shift changes the peak of the 1990-93 cycle to December 1989, the sample in row 10 includes only SIC-based reallocation for the 1990-93 cycle. The change from NAICS-based reallocation for most observations in this period, for which we use the LDB data without disclosure limitations, explains the fall in the sample size of roughly 60 observations in the left panel.

month following a peak in which the cyclical component of an HP filter (smoothing parameter 129,600) of national private sector employment turns positive. This change has almost no effect on the results for either the recession-recovery or expansion results. In row 12, we define the last peak as the first month in which the national unemployment rate falls to or below the Congressional Budget Office estimate of the NAIRU.[21] Again, the results change little.

Finally, row 13 completely redefines the reallocation timing to measure reallocation between two national employment peaks.[22] On the one hand, this peak-to-peak measure obviates the problem of identifying the end of a recession-recovery cycle for purposes of the Abraham and Katz (1986) critique. It also addresses the possibility of some industries leading and others lagging the aggregate cycle, such that our timing would count as secular reallocation the cyclical changes in shares between the leaders and laggards.[23] Third, it provides another means for addressing the in-sample patterns of construction and health care. On the other hand, peak-to-peak reallocation may provide a noisy measure of the actual reallocation occurring during the recession-recovery and expansion parts of the cycle. We would not want to conclude that reallocation affects aggregate outcomes only during the recession-recovery part of the cycle because our measure of reallocation correlates well with actual reallocation during that part of the cycle and poorly during the expansion part. In the event, using peak-to-peak reallocation results in coefficients nearly identical to those in the baseline specification.

6. Model

We summarize the results of the previous section in four stylized facts: (1) dispersion in industry labor demand causes a decline in local aggregate employment during recession-recovery cycles; (2) the decline is persistent; (3) out-migration of the working age population accounts for part of the decline; and (4) dispersion in industry labor demand has little to no effect on local aggregate employment during expansions. We now build a model of the labor market which generates this pattern and use it to draw inferences for the national economy.

The model economy consists of multiple geographic areas and industries linked by a demand structure similar to that described in section 2.2.1. We add three essential features which give rise to a quantitatively meaningful role for reallocation. First, each industry in an area contains a frictional labor market with both employed workers and unemployed searching for work in that industry. Second, each period only some agents may change industry or geographic

[21] For the 2008-14 cycle, we define the last peak as December 2014, the last month in our sample.

[22] For the 2008-14 cycle, we measure reallocation between the peak in January 2008 and December 2014. The dependent variables in row 13 are the same as in the baseline, i.e. they measure the employment change between the peak and last-peak (recession-recovery) or last-peak and peak (expansion).

[23] The specifications with area fixed effects already suggested this sort of variation does not drive the results, at least insofar as some industries always lead and lag the cycle and some areas have permanently higher exposure to those industries.

area or both. The reallocation decision depends on the returns to searching in a particular area and industry as well as on personal idiosyncratic preferences. These two features alone produce the qualitative pattern of aggregate area employment falling in response to an industry dispersion shock and workers moving away from areas undergoing more reallocation. The third feature delivers the quantitative difference between the effects of reallocation occurring during a downturn and an expansion. We embed the industry structure in a general equilibrium framework with sticky prices and downward nominal wage rigidity. The wage rigidity allows the model to overcome the unemployment volatility puzzle highlighted by Shimer (2005). More subtlety, the downward wage rigidity serves to compress the distribution of industry wages during recessions. The compression of wages forces more of the adjustment into differences in labor market tightness across industries, resulting in a larger unemployment response in recessions than in expansions. We provide auxiliary evidence in support of this mechanism in the form of regressions of changes in wage premia on employment share changes.

6.1. Setup

We develop a model in discrete time. The economy consists of A islands, each of which has up to I industries. Our calibration will feature $A = 2$, with one small island representative of an individual CSA/MSA and one large island representative of the rest of the economy, but it proves simpler to describe the model's equations without specializing to this case.

6.1.1. Labor market The labor market in each area-industry operates according to search and matching principles. At the beginning of period t, industry i in area a contains $(1-\delta_{t-1})e_{a,i,t-1}$ workers employed in the previous period and still attached to their firm, $x_{a,i,t}$ workers searching for a job, and $v_{a,i,t}$ job vacancies. Hiring occurs at the beginning of the period, with $n_{a,i,t}$ new matches formed. The $e_{a,i,t} = (1-\delta_{t-1})e_{a,i,t-1} + n_{a,i,t}$ workers employed in t engage in production. At the end of the period, $\delta_t e_{a,i,t}$ of the employed workers exogenously separate from their employer. We let $u_{a,i,t} = x_{a,i,t} - n_{a,i,t}$ denote the number of unemployed workers in period t after the matching process has taken place. Following Christiano et al. (2015), this concept of unemployment allows for job-to-job transitions by workers who separate at the end of $t-1$ but get newly hired at the beginning of t. We let $l_{a,i,t} = e_{a,i,t} + u_{a,i,t} = (1-\delta_{t-1})e_{a,i,t-1} + x_{a,i,t}$ denote the total labor force in industry i in area a at time t. We fix the economy-wide labor force at $\sum_{a=1}^{A}\sum_{i=1}^{I} l_{a,i,t} = 1$.

The firm vacancy posting condition and matching process are standard. Firms post $v_{a,i,t}$ vacancies in industry i at cost κ per vacancy. A free entry condition drives the expected value of a vacancy to zero. The matching function takes the Cobb-Douglas form $n_{a,i,t} = M v_{a,i,t}^{1-\alpha} x_{a,i,t}^{\alpha}$. Letting $\theta_{a,i,t} = \frac{v_{a,i,t}}{x_{a,i,t}}$ denote the vacancy-searcher ratio, or industry labor market tightness, searching workers find jobs at rate $f_{a,i,t} = M\theta_{a,i,t}^{1-\alpha}$, and firms fill vacancies at rate $q_{a,i,t} = M\theta_{a,i,t}^{-\alpha}$.

Unemployed workers search in one industry and one area at a time. Their choice of where to

search plays an important role. In line with recent literature, we assume semi-directed search (Kline, 2008; Artuç, Chaudhuri, and McLaren, 2010; Kennan and Walker, 2011; Pilossoph, 2014; Dvorkin, 2014). Specifically, at the end of period t, employed workers transition into unemployment in their same industry at rate $\delta_t - \lambda_t$. Both unemployed and employed workers receive an industry reallocation shock at exogenous rate $\lambda^I_{a,t}$ and an area reallocation shock at rate $\lambda^A_{a,t}$, with $\lambda_t = \lambda^I_{a,t} + \lambda^A_{a,t}$. An industry reallocation shock consists of an immediate job separation if previously employed, and a draw of I idiosyncratic taste shocks $\{\varepsilon_j\}^I_{j=1}$ from a distribution $F^I(\varepsilon)$. These taste shocks enter additively into the worker's value function for searching in each sector $j = 1, \ldots, I$ in the worker's initial area a. An area reallocation shock has two parts. First, the worker draws A idiosyncratic shocks $\{\varepsilon_b\}^A_{b=1}$ from a distribution $F^A(\varepsilon)$, which enter additively into the worker's value function for searching in area $b = 1, ..., A$. After choosing a location, she then draws idiosyncratic industry taste shocks $\{\varepsilon_j\}^I_{j=1}$ to determine her new industry. We parameterize $F^I(\varepsilon)$ and $F^A(\varepsilon)$ as Type I Extreme Value.

Reallocation shock frequencies $\lambda^A_{a,t}$ may be area-specific. We let $\bar{\lambda}^A_t$ denote the average area reallocation shock across islands. In our calibration, $\lambda^A_{a,t}$ will vary inversely with initial area size to ensure balanced migration flows in steady state. As a corollary, workers in small areas disproportionately receive taste shocks ε_b, which raises their utility. Offsetting this, residents of area a enjoy an amenity $(\bar{\lambda}^A_t - \lambda^A_{a,t})\,(\mathbf{E} \max_b \varepsilon_b)$, where $\mathbf{E}$ denotes the expectation operator and the absence of a time subscript denotes the steady state value, so that the option value of moving does not vary with island size. The assumption that workers receive amenities from living in areas with greater population density has some empirical support (Diamond, 2015).

We comment briefly on the industry and geographic mobility assumptions. The parameter λ^I determines the share of unemployed re-optimizing their industry search market. Holding this share below unity provides one important friction allowing reallocation shocks to affect employment. Having both employed and unemployed workers receive reallocation shocks at the same rate λ^I removes any option value of remaining unemployed in order to potentially switch sectors. One may also interpret λ^I as a stochastic death or retirement shock, with a new generation of workers of mass λ^I born each period and choosing afresh their industry of search. The ε_j shocks have the interpretation of taste shocks which make some individuals prefer to work in certain sectors, or of noise shocks which give individuals private (mis)information about the returns to searching in each sector. Inclusion of these shocks generates two-way gross labor flows across industries. Existence of gross flows in excess of the net reallocation flows induced by non steady-state dynamics captures an important feature of reality and facilitates computation by ensuring an interior solution in cross-industry flows at each moment in time. The level of λ^I and the volatility of the process generating ε_j together govern the directness of search across industries. Analogously, the magnitude of area reallocation shocks λ^A_a and the volatility of area taste shocks determine the directness of search across locations.

We denote the transition probability from industry i to industry j conditional on an industry reallocation shock by $\pi^I_{a,ij,t}$. This probability does not depend on the worker's previous employment status or industry, $\pi^I_{a,ij,t} = \pi^I_{a,kj,t} = \pi^I_{a,j,t}$. We denote the transition probability from area a to area b conditional on an area reallocation shock by $\pi^A_{ab,i,t}$ for a worker starting in industry i. Upon entering a new area b, the worker chooses industry j with probability $\pi^I_{b,j,t}$. Area reallocation shocks are then also independent of the worker's employment status, initial area and initial industry, $\pi^A_{ab,i,t} = \pi^A_{cb,j,t} = \pi^A_{b,t}$. We have three laws of motion for the evolution of job seekers, employment, and unemployment:

$$
\begin{aligned}
x_{a,i,t} &= \delta_{t-1} e_{a,i,t-1} + u_{a,i,t-1} - \lambda_{t-1} l_{a,i,t-1} + \pi^I_{a,i,t-1} \left[\lambda^I_{a,t-1} l_{a,t-1} + \pi^A_{a,t-1} \sum_{b=1}^{A} \lambda^A_{b,t-1} l_{b,t-1} \right], \\
e_{a,i,t} &= (1 - \delta_{t-1}) e_{a,i,t-1} + f_{a,i,t} x_{a,i,t}, \\
u_{a,i,t} &= (1 - f_{a,i,t}) x_{a,i,t}.
\end{aligned}
$$

We assume no aggregate uncertainty and perfect consumption insurance within (but not across) islands. Thus, workers and firms in an area both evaluate the future with the discount factor $m_{a,t,t+1}$. Let $p_{a,i,t}$ denote the real marginal product of a match, $w_{a,i,t}$ the real wage payment to the worker, z the worker's flow opportunity cost of employment, $J_{a,i,t}$ the value of a filled job to a firm, $W_{a,i,t}$ the value of a filled job to a worker, and $U_{a,i,t}$ the value of unemployment in industry i of area a to a worker. The following three Bellman equations and free entry condition summarize the labor market block of the model:

$$
J_{a,i,t} = (p_{a,i,t} - w_{a,i,t}) + (1 - \delta_t) m_{a,t,t+1} J_{a,i,t+1}, \tag{16}
$$

$$
\begin{aligned}
W_{a,i,t} = {} & w_{a,i,t} + m_{a,t,t+1} \left(\bar{\lambda}^A_t - \lambda^A_{a,t} \right) \left(\mathbf{E} \max_b \varepsilon_b \right) \\
& + m_{a,t,t+1} \Bigg\{ \left[(1 - \delta_t) + (\delta_t - \lambda_t) f_{a,i,t+1} \right] W_{a,i,t+1} + (\delta_t - \lambda_t) (1 - f_{a,i,t+1}) U_{a,i,t+1} \\
& \quad + \lambda^I_{a,t} \left(\mathbf{E} \max_j \left\{ (1 - f_{a,j,t+1}) U_{a,j,t+1} + f_{a,j,t+1} W_{a,j,t+1} + \varepsilon_j \right\} \right) \\
& \quad + \lambda^A_{a,t} \left(\mathbf{E} \max_b \left\{ \max_j \left[(1 - f_{b,j,t+1}) U_{b,j,t+1} + f_{b,j,t+1} W_{b,j,t+1} + \varepsilon_j \right] + \varepsilon_b \right\} \right) \Bigg\},
\end{aligned} \tag{17}
$$

$$
\begin{aligned}
U_{a,i,t} = {} & z + m_{a,t,t+1} \left(\bar{\lambda}^A_t - \lambda^A_{a,t} \right) \left(\mathbf{E} \max_b \varepsilon_b \right) \\
& + m_{a,t,t+1} \Bigg\{ (1 - \lambda_t) \left[f_{a,i,t+1} W_{a,i,t+1} + (1 - f_{a,i,t+1}) U_{a,i,t+1} \right] \\
& \quad + \lambda^I_{a,t} \left(\mathbf{E} \max_j \left\{ (1 - f_{a,j,t+1}) U_{a,j,t+1} + f_{a,j,t+1} W_{a,j,t+1} + \varepsilon_j \right\} \right) \\
& \quad + \lambda^A_{a,t} \left(\mathbf{E} \max_b \left\{ \max_j \left[(1 - f_{b,j,t+1}) U_{b,j,t+1} + f_{b,j,t+1} W_{b,j,t+1} + \varepsilon_j \right] + \varepsilon_b \right\} \right) \Bigg\},
\end{aligned} \tag{18}
$$

$$
\kappa = q_{a,i,t} J_{a,i,t}. \tag{19}
$$

Wages follow a Nash bargain between the firm and worker, subject to exogenously imposed downward nominal wage rigidity. This rigidity takes the form

$$w_{a,i,t} = \max\{w^*_{a,i,t}, (1-\chi^w)\, w_{a,i,t-1}/\Pi_{a,t}\}, \tag{20}$$

where $w^*_{a,i,t}$ is the Nash bargained real wage, $\Pi_{a,t}$ is gross producer price inflation, and χ^w is a parameter specifying the maximum permitted decline in the nominal wage. The assumption of downward nominal rigidity has some support in the empirical literature (Kahn, 1997), and following Hall (2005) and Chodorow-Reich and Karabarbounis (2015) allows the model to generate realistic unemployment fluctuations without violating bilateral efficiency conditions or requiring counterfactual assumptions on the sources of wage rigidity.

6.1.2. General equilibrium We embed the industry structure in a demand framework similar to that in section 2.2.1. Output of industry i in area a is

$$Q_{a,i,t} = \eta_{i,t} e_{a,i,t}, \tag{21}$$

where $\eta_{i,t}$ is (strictly exogenous) labor productivity in industry i which by free flow of information does not vary across islands. Industry output is sold under perfect competition at real price $P^Q_{a,i,t}$ to a wholesaler. The wholesaler combines local industry output into an area-specific good $Q_{a,t}$ using the technology

$$Q_{a,t} = \left[\sum_i \tau_{a,i,t}^{\frac{1}{\zeta}} Q_{a,i,t}^{\frac{\zeta-1}{\zeta}}\right]^{\frac{\zeta}{\zeta-1}}, \tag{22}$$

giving rise to a downward sloping industry-level demand curve $Q_{a,i,t} = \tau_{a,i,t}\left(\frac{P^Q_{a,i,t}}{P^Q_{a,t}}\right)^{-\zeta} Q_{a,t}$, and where $\zeta \geq 1$ and $P^Q_{a,t} = \left[\sum_i \tau_{a,i,t}(P^Q_{a,i,t})^{1-\zeta}\right]^{\frac{1}{1-\zeta}}$. In our calibration, we vary the parameters $\{\tau_{a,i,t}\}$ across islands to generate variation in steady state employment shares.

The real marginal revenue product $p_{a,i,t}$ arising in equation (16) is the product of industry productivity and the real price of industry i's good:

$$p_{a,i,t} = \eta_{i,t} P^Q_{a,i,t}. \tag{23}$$

With downward sloping demand, the decline in output engendered by a decline in $\eta_{i,t}$ induces a rise in the real price $P^Q_{a,i,t}$, such that following a negative productivity shock the marginal revenue product $p_{a,i,t}$ changes little but output and employment in sector i fall.

Closing the model requires specifying the determination of the set of real industry prices $P^Q_{a,i,t}$, overall inflation, and the discount factor $m_{a,t,t+1}$. We do so by incorporating standard elements from a New Keynesian currency union model. In short, a continuum of retailers combine the Q good with capital to produce differentiated final goods which they sell to households

across all islands, subject to Rotemberg pricing adjustment costs and home bias in the goods market. An Euler equation for each household determines how much of the final goods to consume and how much to invest. While agents enjoy perfect consumption insurance within an area, asset markets across areas allow only for trade of a nominal bond. A central bank follows a standard interest rate rule that satisfies the Taylor principle. Finally, we allow for a wedge μ_t between the policy interest rate and the interest rate faced by households, and use an increase in the wedge to generate a demand-induced recession. We provide a detailed discussion and formal statement of the equations of the remainder of the model in appendix F.

6.2. Model Intuition

To understand why this model generates additional unemployment during a recession in areas undergoing secular reallocation, it helps to consider an economy with a single area and with two industries indexed as 1 and 2. Suppose the industries begin in an initial, symmetric, steady state, with $\eta_{1,t-1} = \eta_{2,t-1} = 1$, $\tau_{1,t-1} = \tau_{2,t-1} = s_{1,t-1} = s_{2,t-1}$, and where as above $s_{i,t}$ denotes the employment share. Two shocks occur simultaneously at time t. First, the steady state productivities undergo a permanent, mean-preserving spread, with $\eta_{1,t+\infty} > 1 > \eta_{2,t+\infty}$. Second, a temporary increase in the interest rate spread μ_t generates a recession.

Due to reallocation frictions, the marginal revenue products $p_{1,t+\infty}$ and $p_{2,t+\infty}$ initially diverge and then compress again as employment moves out of sector 2 and into sector 1. Along the transition, labor market tightness diverges, as firms concentrate their vacancy postings in the growing industry, while the search and reallocation frictions keep a mass of workers searching in the declining industry. The recession causes the downward wage rigidity to bind especially tightly on the declining industry. The downward rigidity reduces the firm's share of match surplus in the declining industry, further disincentivizing vacancy creation, while the compression of wage differentials keeps additional workers searching in the declining industry. This further divergence in vacancy posting and unemployment implies a much larger difference in labor market tightness between the two industries than would occur without the recession.

With a constant returns to scale matching function, differential labor market tightness generates unemployment.[24] Fixing $\delta_t = \delta$, fluctuations in unemployment depend on fluctuations in job finding rates. In the two sector version, the economy-wide average job finding rate f_t equals $M\left(\frac{x_{1,t}}{x_t}\theta_{1,t}^{1-\alpha} + \frac{x_{2,t}}{x_t}\theta_{2,t}^{1-\alpha}\right)$. Holding the distribution of vacancies fixed and maximizing f_t over the allocation of job seekers requires setting $\theta_{1,t} = \theta_{2,t}$. The sharp divergence in θ_1 and θ_2 caused by reallocation thus generates unemployment. This reallocation unemployment closely resembles the mismatch unemployment in Sahin et al. (2014) and Barnichon and Figura (Forthcoming).[25] The addition of nominal rigidities, capital, and labor mobility across areas

[24] Pissarides (2000) surveys the properties of matching functions and finds broad empirical support for and almost universal adoption of the assumption of constant returns to scale.

[25] Our focus on wage rigidity does not mean that we reject other possible mechanisms. Recent work has

amplify this mechanism and as we show next allow the model to quantitatively match our empirical estimates.

6.3. Quantitative Results

We calibrate a version of the model with two areas and twenty-three industries, $A = 2$ and $I = 23$, at monthly frequency. The two areas allow for one small area which we treat as representative of a single local CSA/MSA, and one large area representative of the rest of the economy. Three industries is the minimum number necessary to have different amounts of reallocation result from the same mean-preserving spread in industry productivities, but a larger number of industries more closely approximates the empirical specification.[26] Appendix F provides details of our calibration and procedure for finding the model steady state. Appendix table F.1 provides a summary of the calibrated parameters and moments matched. We solve the perfect foresight transition paths by reversing the shocks after 700 periods and working backward from the initial steady state equilibrium.[27]

Reallocation is induced by an unexpected, gradual mean-preserving spread in industry productivities. These are constructed as follows: We take average observed sector share changes in the NAICS-3 data, sort them into 23 quantiles, and take the mean in each quantile. The productivity changes are fixed multiples of the sector share changes chosen to match the standard deviation of industry share changes in the data. Thus, the log productivity changes have mean zero. The productivity changes occur gradually, beginning in period 0 and reaching full spread after 48 months. We order the industries in decreasing order of productivity change, such that industry 1 experiences the largest productivity increase, industry 2 the second largest increase, etc. down to industry 23 which experiences the largest productivity decline. The left panel of figure 6 plots these perfect-foresight paths.

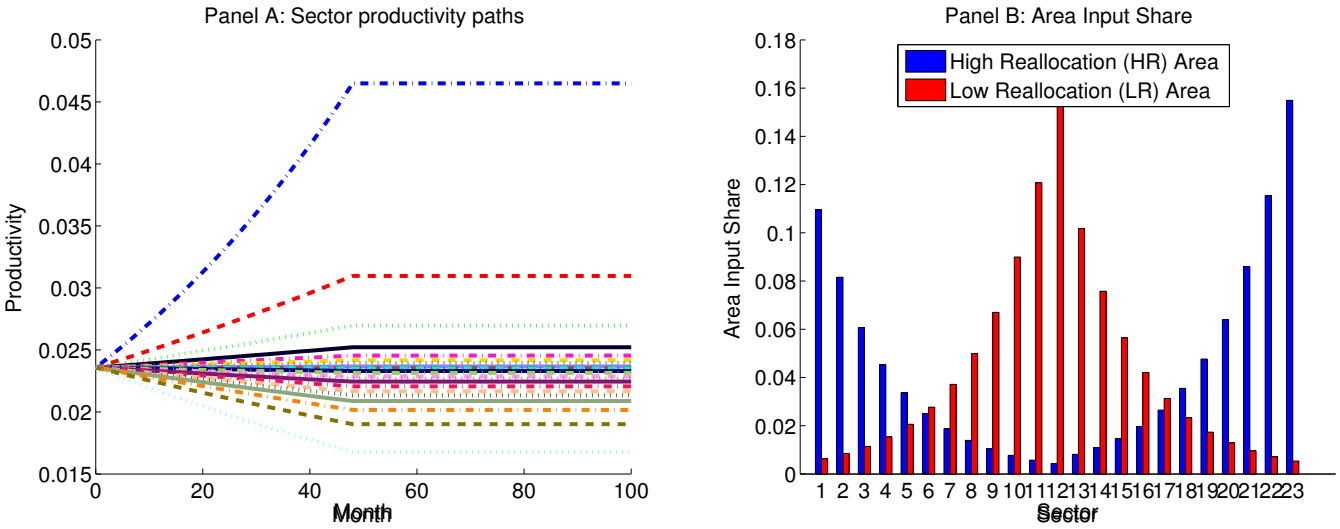
We solve the model separately for a small area with low reallocation (LR) and with high reallocation (HR). In the LR area, we start from input weights $\tau_{a,j} = 2^{-0.425|j-12|} / \sum_{j=1}^{23} 2^{-0.425|j-12|}$, while in the HR area we start from input weights $\tau_{a,j} = 2^{0.425|j-12|} / \sum_{j=1}^{23} 2^{0.425|j-12|}$. We then scale the weights on the declining sectors such that the input-weighted log productivity change

highlighted two: the possibility of search inefficiency or job retraining associated with industry switchers (e.g. Jaimovich and Siu, 2014), and changes in the pool of job seekers during recessions (e.g. Hall and Schulhofer-Wohl, 2015; Ahn and Hamilton, 2014). Pilossoph (2014) and Dvorkin (2014) criticize the inefficient search explanation on the grounds that gross flows of workers across industries vastly exceed net flows and appear mostly unresponsive to changes in net flows, a fact consistent with our model. We also note that on at least some observable dimensions such as educational attainment, the composition of workers switching industries appears roughly acyclical. We provide direct evidence of the wage compression channel in table 9.

[26] Twenty-three industries corresponds to the number at the NAICS2 level. While the model allows for an arbitrary, finite number of sectors, we found it computationally infeasible to scale our model to the $I \approx 80$ industries at the SIC2/NAICS3 level. Increasing the number of industries from three to twenty-three slightly amplifies the marginal effect of reallocation on employment, suggesting that we do not bias toward larger responses by restricting to small I.

[27] Allowing for a longer time before reversal has little effect on our results.

Figure 6 – Productivity Paths and Local Input Shares



Notes: Panel A displays the perfect-foresight productivity paths η_i of each of the $i = 1, ..., 23$ sectors. The productivity paths are mean-preserving in logs. The right panel displays the input shares τ_{ai} for the high-reallocation area (blue) and the low-reallocation area (red).

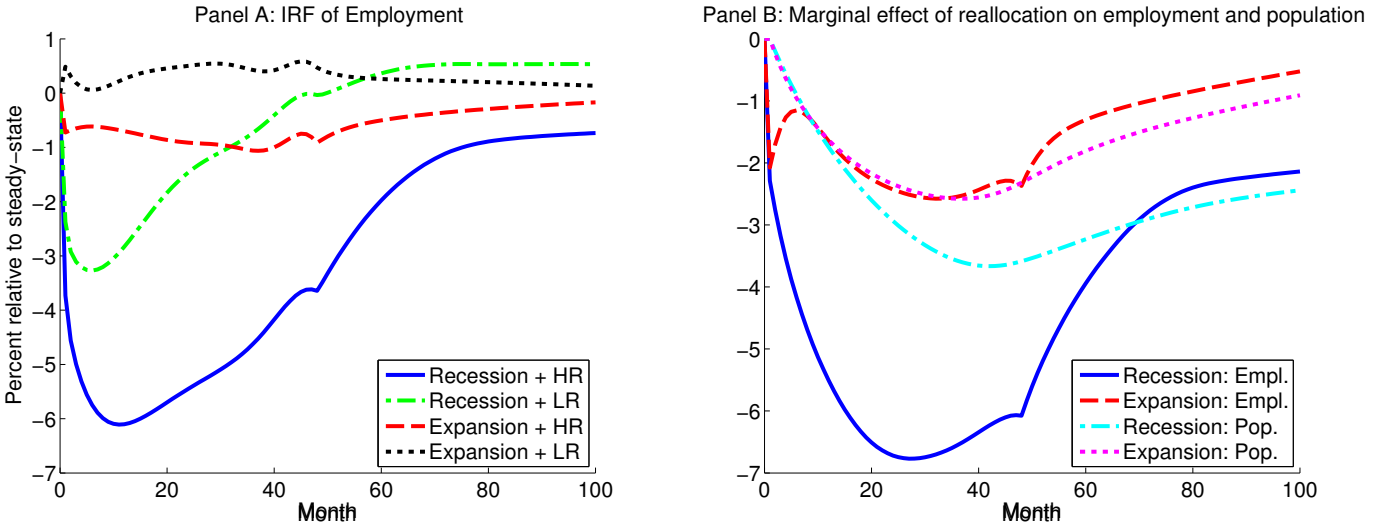
in each area is mean-zero. The right panel of figure 6 plots these input shares. The low reallocation area has predicted reallocation of 0.7%, approximately a one-and-a-half standard deviation below median reallocation in our sample. The input shares of the high reallocation area are more concentrated among the expanding and contracting sectors and give predicted reallocation of 1.3%, corresponding to a one-and-a-half standard deviation increase in reallocation relative to the median area in the data. In both versions, the large area has uniform input shares $\tau_{b,j} = 1/23$. Since this area is large, it will behave identically regardless of the reallocation in the small area. Thus, this set-up allows us to compare outcomes under high and low reallocation for members of the same currency union.

We consider two scenarios. A pure reallocation shock consists of only the mean-preserving spreads in productivities and corresponds to the expansion results in section 5. The second experiment combines the reallocation shock with a temporary increase in the wedge between the central bank's policy interest rate and the interest rate available to households. This wedge increase simulates a demand-driven recession. We set the increase in the wedge to 0.8 p.p. with subsequent decay of 4.5% per month to generate an employment contraction of approximately 4% in the large area. The recession *cum* reallocation shock corresponds to the recession-recovery results in section 5.[28]

The left panel of figure 6 plots the employment rates in the small area in the high and low reallocation version and for each scenario. In the expansion scenario, the low reallocation area

[28] A pure recession shock has nearly identical effects in both areas. Thus, any differential exposure to a recession comes from the *combination* of the recession shock with the reallocation shock.

Figure 7 – Model Impulse Response Function and Marginal Effect



Notes: Panel A the impulse response functions of employment relative to steady-state in the small member of the currency union. HR correspond to initial area shares given by the blue bars in panel B of figure 6 (high reallocation). LR corresponds to the red bars (low reallocation). A recession combines the mean-preserving productivity paths with an AR(1) interest rate wedge of 0.8p.p. with persistence 0.955. Panel B displays the marginal effect of reallocation in recessions and expansions based on equation (24). The marginal effect is the difference in employment/population between the high-reallocation and low-reallocation area divided by the difference in predicted reallocation.

has essentially no employment change, while employment in the high reallocation area falls by approximately 1% after 40 months. A recession induces a much larger decline in employment. In the low reallocation area the decline reaches around 3% after 6 months, while in the high reallocation area the trough is 6% after 10 months.

The right panel of figure 7 plots the marginal effect of reallocation on unemployment for the expansion and recession-recovery scenarios. The marginal effect is the ratio of the difference in employment change to the difference in predicted reallocation,

$$D_{R,t} = \frac{e_{HR,t} - e_{LR,t}}{R^b_{HR,0,T} - R^b_{LR,0,T}}, \tag{24}$$

and corresponds to the marginal effects of reallocation on employment that we estimate in columns (5) through (8) in tables 5 and 7 and plot in figure 5.[29] The timing of the Bartik

[29]We have also estimated regressions of local area employment growth on Bartik growth and Bartik reallocation in simulated data. Specifically, we ran 100 simulations of a 5-sector partial equilibrium version of the model corresponding to the set-up in section 6.1.1 but with exogenous marginal revenue product paths, a constant discount factor, and a constant zero inflation rate. Relative to the full quantitative model, which takes approximately five hours to solve, this model is solved significantly faster. For each simulation, we drew a different initial local productivity distribution over the five sectors, giving rise to a different initial industry employment distribution. We then regressed area employment growth on Bartik reallocation and Bartik growth. The coefficient on reallocation aligns very closely with the result from simulating two local areas that have the same predicted Bartik growth but different levels of reallocation and calculating the marginal effect as in equation (24).

reallocation variable depends on the large area's recession-recovery cycle, which is just over 5 years long, $T = 69$. We use the same duration for the expansion cycle. In the model, the marginal effect of reallocation on unemployment is significantly larger and more persistent in recession-recoveries than in expansions. The maximum marginal effect during the recession-recovery is about -7, only slightly below the range of the estimates reported in table 5 from the data. The effect is persistent and traces a path similar to the manufacturing-trimmed impulse response function in figure 5. In contrast, the maximum marginal effect during the expansion is about -2.5, which is within the range of point estimates reported in table 7 and substantially below the effect during recession-recoveries. The share of the employment response accounted for by population migration also appears roughly consistent with the data.

6.4. Discussion

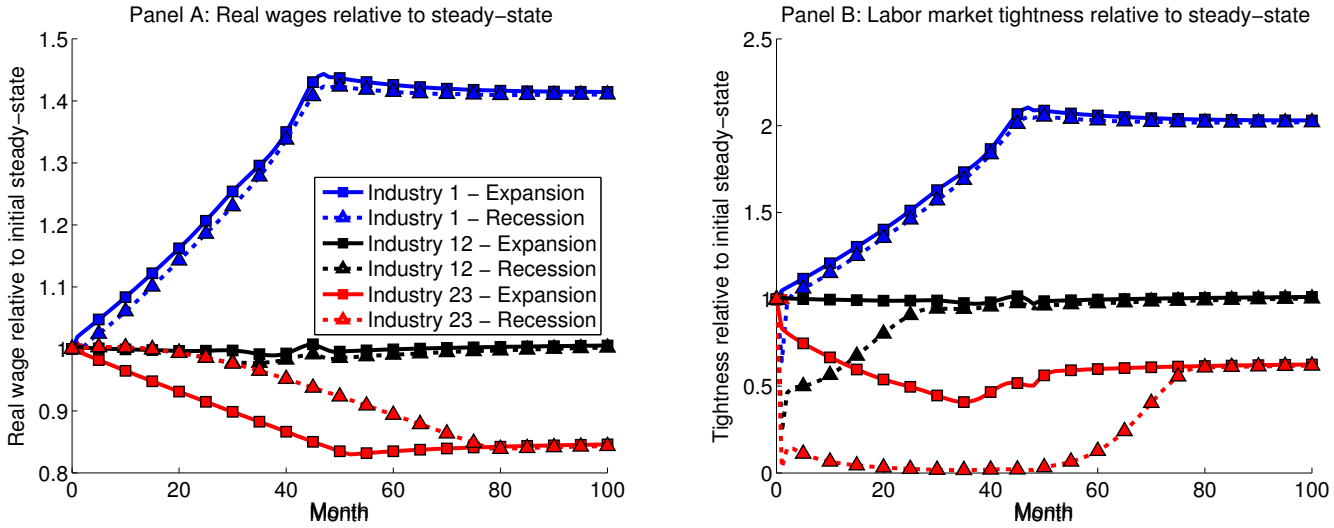
The model replicates the finding that reallocation generates a quantitatively large decline in aggregate employment if it occurs during a recession but not if it occurs during an expansion. Compression of the wage distribution during recessions forms a key mechanism underlying this difference. In expansions the wage constraint does not bind, so real wages rise in the expanding sector and fall in the declining sector commensurate with the rate of the productivity change. This is shown by the squared-hash lines in the left panel of figure 8, where, for ease of exposition, we only show the most productive, least productive and median sector. Due to imperfect labor mobility, tightness diverges as shown in the right panel of figure 8, causing a decrease in aggregate matching efficiency and reducing employment. However, the divergence in tightness is small as higher wages draw workers into the expanding sectors, increasing the number of job seekers in the expanding sectors and reducing the number in the contracting sectors.

In a recession the downward nominal wage constraint becomes binding. However, the extent to which it binds differs across sectors, with the sharpest effect on sectors contracting due to both a temporary recession and permanent productivity decline, as illustrated by the triangle-hash lines in the left panel of figure 8. The increase in the workers' share of match surplus in the contracting sectors leads to greater divergence in allocation of both vacancies and job seekers, magnifying the divergence in labor market tightness.

To illustrate quantitatively the model's key mechanisms, figure 9 plots the baseline marginal coefficient impulse response function (the solid blue line) along with counterfactual marginal coefficients obtained by selectively changing features of the model. The dash-dot light blue line shows the marginal effects in recession-recoveries with wages determined by Nash bargaining in each period. Without downward wage rigidity, the effects are much smaller and of similar magnitude independent of the state of the business cycle.[30]

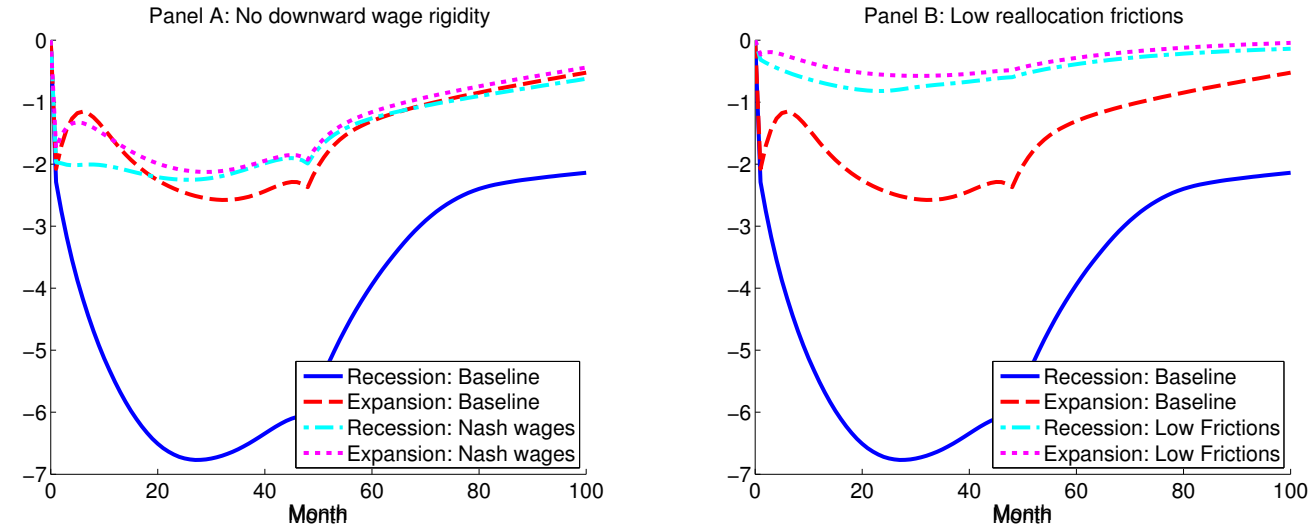
[30] A potential concern with the downward nominal wage rigidity mechanism is whether it applies in the 1980s when trend inflation was higher. However, this period also featured higher average reallocation and higher inter-regional mobility. When we re-calibrate our model to match these three features of the data

Figure 8 – Model Impulse Response Functions of Real Wages and Labor Market Tightness



Notes: Panel A plots the impulse response functions of real wages relative to the initial steady-state in the large area of the currency union. For ease of exposition, only the highest productivity (sector 1), median productivity (sector 12) and lowest productivity sector (sector 23) are plotted. Panel B displays the corresponding impulse response functions for labor market tightness θ_{bi}.

Figure 9 – Alternative Model Specifications



Notes: Panel A displays the marginal effect of reallocation on employment in the baseline model and and when period-by-period Nash bargaining replaces the downward nominal wage rigidity assumption. Panel B contrasts the baseline model with a parameterization that sets the area and industry reallocation noise parameters to 20 percent of their baseline values, which weakens reallocation frictions and makes labor search more directed. In this case the productivity paths are re-parameterized to yield the same amount of average reallocation as in the baseline calibration.

The asymmetry of wage compression during recessions and expansions has a counterpart

($\chi = 0.0052$, $\lambda_a^A = 0.005$, $R_{b,t,t+T} = 1.3\%$), the peak marginal effect of reallocation on local employment is around 6, close to our baseline effect. Moreover, in unreported results we find industry wage-compression in the 1980s recession of similar magnitude to that in the full sample reported in table 9.

Table 9 – Recession Wage Compression in the Data

	Dep. var.: change in industry wage premium	
	NAICS 2	SIC 2/NAICS 3
	(1)	(2)
Right hand side variables:		
Share change growth rate $\left(\frac{12}{j}\frac{2\Delta s_{i,t}}{s_{i,t}+s_{i,t-T}}\right)$	0.39^{+}	0.35^{*}
	(0.19)	(0.15)
Recession X $\frac{12}{j}\frac{2\Delta s_{i,t}}{s_{i,t}+s_{i,t-T}}$	-0.43^{*}	-0.35^{*}
	(0.20)	(0.17)
Employment share weighted	No	Yes
Industry clusters	17	143
Observations	102	492

Notes: The dependent variable is the change in the industry wage premium over the recession or expansion episode. The wage premium is a centered twelve month moving average of the industry fixed effect in a regression in the CPS ORG data of the log hourly wage on categorical variables for industry, race, 5 year age bin, gender, educational attainment, state, rural, and occupation. The variable $\frac{12}{j}\frac{2\Delta s_{i,t}}{s_{i,t}+s_{i,t-T}}$ is the annualized symmetric growth rate of the industry employment share during the expansion or the recession-recovery containing the recession in the QCEW data. Standard errors in parentheses and clustered by industry.

moment in the data. Table 9 tests for this asymmetry using national hourly wages by industry from the CPS and QCEW employment share changes during recession-recoveries and expansions. The table reports regressions where each observation is a national NAICS 2, SIC 2, or NAICS 3 digit industry during a national recession or expansion episode. The dependent variable is the change in the industry wage premium during a recession or expansion.[31] The regressors include the growth rate of the employment share in the industry during the expansion or recession-recovery containing the recession, and the growth rate interacted with the state of the business cycle.

In the data, industries with rising employment shares have rising wage differentials during expansions. In contrast, there is no economically or statistically significant relationship between the change in the wage premium during a recession and industry share growth. The data reject equality of coefficients during expansions and recessions at the 5% level. Because realized reallocation and wage differentials may be jointly determined, we do not read causality into these results. Nonetheless, they provide evidence of wage compression between expanding and contracting industries during recessions but not during expansions, consistent with the

[31]See appendix G for details of the construction of the dependent variable. By construction, the dependent variable has essentially zero mean across industries in a given time period, and the changes in industry shares also have essentially zero mean within a time period. We therefore omit time fixed effects from the regressions for parsimony. We weight the SIC 2/NAICS 3 digit regressions by employment share because smaller industries have greater measurement error in the industry wage premia. The CPS provides too small a sample to repeat the exercise at the CSA/MSA level.

mechanism in the model.[32]

Panel B of figure 9 illustrates the importance of reallocation frictions. The teal and pink dashed lines report the impulse response functions of the marginal coefficients from a calibration with the variance of the reallocation taste shocks $\varepsilon_j, \varepsilon_b$ reduced to 20% of their baseline values. This change makes the industry and location choice much more directed, and so in these experiments we feed in much less dispersion in productivity in order to generate the same amount of labor reallocation as in the baseline calibration. While the directedness of search per se has only a minor effect on the employment response to a given set of productivity processes, the lower productivity dispersion required to match the actual amount of reallocation in the data results in marginal coefficients of nearly zero in both expansions and recessions.

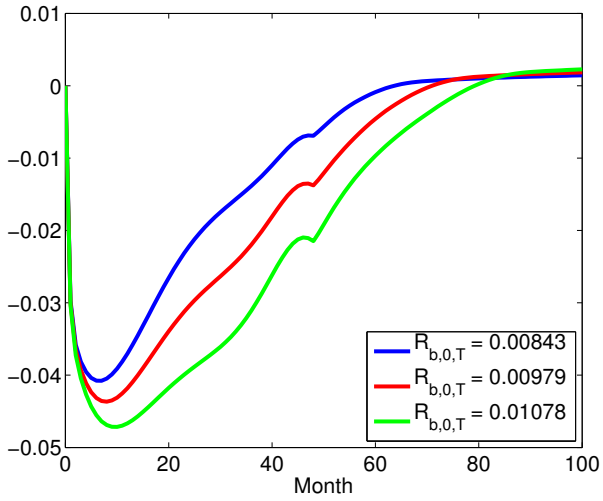
6.5. Aggregate Importance

Armed with our calibrated benchmark model, we can ask what it implies for the importance of reallocation in propagating and amplifying business cycles at the national level. To do so, we report employment responses in the large area to different amounts of reallocation generated by scaling up and down the mean-preserving productivity shocks shown in the left panel of figure 6. This experiment differs from the cross-sectional empirical and model-based results, which are based on *differential* exposure in small areas to the *same* mean-preserving productivity shocks, but perhaps better captures the source of time series variation in national reallocation. Likewise, the large area employment response excludes channels such as population mobility which affect local areas but not the national economy. Figure 10 shows how reallocation amplifies and propagates national recessions in the model. The effects can be substantial.

The model's implications for the relationship of secular industry labor reallocation to the behavior of vacancies and worker flows also merit comment. In their article, Abraham and Katz (1986) suggest looking to vacancies to ascertain the importance of reallocation shocks in driving the business cycle. Their intuition rests on a downward sloping Beveridge curve absent sectoral shifts, that is, a negative relationship between unemployment and vacancies, while reallocation shocks shift the Beveridge curve out and therefore induce positive comovement of vacancies and unemployment. Since vacancies fall in recessions, they argue that reallocation is not a major driver of business cycles. However, in our model reallocation reduces vacancies; comparing across the episodes shown in figure 10, a larger reallocation shock generates a larger decline in vacancies. The key is again the downward-rigid wage constraint. Vacancy creation in the declining sector plummets because of rigid wages, while the rising wages in the expanding sector limit vacancy creation there. Thus, our finding that reallocation can be important for

[32] The magnitude of the coefficient on share change growth during expansions is smaller than the corresponding magnitude in our calibrated model. This difference may partly reflect unmodeled labor market institutions such as minimum wage increases which might generate non-random changes in wage premia without corresponding changes in employment share.

Figure 10 – Model Impulse Response Function of Aggregate Employment in Recessions for Different Reallocation Shock Sizes



Notes: Each line scales the dispersion in productivities in panel A of figure 6 and thus the amount of labor reallocation at the national level. The top line corresponds to 80 percent of the baseline dispersion. Each further line adds increments of 20 percent, so the bottom line corresponds to 120 percent of the baseline dispersion. All reallocation shocks occur concurrently with an AR(1) interest rate wedge of 0.8p.p. with persistence 0.955.

aggregate fluctuations is consistent with a downward sloping Beveridge curve.

In interpreting our results it is important to distinguish between industry labor demand shocks and worker fluidity shocks as sources of worker mobility. Worker fluidity is the sum of the separation and hiring rates. Using a state-year panel, Davis and Haltiwanger (2014) find that an exogenous increase in fluidity raises aggregate employment. We simulate an increase in fluidity in the model by simultaneously increasing matching efficiency and the separation rate δ such that both δ and the job finding rate f rise by the same amount in percent terms. This shock raises both worker fluidity and aggregate employment, but leaves our measure of cross industry reallocation unchanged. The increase in aggregate employment follows immediately from the steady state approximation $u = \frac{\delta(1-f)}{\delta(1-f)+f}$. Moreover, following an industry reallocation shock, the separation rate remains unchanged but the aggregate hiring rate falls commensurately with the decline in aggregate matching efficiency. Thus, an increase in industry reallocation generates a decline in worker fluidity along with a decline in employment, again reproducing the positive correlation found by Davis and Haltiwanger (2014). Perhaps more importantly, higher exogenous fluidity in the model mitigates the adverse employment consequences of a reallocation shock by making industry transitions easier.

7. Conclusion

Reallocation of workers across industries matters to aggregate labor market outcomes during recessions. To establish this fact, we develop a methodology to isolate the mean-preserving

spread component of secular industry-specific shocks. The methodology identifies reallocation in a local area predicted by the local area's industry composition and national trends, and orthogonalizes this predicted reallocation with respect to the direct growth rate consequences of the local industry composition. We apply the methodology to local labor markets over a 35 year period. We find a robust effect of reallocation on aggregate labor market outcomes during periods when national employment is depressed, but not during national expansions.

We interpret the empirical results through the lens of a model of the labor market featuring decisions by job-seekers of which industry to search for work. The model delivers the following insights. First, absent labor market frictions or mechanisms which generate sticky wages, reallocation does not affect aggregate outcomes. Intuitively, when workers can seamlessly transition across industries, dispersion shocks result in immediate transitions to a new steady state. Second, plausible frictions result in marginal reallocation effects of similar magnitude to those found in our empirical results. Third, compression of wage differentials during recessions can explain the asymmetric response of aggregate employment to reallocation during recession-recoveries and expansions. Fourth, secular labor reallocation can have substantial effects on the persistence and amplification of national as well as local business cycles.

Our analysis has implications for policy not explored in this paper. The idiosyncratic industry shocks which underly secular industry reallocation include real shocks such as dispersion in productivity levels or consumer taste trends. Nonetheless, the ease with which reallocation occurs appears to depend sharply on the state of the business cycle. This interaction suggests a possible role for monetary policy in accommodating such shocks to "grease the wheels" of reallocation with higher inflation. Likewise, the interaction between worker fluidity and industry reallocation in our model suggests a possible role for policy in increasing the fluidity of labor markets to mitigate against the consequences of industry reallocation shocks. We leave further development of these conjectures and other implications to future work.

References

Aaronson, Daniel, Ellen Rissman, and Daniel Sullivan, "Can Sectoral Reallocation Explain the Jobless Recovery?," *Federal Reserve Bank of Chicago Economic Perspectives*, 2004, (2Q), 36–49.

Abraham, Katharine and Lawrence F Katz, "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?," *Journal of Political Economy*, 1986, *94*.

Ahn, Hie Joo and James D Hamilton, "Heterogeneity and Unemployment Dynamics," 2014.

Artuç, Erhan, Shubham Chaudhuri, and John McLaren, "Trade Shocks and Labor Adjustment: A Structural Empirical Approach," *American Economic Review*, 2010, *100* (3), 1008–45.

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

B., Jr. Long John and Charles I. Plosser, "Real Business Cycles," *Journal of Political Economy*,

1983, *91* (1), pp. 39–69.

Baqaee, David Rezza, "Asymmetric Inflation Expectations, Downward Rigidity of Wages, and Asymmetric Business Cycles," *Downward Rigidity of Wages, and Asymmetric Business Cycles (September 18, 2014)*, 2014.

Barnichon, Regis and Andrew Figura, "Labor Market Heterogeneity and the Aggregate Matching Function," *AEJ: Macroeconomics*, Forthcoming.

Barro, Robert J, "Long-term contracting, sticky prices, and monetary policy," *Journal of Monetary Economics*, 1977, *3* (3), 305–316.

Bartik, Timothy J., *Who Benefits from State and Local Development Policies?*, W.E. Upjohn Institute for Employment Research, 1991.

Berger, David, "Countercyclical Restructuring and Jobless Recoveries," 2014.

Blanchard, Olivier Jean and Lawrence F. Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, *1992* (1), pp. 1–75.

Brainard, S. Lael and David M. Cutler, "Sectorial Shifts and Cyclical Unemployment Reconsidered," *The Quarterly Journal of Economics*, 1993, *108* (1), pp. 219–243.

Caballero, Ricardo J. and Mohamad L. Hammour, "The Cost of Recessions Revisited: A Reverse-Liquidationist View," *The Review of Economic Studies*, 2005, *72* (2), pp. 313–341.

Charles, Kerwin, Erik Hurst, and Matt Notowidigdo, "Masking and Aggregate Employment Changes: Housing Booms and Manufacturing Decline During the 2000s," 2014.

Chetty, Raj, "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 2009, *1* (1), 451–488.

Chodorow-Reich, Gabriel and Loukas Karabarbounis, "The Cyclicality of the Opportunity Cost of Employment," 2015.

Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, "Unemployment and Business Cycles," 2015.

Daly, Mary C and Bart Hobijn, "Downward nominal wage rigidities bend the Phillips curve," *Journal of Money, Credit and Banking*, 2014, *46* (S2), 51–93.

Davis, Steven J. and John C. Haltiwanger, "Labor Market Fluidity and Economic Performance," 2014. Jackson Hole Economic Policy Symposium.

_ and John Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *The Quarterly Journal of Economics*, 1992, *107* (3), pp. 819–863.

DeLong, J. Bradford, "Identifying Cyclical vs. Structural Unemployment," Blog post August 2010.

_ and Lawrence H. Summers, "How Does Macroeconomic Policy Affect Output?," *Brookings Papers on Economic Activity*, 1988, *1988* (2), pp. 433–494.

Diamond, Rebecca, "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," 2015. Unpublished manuscript.

Dvorkin, Maximiliano, "Sectoral Shocks, Reallocation and Unemployment in Competitive Labor Markets," 2014. Unpublished manuscript.

Farmer, Tracy and Michael Searson, "Use of Administrative Records in the Bureau of Labor Statistics' Covered Employment and WWage (ES-202) Program," Technical Report, Bureau of the Census Annual Research Conference 1995.

Foster, Lucia, Cheryl Grim, and John Haltiwanger, "Reallocation in the Great Recession: Cleansing or Not?," 2014. NBER Working Paper No. 20427.

Fujita, Shigeru and Giuseppe Moscarini, "Recall Unemployment," 2013.

Gabaix, Xavier, "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 2011, *79*

(3), 733–772.

Ganong, Peter and Daniel Shoag, "Why Has Regional Income Convergence in the U.S. Declined?," 2015. Harvard University mimeo.

Garin, Julio, Mike Pries, and Eric Sims, "Reallocation and the Changing Nature of Economic Fluctuations," 2013.

Groshen, Erica and Simon Potter, "Has structural change contributed to a jobless recovery?," *Current Issues in Economics and Finance*, 2003, *9* (8).

Hall, Robert, "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 2005, *95* (1), 50–65.

_ and Sam Schulhofer-Wohl, "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers," 2015. NBER Working Paper No. 20939.

Hamilton, James D. and Michael T. Owyang, "The Propagation of Regional Recessions," *The Review of Economics and Statistics*, 2012, *94* (4), pp. 935–947.

Jaimovich, Nir and Henry Siu, "The Trend is the Cycle: Job Polarization and Jobless Recoveries," 2014.

Kahn, Shulamit, "Evidence of Nominal Wage Stickiness from Microdata," *The American Economic Review*, 1997, *87* (5), pp. 993–1008.

Kapan, Greg and Sam Schulhofer-Wohl, "Understanding the Long-Run Decline in Interstate Migration," *International Economic Review*, Forthcoming.

Kennan, John and James R. Walker, "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 2011, *79* (1), 211–251.

Kline, Patrick, "Understanding Sectoral Labor Market Dynamics: An Equilibrium Analysis of the Oil and Gas Field Services Industry," 2008. Unpublished manuscript.

Koenders, Kathryn and Richard Rogerson, "Organizational dynamics over the business cycle: a view on jobless recoveries," *Review*, 2005, (Jul), 555–580.

Krugman, Paul, "Jobs and Skills and Zombies," Blog post March 2014.

Lilien, David M, "Sectoral shifts and cyclical unemployment," *The Journal of Political Economy*, 1982, pp. 777–793.

Mehrotra, Neil R and Dmitriy Sergeyev, "Sectoral Shocks, the Beveridge Curve and Monetary Policy," *mimeo*, 2012.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak, "Internal Migration in the United States," *Journal of Economic Perspectives*, 2011, *25* (3), 173–96.

Murply, Kevin and Robert Topel, "The Evolution of Unemployment in the United States: 1968-1985," in "in" 1987, pp. 11–68.

Pilossoph, Laura, "Sectoral Shocks and Move Unemployment," 2014. Unpublished manuscript.

Pissarides, Christopher A., *Equilibrium Unemployment Theory*, The MIT Press, 2000.

Sahin, Aysegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante, "Mismatch Unemployment," *American Economic Review*, 2014, *104* (11), 3529–64.

Schmitt-Grohé, Stephanie and Martın Uribe, "Downward nominal wage rigidity, currency pegs, and involuntary unemployment," 2013.

Schumpeter, Joseph A., *Capitalism, Socialism, and Democracy* 1942.

Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *The American Economic Review*, 2005, *95* (1), pp. 25–49.

Walsh, Carl E, "Labor Market Search, Sticky Prices, and Interest Rate Policies," *Review of economic Dynamics*, 2005, *8* (4), 829–849.