

Secular Labor Reallocation and Business Cycles

Gabriel Chodorow-Reich

Harvard University and National Bureau of Economic Research

Johannes Wieland

University of California San Diego, Federal Reserve Bank of Chicago, and National Bureau of Economic Research

We revisit an old question: does industry labor reallocation affect the business cycle? Our empirical methodology exploits variation in a local labor market's exposure to industry reallocation on the basis of the area's initial industry composition and national industry employment trends for identification. Applied to confidential employment data over 1980–2014, we find sharp evidence of reallocation contributing to higher local area unemployment if it occurs during a national recession but little difference in outcomes during an expansion. A multi-area, multisector search-and-matching model with imperfect mobility across industries and downward nominal wage rigidity can reproduce these cross-sectional patterns.

I. Introduction

Industries 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

We thank David Berger, Paul Goldsmith-Pinkham, Gordon Hansen, Larry Katz, Pat Kline, Giuseppe Moscarini, Laura Pilossoph (discussant), Valerie Ramey, Lawrence Uren (discussant), Gianluca Violante, Ivan Werning, four anonymous referees, Harald Uhlig (editor), and numerous seminar and conference participants for their comments. This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. The views expressed

Electronically published April 23, 2020

[*Journal of Political Economy*, 2020, vol. 128, no. 6]

© 2020 by The University of Chicago. All rights reserved. 0022-3808/2020/12806-0005\$10.00

cycles, labor markets, and the scope for policy. Yet the issue remains unsettled.

We study the consequences of secular labor reallocation, defined as the change in an economy's allocation of labor in response to mean-preserving, long-lasting idiosyncratic industry shocks. We make two main contributions. First, we propose a novel method to estimate how secular labor reallocation affects local labor markets and implement it using confidential administrative employment data. Different from much of the literature exploring the "sectoral shifts" hypothesis, we examine whether the consequences of reallocation depend on the phase of the business cycle. We find that they do. More reallocation implies higher unemployment during a recession but implies roughly neutral effects when it occurs during an expansion. Second, we show that a multiarea, multisector search-and-matching model featuring realistic frictions to sectoral mobility and downward wage rigidity can rationalize this result.

Our analysis starts in section II with a description of our empirical identification strategy. A number of challenges arise. 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, realized 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). Third, we generally do not observe pure cross-industry dispersion shocks that do not also affect the mean of variables such as productivity. 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 isolate long-lasting shocks, our metric of reallocation sums the absolute value of industry employment share changes between the start and the end of a recession-to-recovery or expansion cycle, thereby filtering out cyclical changes that occur during a recession but reverse during a recovery. We address the endogeneity of reallocation to local conditions by developing a Bartik-style measure of predicted reallocation on the basis of a local area's initial industry composition and industry employment changes in the rest of the country and use this measure as an instrument for actual reallocation. Finally, we account for non-mean-preserving industry shifts by directly controlling for the Bartik predicted employment growth rate given an area's industry composition. Thus, our empirical specification regresses local area unemployment on

here do not necessarily reflect the views of the BLS or the US government. We are grateful to Jessica Helfand and Michael LoBue of the BLS for their help with the Longitudinal Database. The views expressed in this paper do not necessarily reflect the views of the Federal Reserve Bank of Chicago or of the Federal Reserve System. Data are provided as supplementary material online.

local area reallocation, controlling for predicted growth and with reallocation instrumented using predicted reallocation. Intuitively, the research design compares outcomes in areas with the same predicted employment growth but different predicted reallocation.

We implement our exercise using confidential employment data by local area and industry from the Bureau of Labor Statistics (BLS) Longitudinal Database (LDB) merged with the public-use counterpart of these data, the Quarterly Census of Employment and Wages (QCEW). We use the public-use version to extend the analysis back to 1979. The resulting data set tracks industry reallocation in more than 200 urban local labor markets. We describe these data in detail in section III.

Section IV contains our empirical analysis. Predicted reallocation is a strong instrument for actual reallocation except around the 1990 recession because of a change in the data collection procedure at that time. We therefore introduce a two-sample two-stage least squares (TS2SLS) design where we estimate the first stage excluding the 1990 episode but include it in the second stage and derive the standard error formula appropriate to this setting. Our formula should prove useful in other settings where researchers encounter missing data.

We obtain two main empirical results. First, higher reallocation causes higher unemployment. Second, this average response masks a crucial asymmetry. During a recession-to-recovery cycle, a 1 standard deviation increase in predicted reallocation raises unemployment by roughly 0.5 percentage points at the national recession trough, with the effect then dissipating over the subsequent recovery. In contrast, reallocation does not affect unemployment during an expansion. These results are statistically strong, are not driven by particular sectors or areas, and are robust to inclusion of local area time-varying control variables or local area fixed effects.

Section V introduces a multisector, multiarea model of reallocation and unemployment to provide a structural interpretation of our results. Each area in the model contains a number of industries consisting of firms and workers who interact according to a search-and-matching framework subject to a downward nominal wage constraint. The shares of workers and firms in each industry depend on industry-specific productivity and consumer preferences. In line with the data, the model features two-way gross flows of workers across industries in each period. We shock the model with an increase in the cross-sectional variance of industry-level productivities and estimate the same regression in the model as in the data. Specifically, we estimate the marginal effect of reallocation on unemployment during an expansion—in which the increase in the cross-sectional variance of industry productivity constitutes the only set of shocks—and a recession—in which borrowers simultaneously face an increase in the interest rate.

Without any labor market or wage-setting frictions, reallocation across industries would occur instantaneously and without generating any unemployment. Allowing for only within-industry search-and-matching frictions, the mean-preserving spread in industry productivities generates a small increase in unemployment regardless of whether it occurs coincident with a demand-induced recession. Incorporation of frictions to moving across industries and empirically plausible downward wage rigidity breaks this symmetry. Intuitively, during expansions higher wages draw job seekers into the expanding sectors, while wage compression during recessions pushes the adjustment into a greater difference in job-finding rates. We construct impulse response functions of the cross-area marginal effect of reallocation on unemployment in model-simulated data and find that they accord well with our empirical results.

Related literature.—This 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 Murphy and Topel (1987). Their critiques inform our methodological approach. Debate over the importance of sectoral reallocation has been renewed in the context of the slow recoveries from the two most recent national recessions.¹ Different from the Lilien (1982) question of whether business cycle downturns coincide with more restructuring, we find a connection between reallocation and the business cycle because of the greater ease with which labor markets absorb a given amount of reallocation when it occurs during an expansion rather than a recession-to-recovery cycle.

Methodologically, our paper follows Autor, Dorn, and Hanson (2013) and Charles, Hurst, and Notowidigdo (2014) in using industry shocks to local labor markets. Our paper differs in its focus on business cycle outcomes. As such, we construct a Bartik measure that does not rely on a specific source of sectoral reallocation. We validate our identification strategy following recommendations in Borusyak, Hull, and Jaravel (2018) and Goldsmith-Pinkham, Sorkin, and Swift (2018). Our results complement work on the consequences of reallocation at the worker level (Fujita and Moscarini 2013; Davis and Haltiwanger 2014; Jaimovich and Siu 2014).

¹ See, e.g., Groshen and Potter (2003); Koenders and Rogerson (2005); Mehrotra and Sergeyev (2012); Garin, Pries, and Sims (2013); and Berger (2014) for papers highlighting the importance of input reallocation, and see Aaronson, Rissman, and Sullivan (2004); Dvorkin (2014); Pilossoph (2014); and Hall and Schulhofer-Wohl (2015) for an opposing view. Sahin et al. (2014) stake a middle ground using an empirical decomposition. A related literature views secular sectoral shocks as inevitable and a diversified industrial base as a necessary condition for a city to be able to reinvent itself when such shifts occur (Glaeser 2005). Our results do not dispute the long-run benefits of having a diversified industrial base but instead point out that the cost of undergoing such a reinvention depends on the phase of the business cycle.

Our general equilibrium search-and-matching model with nominal frictions builds on Christiano, Eichenbaum, and Trabandt (2016) and earlier work by Walsh (2005). We incorporate an industry structure and labor reallocation frictions following Kline (2008), Dvorkin (2014), and Pilosoph (2014). Downward nominal wage rigidity has recently been emphasized by Daly and Hobijn (2014) and Schmitt-Grohé and Uribe (2016), and following Hall (2005) our implementation does not violate bilateral efficiency conditions.

The importance of wage rigidity in our model leads to conclusions that differ from existing literature. A popular account suggests that desired reallocation 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 desired reallocation is magnified. Reallocation in the model also causes total vacancies to fall, a result at odds with the claim in Abraham and Katz (1986) that rising vacancies are the signature of reallocation. Closer to our mechanism, Jackman and Roper (1987), Shimer (2007), and Sahin et al. (2014) also emphasize “mismatch unemployment” caused by a dispersion in job-finding rates across sectors.

II. Measurement and Empirical Strategy

We define a measure of reallocation across industries and then discuss our empirical strategy.

A. Measure of Reallocation

We define an index of reallocation based on the dispersion in industry employment growth rates as in Lilien (1982). The economy consists of A areas, each with I industries. Let $e_{a,i,t}$ denote employment in area a and industry i at time t , $e_{a,t} = \sum_{i=1}^I e_{a,i,t}$ total employment in the area, $s_{a,i,t} = e_{a,i,t}/e_{a,t}$ industry i 's employment share, $g_{a,i,t,t+j} = e_{a,i,t+j}/e_{a,i,t} - 1$ the area-industry employment growth rate, and $g_{a,t,t+j} = e_{a,t+j}/e_{a,t} - 1$ total local area employment growth. Reallocation in area a between months t and $t + j$ is

$$R_{a,t,t+j} = \frac{12}{j} \frac{1}{2} \sum_i s_{a,i,t} \left| \frac{1 + g_{a,i,t,t+j}}{1 + g_{a,t,t+j}} - 1 \right|. \quad (1)$$

The measure $R_{a,t,t+j}$ is easily interpreted. The term $1/2 \sum_{i=1}^I |(1 + g_{a,i,t,t+j})/(1 + g_{a,t,t+j}) - 1|$ equals zero if employment grows at an identical rate in every industry between t and $t + j$ and 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 ratio $12/j$ translates the reallocation between t and $t + j$ into an annualized monthly flow, such that $R_{a,t,t+j} \subseteq [0, 12/j]$.²

B. Econometric Approach

We assume that unemployment and reallocation in local area a evolve according to

$$\Delta u_{a,t,t+j} = \beta_j R_{a,t,t+k} + \mathcal{G}(\{s_{a,i,t}\}; \{\eta_{i,t,t+j}\}) + \Gamma^u X_{a,t} + \varepsilon_{a,t,t+j}, \quad (2)$$

$$R_{a,t,t+k} = \mathcal{C}(u_{a,t}, u_{a,t+1}, \dots, u_{a,t+k}) + \mathcal{R}(\{s_{a,i,t}\}; \{\eta_{i,t,t+k}\}) + \Gamma^R X_{a,t} + \nu_{a,t,t+k}, \quad (3)$$

where $\Delta u_{a,t,t+j} = u_{a,t+j} - u_{a,t}$, $\mathcal{G}(\{s_{a,i,t}\}; \{\eta_{i,t,t+j}\})$, and $\mathcal{R}(\{s_{a,i,t}\}; \{\eta_{i,t,t+k}\})$ are functions that relate cumulative idiosyncratic industry-level shocks $\{\eta_{i,t,t+k}\}$ and local area employment shares to local area unemployment and reallocation; $X_{a,t}$ contains time t measurable observed variables that affect unemployment or reallocation; and $\varepsilon_{a,t,t+j}$ and $\nu_{a,t,t+k}$ are unobserved area-specific determinants. We write $\eta_{i,t,t+k}$ without an a subscript to emphasize that these shocks occur at the national level; shocks to an industry specific to a particular area are subsumed in $\nu_{a,t,t+k}$. While we do not need to specify the functional forms of $\mathcal{G}$ and $\mathcal{R}$, it is natural and convenient to think of $\mathcal{G}$ as a function solely of the weighted mean shock in an area $\sum_i s_{a,i,t} \eta_{i,t,t+j}$ and $\mathcal{R}$ as a function of the weighted dispersion in $\{\eta_{i,t,t+k}\}$.

Our object of interest is β_j , the effect of reallocation on unemployment. Previous literature has emphasized two sources of causality running instead from unemployment to reallocation, captured by the $\mathcal{C}(\cdot)$ function in equation (3). First, a low opportunity cost of restructuring during periods of high unemployment and weak demand may lead to higher reallocation (Schumpeter 1942; Berger 2014). Second, a demand-induced recession can cause cyclical reallocation across industries if industries differ in their cyclical sensitivities (Abraham and Katz 1986). More generally, both unemployment and reallocation may depend on common determinants, as captured by the common arguments in the $\mathcal{G}$ and $\mathcal{R}$ functions. For example, an area with a large manufacturing base in 1980 may have

² The measure defined in eq. (1) has an equivalent representation in terms of changes in industry employment shares, $R_{a,t,t+j} = (12/j)(1/2\sum_{i=1}^I |s_{a,i,t+j} - s_{a,i,t}|)$. The measure differs slightly from Lilien (1982); Lilien measures reallocation only period by period, corresponding to $R_{a,t,t+1}$, and Lilien sums the squares of industry growth rate dispersion, whereas we sum absolute values to reduce sensitivity to outliers. The measure $R_{a,t,t+1}$ also equals the Davis and Haltiwanger (1992) cross-industry job-reallocation rate if total employment remains unchanged between the two periods. Appendix B (apps. B–F are available online) contains additional details.

had high unemployment directly as a result of concentrating in industries that had negative labor demand shocks while also experiencing substantial reallocation over the past few decades.

Double Bartik.—We introduce a “double Bartik” strategy to overcome these difficulties. Following Bartik (1991) and a large subsequent literature, we define Bartik predicted employment growth as the employment growth in local area a if employment in each local industry grew at exactly the same rate as employment in that industry in the rest of the country,

$$g_{a,t,t+j}^b = \frac{12}{j} \sum_{i=1}^I s_{a,i,t} g_{-a,i,t,t+j}^-, \quad (4)$$

where $e_{-a,i,t}$ denotes employment in industry i at time t summing over all areas other than area a and $g_{-a,i,t,t+j}^- \equiv e_{-a,i,t,t+j}/e_{-a,i,t} - 1$ is the leave-one-out employment growth rate for industry i . Substituting $g_{-a,i,t,t+j}^-$ and $g_{a,t,t+j}^b$ into equation (1), we analogously define Bartik predicted reallocation as the reallocation that would be obtained in area a if employment in each local industry grew at exactly the same rate as employment in that industry in the rest of the country,

$$R_{a,t,t+k}^b = \left(\frac{12}{k} \right) \left(\frac{1}{2} \sum_{i=1}^I s_{a,i,t} \left| \frac{1 + g_{-a,i,t,t+k}^-}{1 + g_{a,t,t+k}^b} - 1 \right| \right). \quad (5)$$

We believe that this second Bartik measure is original.³

We use Bartik predicted reallocation $R_{a,t,t+k}^b$ as an excluded instrument for $R_{a,t,t+k}$ and Bartik predicted employment growth as a control (i.e., included instrument) in a specification including period fixed effects. To understand this specification, suppose that we knew the functions $\mathcal{G}$ and $\mathcal{R}$ and observed $\{\eta_{i,t,t+k}\}$. Then, so long as $\mathcal{G}$ and $\mathcal{R}$ were not collinear, estimating equations (2) and (3) with $\mathcal{R}$ as an excluded instrument and $\mathcal{G}$ as an included instrument would identify β_j from the part of local reallocation due to $\mathcal{R}$ and after controlling for the direct effect of industry shocks on local area unemployment that occur through the function $\mathcal{G}$. Since we do not actually observe these shocks or know these functions, we assume instead that Bartik growth and reallocation, measured over appropriate time spans, can proxy for the functions $\mathcal{G}$ and $\mathcal{R}$. Indeed, Bartik reallocation being a good proxy for $\mathcal{R}$ implies the relevance condition $\text{Cov}(R_{a,t,t+k}, R_{a,t,t+k}^b | g_{a,t,t+k}^b, X_{a,t}, \alpha_t) > 0$.⁴ The exclusion restriction

³ The closest antecedent of which we are aware comes from Davis and Haltiwanger (2014), who develop an instrument for cross-establishment job reallocation based on the interaction of lagged industry employment shares and the job-reallocation rate within each industry.

⁴ As we show below, the first-stage coefficient on Bartik reallocation is close to one. Therefore, our empirical results can also be interpreted as arising from $\mathcal{R}$ entering directly into eq. (2), as occurs in the model in sec. V, where $\mathcal{R}$ has the interpretation of the “desired” amount of reallocation.

$\text{Cov}(\varepsilon_{a,t,t+j}, R_{a,t,t+k}^b | g_{a,t,t+k}^b, X_{a,t}, \alpha_t) = 0$ requires that Bartik growth absorb all direct effects of industry shocks on unemployment that occur through $\mathcal{G}$. Furthermore, specifying predicted reallocation as an excluded instrument solves the problem of feedback from local unemployment to reallocation because (with time fixed effects) predicted reallocation does not depend on local outcomes after time t , while including $g_{a,t,t+k}^b$ controls for both the area's industrial cyclical sensitivity and the possibility that predicted reallocation concentrates in areas also undergoing secular decline or expansion. Intuitively, the research design compares areas with the same predicted growth but different predicted reallocation.

Reallocation timing.—If labor market frictions impede reallocation, then the pattern of national industry employment growth could lag shifts in $\{\eta_{i,t,t+k}\}$ and make Bartik growth and reallocation, which depend on realized national industry employment growth rates, poor proxies for the functions of the underlying shocks. To address this potentiality, we measure both actual and predicted reallocation over two separate, multi-period windows. The first window begins at a national employment peak and lasts through the course of a national recession and subsequent labor market recovery. The second window begins when the labor market has fully recovered and ends at the start of the next recession. Therefore, we assume that actual national employment reallocation between the start and the end of a recession-to-recovery cycle and over an expansion fully reflects the reallocation that would eventually occur as a result of the idiosyncratic industry shocks, so that the Bartik reallocation instrument embodies the reallocation in an area implied by the national industry shocks.⁵ This timing also filters out temporary reallocation induced by differing cyclical sensitivities.

We define a national labor market recession as the period between a private sector employment peak and the employment trough, a recovery as the period from the trough until the economy regains its previous peak level, and an expansion as the period between the end of a recovery and the start of the next recession. Thus, we measure reallocation over the 2005–8 expansion using the growth rates of industry employment between June 2005 and the private sector employment peak in January 2008 and reallocation during the 2008–14 recession-to-recovery cycle using the growth rates of industry employment between January 2008 and March 2014 when employment first regains its January 2008 level. Figure 1 illustrates the labor market recessions, recoveries, and expansions in our sample. (We treat the 1980–82 period as a single long recession.) The view of cyclical tightness as the same at the start and the end of each recession-to-recovery and expansion cycle echoes the “gaps” view of business cycles advocated by DeLong and Summers (1988). Our main results are not

⁵ This is the case in our model, as we show in fig. D.1 (figs. D.1–D.4 are available online).




FIG. 1.—National business cycle partition. A color version of this figure is available online.

sensitive to this particular partitioning.⁶ We apply the same national timing to all local areas to compute predicted and realized reallocation, allowing for the interpretation of our regressions as pooled cross sections.⁷

By construction, national reallocation during a recession-to-recovery cycle is mean preserving in overall employment. Measuring reallocation between two periods when total employment remains unchanged facilitates a natural economic interpretation, since

$$R_{a,i,t,t+T} \Big|_{e_{a,i,t} = e_{a,i,t+T}} = \frac{12}{T} \frac{1}{2} \sum_{i=1}^I \left| \frac{e_{a,i,t+T} - e_{a,i,t}}{e_{a,i,t}} \right|. \quad (6)$$

Equation (6) rewrites $R_{a,i,t,t+T}$ as the minimum fraction of total period t employment that changes industries between t and $t + T$, expressed

⁶ We use the term “recession” to refer to the period between the private sector employment peak and the trough with the understanding that this definition differs from the periods designated by the National Bureau of Economic Research. We report robustness to other timing conventions in sec. IV.C. We prefer the timing procedure described above for two reasons. First, when $e_t = e_{t+j}$, the predicted reallocation measure $R_{a,i,t,t+j}^b$ has a natural interpretation as described shortly. Second, we do not see an obvious alternative for how to adjust for demographic trends. For example, not only had the national employment-to-population ratio not recovered its prerecession level as of the end of 2016, but the peak of the series predates the 2000 recession as well. Similarly, the employment-to-population ratio for prime-aged males did not regain its previous peak following any downturn since 1975.

⁷ Using local business cycle timing could induce a feedback from local demand/supply shocks to predicted reallocation through the length of the local cycle and thereby reintroduce reverse causality. Because national and local cycles are highly correlated, our partitioning captures much of the variation in local business cycles. For example, applying the same recession-to-recovery/expansion partition to local areas, in our sample a monthly average of 75% of local areas are in a local recession-to-recovery cycle when the national economy is in a recession-to-recovery cycle and 68% of local areas are in a local expansion when the national economy is in an expansion.

as a monthly flow at an annual rate. When $e_{-a,t} = e_{-a,t+T}$, a derivation similar to equation (6) shows that predicted reallocation $R_{a,t,t+T}^b$ has the interpretation of the predicted net quantity of industry employment reshuffling between t and $t + T$ as a share of total employment at t , expressed at an annual rate.

Other specification details.—An important element of our analysis will be to allow the effect of reallocation on unemployment to vary by the phase of the business cycle. We argue in section V that the differences in β across business cycle phases are informative about the underlying economic mechanisms.

Finally, varying the horizon of the unemployment response and holding fixed the horizon over which reallocation occurs traces out an impulse response function. That is, for months t that mark the start of a recession or an expansion and letting T denote the length of the recession-to-recovery cycle or the expansion, we estimate for $c \in \{\text{Recession-to-recovery}, \text{Expansion}\}$,

$$\Delta u_{a,t,t+j} = \beta_{j,c} R_{a,t,t+T} + g_{a,t,t+T}^b + \Gamma_{j,c}^u X_{a,t} + \varepsilon_{a,t,t+j}. \quad (7)$$

C. Discussion

The Bartik research design has the advantage of not requiring the researcher to take a stand on the deep industry-level determinants of reallocation in any given period (the η_i in our notation), such as changes in technology, consumer tastes, exchange rates, or trade policy. Rather, the evolution of employment nationally summarizes the consequences of the combination of these deep determinants for reallocation. The Bartik approach simply requires that the deep determinants produce a common component of industry employment growth across areas and that, after residualizing with respect to predicted area growth, these determinants affect local areas only through their effect on reallocation. Not needing to link reallocation to a primitive shock makes this approach well suited to the study of business cycle frequency outcomes that span multiple cycles, each with its own unique deep determinants of reallocation.

This aspect of the research design also introduces two important limitations. First, we study the consequences of reallocation but do not attempt to identify its primitive causes and therefore cannot answer how a policy maker might manipulate reallocation if desired. Related, the representation (2) and (3) may not uniquely characterize the system, as unemployment ultimately depends on the underlying shocks. Nonetheless, our results shed light on a long-running debate about whether the frictional reallocation of labor exerts an independent impact on unemployment and the frictions that give rise to such effects. Second, despite the common usage of the phrase “Bartik shock” (which we have purposefully

avoided), neither Bartik predicted employment growth nor predicted reallocation necessarily constitutes a shock in the standard meaning of being unanticipated and orthogonal to previous outcomes. In our setting, anticipation effects—knowledge of the industry shocks $\{\eta_{i,t,t+k}\}$ before time t —would complicate the interpretation of β if unemployment begins to rise before the reallocation occurs. In our empirical work, we test for differential pretrends to diagnose such anticipation effects and find little evidence for them. The model in section V expands on both of these issues by providing a concrete example of a set of primitive shocks $\{\eta_{i,t,t+k}\}$ that give rise to our measure of reallocation. Finally, two recent contributions—Borusyak, Hull, and Jaravel (2018) and Goldsmith-Pinkham, Sorkin, and Swift (2018)—formalize identification arguments and recommend validation procedures in a Bartik setting.⁸ We carry out their recommendations as part of our robustness exercises and find results favorable to our identification strategy.

III. Data and Summary Statistics

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

A. Data

Data on employment by county and industry come from the BLS LDB and the QCEW. The LDB reports employment by establishment and month starting in 1990. The source data come from quarterly reports that employers file 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 six-digit North American Industry Classification System (NAICS) code associated with its primary activity. Our LDB sample contains 42 states that allow access to their data through the BLS visiting researcher confidential data access program. These data are uniquely suited to measuring reallocation because they do not contain sampling error, which would artificially increase reallocation rates.

⁸ Goldsmith-Pinkham, Sorkin, and Swift (2018) use an equivalence result between a shift-share instrumental variable (IV) and a weighted sum of coefficients from just-identified IV regressions, each with the area share in one industry as the excluded instrument, to characterize exogeneity of local industry shares as a sufficient condition for identification. Borusyak, Hull, and Jaravel (2018) show that exogeneity of national industry trends is also sufficient for validity of a Bartik estimator. The equivalence results in these papers do not directly apply to the reallocation Bartik defined in eq. (5) because $R_{n,t,t+k}^b$ does not have a shift-share structure due to the location-specific subtraction of the Bartik growth rate inside the absolute value. They do apply to a close variant that contains only the national growth rate inside the absolute value in eq. (5).

The QCEW is the public-use version of the LDB. It reports monthly employment at the industry-county level for all 50 states starting in 1975, subject to disclosure limitations to prevent the release of identifying information regarding single establishments.⁹ We use the QCEW to extend the sample back to 1979 and to fill in states not in our LDB sample.

Two details of the data collection procedure merit mention as they affect 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 1980.¹⁰ We exclude data before 1978 because the staggered implementation of the coverage expansion across states produces substantial measurement difficulties during that period. In effect, we exclude the 1976–80 expansion from the analysis. Second, in 1990 and 1991 the BLS lowered the threshold requirements for multiestablishment employers to report employment by single establishment (Farmer and Searson 1995). As a result, an unusually high number of establishments changed industry code during those years. While predicted reallocation during the 1990–93 recession-to-recovery cycle 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.¹¹ We instead develop a TS2SLS estimator, where we estimate the first stage excluding the 1990–93 period as described further below.

We combine the LDB data with NAICS three-digit employment from the QCEW for counties in states that are not in the LDB and with two-digit Standard Industrial Classification (SIC) data for 1975–2000.¹² We seasonally adjust all series at the industry-county level using the multistep moving average approach contained in the Census Bureau's X-11 algorithm. Relative to other data sets with employment by geography and industry, such

⁹ Even at the NAICS two-digit level and with counties already aggregated into metropolitan statistical areas (MSAs), roughly one-fifth of potential cells are suppressed for disclosure reasons; the suppressed share rises to 35% for MSA-industry cells at the NAICS three-digit level.

¹⁰ See <http://www.bls.gov/cew/cewbultncur.htm#Coverage>.

¹¹ 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. Separately, the NAICS version of the QCEW contains a number of transcription errors before 2001 that do not appear in the LDB and that we hand-correct. Appendix F lists these errors.

¹² The QCEW reports employment by county and SIC two-digit industry beginning in 1975 and by three- and four-digit industry for 1984–2000. We date the 1980s expansion as beginning in October 1983, making the introduction in 1984 of the SIC three- and four-digit industry detail redundant for our analysis. The 1987 revision of the SIC made large changes to a handful of industry definitions that, 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. In our analysis, we always interact period fixed effects with the classification (NAICS or SIC) to account for any level differences in reallocation across the two systems.

as the Census Bureau's County Business Patterns or Longitudinal Business Database (LBD), the BLS data have the important advantage for business cycle analysis of providing monthly rather than annual frequency. We choose SIC 2/NAICS 3 as our level of industry detail because our measure of reallocation does not distinguish between movement across similar or dissimilar industries. The SIC 2/NAICS 3 level allows for enough industry detail (roughly 80 industries) to generate variation in reallocation across areas while ensuring that all such reallocation occurs across broadly defined industries. A finer level of detail would also diminish our ability to make any use of the public data.

We aggregate county-level data into core-based statistical areas (CBSAs) using the 2013 Office of Management and Budget (OMB) county classifications. The 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). 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 main outcome variable is the local unemployment rate and comes from the BLS Local Area Unemployment Statistics (LAUS) program. We combine published data starting in 1990 with unpublished data available from the LAUS office for 1976–89. For 1990–present, the LAUS office provides seasonally adjusted data for MSAs; we augment these data by seasonally adjusting the county data using the same procedure described above for the 1976–89 period and for counties not in an MSA and aggregate up to the MSA or CSA level. While the construction of county and MSA unemployment rates involves imputation, any noise is likely to be classical left-hand-side error, and the unemployment rate offers conceptual advantages by reducing the effect of migration on the analysis.

Our final sample includes all MSAs and CSAs containing at least one MSA, with employment of at least 50,000 in 1 month, an agricultural share of employment of less than 20%, and where we observe at least 95% of private-sector employment at the industry level.¹³ The final sample contains 1,314 of the 3,144 counties in the United States and covers 86% of 2013 employment.

¹³ 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 and for the period 1980–89 when we do not have confidential data. As a result, our sample contains fewer CSAs/MSAs in the 1980–89 period than thereafter.

B. *Trends in National Reallocation*

An overview of reallocation at the national level provides useful context for what follows. Table 1 reports national reallocation, $R_{US,i,t+j}$, for each recession-to-recovery cycle and expansion and at various levels of industry aggregation. The rows with all data in boldface indicate the recession-to-recovery episodes. 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 2 with NAICS 3 and group SIC 4 with NAICS 6 on the basis of similarity in the number of industries. Reallocation measures for the overlapping episodes appear roughly comparable across these definitions.

A number of interesting patterns emerge. First, cross-industry reallocation occurs all the time. Since Lilien (1982), a debate has continued about whether sectoral employment shifts concentrate enough during periods of low economic activity to explain fluctuations in the business cycle. The problem identified by Abraham and Katz (1986) of how to account for different industry cyclical sensitivities during recessions makes answering this question difficult. By filtering out cyclical reallocation that occurs during a recession but reverts during a recovery, our timing approach provides one way around the Abraham and Katz (1986) critique.¹⁴ Using our approach, more secular reallocation does occur during episodes containing recessions, qualitatively consistent with the Lilien conjecture.

Second, consistent with a downward trend in a number of measures of labor market flows (Davis, Faberman, and Haltiwanger 2012; Molloy et al. 2016), the rate of reallocation has trended down. For example, 4.6% ($1.29 \times 43/12$) of employment changed SIC two-digit industry between March 1980 and October 1983. The same fraction changed NAICS three-digit industry between January 2008 and February 2014, despite the latter recession-to-recovery 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 for recession-to-recovery cycles. Despite the widespread attention to industry reallocation during the 2008–14 episode, our measure of secular reallocation suggests a decline in reallocation intensity

¹⁴ Alternatively, see Brainard and Cutler (1993); Aaronson, Rissman, and Sullivan (2004); and Mehrotra and Sergeyev (2012) for articles that apply parametric time series models to either the cyclical or the trend component of employment shares to address this question. Note, however, that the comparison of recession-to-recovery cycles and expansions in table 1 does not exclude the possibility that secular reallocation concentrates during recession-to-recovery cycles because of Schumpeterian restructuring. That is, while our timing solves the Abraham and Katz (1986) critique, it does not address other endogeneity concerns. For that, we turn to local variation.

TABLE 1
REALLOCATION BY EPISODE AND INDUSTRY DETAIL

EPISODE	MONTHS	EXPANSION	INDUSTRY DEFINITION						
			SIC 1.5	NAICS 2	SIC 2	NAICS 3	NAICS 4	SIC 4	NAICS 6
March 1980–October 1983	43	No	1.14		1.29			1.68	
October 1983–March 1990	77	Yes	.71		.93				
March 1990–April 1993	37	No	.82	1.04	.97	1.15	1.32	1.34	1.56
April 1993–December 2000	92	Yes	.42	.60	.85	.77	.95	1.14	1.13
December 2000–May 2005	53	No		.80		.97	1.25		1.42
May 2005–January 2008	32	Yes		.60		.73	1.00		1.20
January 2008–March 2014	74	No		.64		.71	.87		1.03
$R^2: \frac{1}{j} \Delta s_{i,t,t+j} = \alpha_i + \varepsilon_{i,t,t+j}$			.72	.71	.73	.59	.65	.75	.65
Industry count			18	20	73	92	305	963	1,028

NOTE.—This table reports values of $R_{US,t,t+j}$ for all complete national recession-to-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. The rows with all data in boldface indicate the recession-to-recovery episodes.

during the Great Recession period (see also Foster, Grim, and Haltiwanger 2016).¹⁵

Third, a large amount of reallocation occurs across broadly defined industries. For example, of the 6.5% (1.06% per year multiplied by 6.08 years) of employment changing six-digit NAICS industry between January 2008 and March 2014, 4.1 percentage points constituted movement across two-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 values from the regression $j^{-1}|\Delta s_{i,t,t+j}| = \alpha_i + \varepsilon_{i,t}$. For example, at the NAICS 3 level, the R^2 value of this regression is 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.

C. Local Reallocation

Figure 2 shows a map of the variation in predicted reallocation during the 2008–14 recession-to-recovery cycle.¹⁶ We split the MSA/CSA observations into quintiles based on their Bartik reallocation and mark higher reallocation levels with darker shading. The map shows that predicted reallocation is not easily explained by geographic factors.

Table 2 reports the pairwise correlations in predicted local reallocation across each national recession-to-recovery and expansion cycle. Predicted reallocation has a modest correlation across most episodes. These patterns inform our analysis in three ways. First, in some specifications we will exploit the absence of perfect serial correlation by including area fixed effects. Second, we will always cluster all standard errors by CSA/MSA to account for arbitrary correlation within an area over time.¹⁷

¹⁵ Interpreted through the lens of our empirical exercises and our model, the decline in reallocation intensity during the Great Recession translates into a similar increase in unemployment due to reallocation. This is because the larger aggregate shock during the Great Recession tightens the downward wage constraint more than in an average recession. Therefore, a given amount of reallocation generates more unemployment during the Great Recession, which compensates for the decline in reallocation intensity.

¹⁶ 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.

¹⁷ Because we control for Bartik growth, which uses all the information in national industry employment growth rates, our standard errors do not suffer from the criticism raised in Adão, Kolesár, and Morales (2018).




FIG. 2.—Map of predicted reallocation per year, 2008–14 cycle. This figure shows the geographic distribution of predicted reallocation per year for the national employment peak in January 2008. Because of disclosure limitations, for this figure we use only data from the public-use QCEW and require a minimum industry employment coverage of 80%. A color version of this figure is available online.

TABLE 2
CORRELATION OF PREDICTED REALLOCATION ACROSS EPISODES

	1980–83	1983–90	1990–93	1993–2000	2000–2005	2005–8	2008–14
1980–83	1.00						
1983–90	.24	1.00					
1990–93	.21	.09	1.00				
1993–2000	–.02	.52	–.16	1.00			
2000–2005	.30	.26	.07	.43	1.00		
2005–8	.17	.35	.04	.43	.52	1.00	
2008–14	.10	.15	.05	.26	.49	.71	1.00

NOTE.—This table reports the correlation of predicted reallocation across different episodes.

Third, in interpreting our findings and in the model in section V, we will not assume that reallocation at the start of a cycle is unanticipated.

IV. Empirical Results

A. First-Stage and Two-Sample 2SLS

Table 3 presents first-stage regressions of actual reallocation on predicted reallocation, controlling for predicted growth and period fixed effects. Column 1 pools over the three recession-to-recovery cycles 1980–83, 2000–2005, and 2008–14. Predicted reallocation has strong explanatory power, with a partial *F*-statistic of 15.5. As discussed in section III.A, the change between 1990 and 1993 in the minimum employment level requiring multiestablishment employers to report employment separately

TABLE 3
FIRST-STAGE REGRESSIONS

DEPENDENT VARIABLE: ACTUAL REALLOCATION, R_a				
	Recession-to-Recovery Cycles		Expansions	Pooled
	(1)	(2)		
Right-hand-side variables:				
Predicted reallocation	.98** (.25)	.21 (.24)	.66** (.22)	.80** (.18)
Episodes	Excluding 1990–93	All	All	Excluding 1990–93
National cycle fixed effects	Yes	Yes	Yes	Yes
Predicted growth	Yes	Yes	Yes	Yes
R^2	.38	.08	.22	.38
CSA/MSA clusters	217	218	218	220
Observations	534	748	557	1,091

NOTE.—Predicted reallocation is the reallocation measure R_a^b . Standard errors are given in parentheses and clustered by CSA/MSA.

** Significant at the 1% level.

for each establishment renders actual reallocation at a local level during that period too noisy to use. Column 2 illustrates the problem. Including the 1990–93 episode in the regression lowers the first-stage coefficient from 0.98 to 0.21 and yields a partial F -statistic below one, far too small to provide reliable second-stage estimates.

To include the 1990–93 episode in the subsequent analysis, we introduce a two-sample IV framework (Angrist and Krueger 1992). TS2SLS estimates the first-stage regression in one sample and applies the estimated first-stage coefficients to the instruments in a second sample used for the second-stage analysis. Thus, for recession-to-recovery episodes, the first-stage regression will exclude the 1990–93 cycle while the second stage will include it. Implicitly, we assume that the true first-stage coefficient does not differ during the 1990–93 cycle. While not directly testable, supportive evidence comes from the tight range of coefficients (between 0.88 and 1.08) obtained from estimating the first-stage regression separately in each of the other recession-to-recovery cycles.

Column 3 shows the first stage for the expansion cycles. The partial F -statistic lies just below the Staiger and Stock (1997) rule of thumb of 10. Including covariates, as we do below, raises the first-stage partial F -statistic substantially. Finally, column 4 shows the first stage pooling over the recession-to-recovery (excluding 1990–93) and expansion episodes.

This setup requires an econometric contribution, which we briefly describe. We differ from the typical TS2SLS framework in which the first- and second-stage samples contain entirely separate observations and the regression residuals are assumed to be independent across the two samples. Appendix A derives the asymptotic variance of the second-stage coefficients when independence does not hold. The expression that we derive is, to our knowledge, original and nests formulas in Inoue and Solon (2010) and Pacini and Windmeijer (2016) when the two samples are independent.

B. *Effects of Reallocation*

We start by reporting the effects of reallocation without distinguishing by the phase of the business cycle. Table 4 reports TS2SLS regressions of the following form:

$$\text{Second stage: } u_{a,t+j} - u_{a,t} = \beta \hat{R}_{a,t,t+T} + \Gamma' X_{a,t} + \alpha_t + \varepsilon_{a,t,t+j}, \quad (8)$$

$$\text{First stage: } R_{a,t,t+T} = \pi_1 R_{a,t,t+T}^b + \pi_2' X_{a,t} + \delta_t + \eta_{a,t,t+T}, \quad (9)$$

where $u_{a,t+j} - u_{a,t}$ is the change in the area a unemployment rate between the national recession peak and trough (recession-to-recovery episode) or during the first 30 months of the national expansion and $\hat{R}_{a,t,t+T} = \hat{\pi}_1 R_{a,t,t+T}^b + \hat{\pi}_2' X_{a,t} + \hat{\delta}_t$ is the cross-sample fitted value of reallocation

TABLE 4
HOMOGENOUS EFFECTS OVER CYCLE

	DEPENDENT VARIABLE: CHANGE IN UNEMPLOYMENT RATE		
	(1)	(2)	(3)
Right-hand-side variables:			
Reallocation	.44 ⁺ (.25)	.39* (.15)	.60** (.18)
National cycle fixed effects	Yes	Yes	Yes
Predicted growth	Yes	Yes	Yes
Additional controls	No	Yes	No
Area fixed effects	No	No	Yes
First-stage coefficient	.80	1.18	1.11
First-stage <i>F</i> -statistic	18.8	67.7	29.5
CSA/MSA clusters	220	220	220
First-stage observations	1,091	1,091	1,091
Second-stage observations	1,305	1,305	1,305

NOTE.—This table reports TS2SLS regressions. The dependent variable is the change in the unemployment rate between the national recession peak and trough (recession-to-recovery episode) or during the first 30 months of the national expansion. Additional controls in col. 2 are lags of employment growth, population growth, and house price growth, each measured from 5 years to 1 year before the cycle start; area size, measured by the log of sample mean employment; and the Herfindahl of industry concentration at the cycle start. Standard errors are given in parentheses and clustered by CSA/MSA.

⁺ Significant at the 10% level.

* Significant at the 5% level.

** Significant at the 1% level.

obtained by applying the first-stage regression coefficients to the variables in the second-stage sample. The 30-month horizon corresponds to the mean peak-to-trough length in our sample. The endogenous variable $R_{a,t,t+T}$ and excluded instrument $R_{a,t,t+T}^b$ measure the monthly flow of reallocation and predicted reallocation, respectively, between the beginning and the end of the national recession-to-recovery or expansion cycle. Both our simulated model in section V and the empirical impulse response function point to the national trough as the point at which the effect of reallocation reaches its maximum impact during a recession-to-recovery cycle, so we start our analysis at this horizon. We will shortly report the impulse response function for recession-to-recovery cycles.

Table 4 shows that on average over the business cycle, reallocation results in higher unemployment. Column 1 is our most parsimonious specification and includes in $X_{a,t}$ only predicted growth variables measured over the same horizon as reallocation and over the same horizon as the dependent variable. In anticipation of our next result showing how the effects of reallocation vary by phase of the business cycle, we interact each covariate with an indicator for national recession-to-recovery or expansion cycle, so that any difference between the effects of reallocation in this table and the next comes only from allowing the effect of reallocation to vary. The coefficient of 0.44 means that a marginal 1 percentage point

TABLE 5
HETEROGENEOUS EFFECTS OVER CYCLE

	DEPENDENT VARIABLE: CHANGE IN UNEMPLOYMENT RATE		
	(1)	(2)	(3)
Right-hand-side variables:			
Reallocation $\times$ expansion	-.40 (.35)	-.10 (.16)	-.32 (.34)
Reallocation $\times$ recession-to-recovery	.87** (.25)	.87** (.23)	.91** (.21)
<i>p</i> -value of equality	.004	.001	.004
National cycle fixed effects	Yes	Yes	Yes
Predicted growth	Yes	Yes	Yes
Additional controls	No	Yes	No
Area fixed effects	No	No	Yes
CSA/MSA clusters	220	220	220
First-stage coefficient, recession	1.01	1.06	1.28
First-stage coefficient, expansion	.66	1.28	.98
First-stage <i>F</i> -statistic, recession	16.7	20.2	13.3
First-stage <i>F</i> -statistic, expansion	8.5	66.7	15.4
First-stage observations	1,091	1,091	1,091
Second-stage observations	1,305	1,305	1,305

NOTE.—This table reports TS2SLS regressions. The dependent variable is the change in the unemployment rate between the national recession peak and trough (recession-to-recovery episode) or during the first 30 months of the national expansion. Additional controls in col. 2 are lags of employment growth, population growth, and house price growth, each measured from 5 years to 1 year before the cycle start; area size, measured by the log of sample mean employment; and the Herfindahl of industry concentration at the cycle start. All controls and fixed effects are interacted with an indicator variable for the recession-to-recovery period. Standard errors are given in parentheses and clustered by CSA/MSA.

** Significant at the 1% level.

of reallocation per year over the course of a cycle causes unemployment to rise by $0.44 \times 2.5 = 1.1$ percentage points during the first 30 months of the cycle. Column 2 adds the following control variables, described in detail in appendix C: lags of employment growth, population growth, and house price growth, each measured from 5 years to 1 year before the cycle start; area size, measured by the log of sample mean employment; and the Herfindahl of industry concentration at the cycle start.¹⁸ Column 3 adds area fixed effects. 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. In columns 2 and 3, the added controls produce stronger first-stage fits and smaller second-stage standard errors. Importantly, the point estimate for the effect of reallocation remains stable across columns.

Table 5 presents the main empirical result of the paper. The specifications mirror those in table 4 except that we allow the effects of reallocation

¹⁸ We further describe and report the coefficients and standard errors on the control variables in app. C.

(and any covariates) to vary according to whether the reallocation occurs during a recession-to-recovery or an expansion cycle. Accordingly, the first-stage regression includes two excluded instruments, one for predicted reallocation during a recession-to-recovery cycle and the other for predicted reallocation during an expansion. All coefficients are exactly equal to those from estimating regressions during recession-to-recovery and expansion phases separately. Reallocation during a recession-to-recovery cycle increases unemployment by an economically large and statistically significant amount. Across columns, the data reject zero effect of reallocation during a recession-to-recovery cycle at the 1% level, with t -statistics ranging from 3.5 to 4.3. In economic magnitude, a 1 standard deviation increase in predicted reallocation causes an unemployment rate 0.5 percentage points higher ($2.5 \text{ years} \times 0.87 \text{ percentage points per year} \times 0.23 \text{ standard deviation}$) by the time national employment reaches its trough. In contrast, the data do not reject zero effect of reallocation on unemployment during expansions, with the point estimates slightly negative. The data therefore strongly reject equality of coefficients during recession-to-recovery and expansion cycles.

Figure 3 shows the full timing of the effects of reallocation on unemployment during a recession-to-recovery cycle. The solid line in figure 3 plots the coefficients β_j from a local projection of the change in the unemployment rate on reallocation—that is, the coefficients β_j from estimating equations (8) and (9), allowing j to vary in 6-month increments. The coefficients trace out a hump-shaped impulse response function. Areas undergoing reallocation during a national recession-to-recovery cycle experience a relative rise in unemployment while national employment is falling. The effect crests at the national employment trough and reverses




FIG. 3.—Timing of effects on unemployment during a recession-to-recovery cycle. The solid line plots the coefficients on reallocation from a regression of the change in the unemployment rate at different horizons on reallocation. The dotted lines plot the 95% confidence interval for each horizon. A color version of this figure is available online.

as the economy recovers. The coefficients for 1 and 2 years before the national peak indicate little evidence of areas with large predicted reallocation during the recession-to-recovery cycle experiencing differential unemployment rate trends immediately before the national peak. The absence of pretrends means that even if reallocation does not come as a surprise shock at the start of a recession, anticipation effects do not contaminate our estimates of what happens during the recession.

C. *Robustness*

Table 6 groups together a number of sensitivity exercises to assess the robustness of the finding that reallocation affects unemployment during recessions. Each row of the table reports for a separate regression the coefficient and standard error for recession-to-recovery periods and the *p*-value of equality of coefficients in a recession-to-recovery cycle and an expansion. The first row, labeled “Baseline,” reproduces column 1 of table 5.

Rows 2–5 further expand the control variables in the regression. Row 2 adds nonparametric controls for the Bartik predicted growth rate over the cycle by adding episode-specific indicator variables for belonging to each of 20 quantiles of predicted growth. This specification compares the evolution of unemployment across areas with different predicted reallocation but in the same vigintile of predicted growth. Row 3 controls for the share of the population in 5-year age bins at the start of the cycle. Rows 4 and 5 allow the coefficients on predicted growth to vary by cycle and by region, respectively. The effects of reallocation during recession-to-recovery cycles increase slightly in each of these specifications, and with the exception of row 2 the difference in coefficients between recession-to-recovery cycles and expansions remains significant at the 5% level.¹⁹

Row 6 removes from the sample areas that contain an industry with employment relative to national employment in that industry above 5%. If employment in an industry concentrates in a few areas (e.g., auto manufacturing employment in Detroit and Birmingham) and if the firms in different areas engage in strategic interaction (e.g., negative shocks to plants in Detroit induce expansion at plants in Birmingham), then local industry employment in an area may correlate with employment elsewhere 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. In practice, the small size of CSA/MSAs relative

¹⁹ The specifications in rows 2, 9, and 15 of the table cannot reject equality of coefficients at the 10% level. In each of these cases, the larger *p*-value reflects a large increase in the standard error for the expansion coefficient rather than convergence of the point estimates, the increase in the standard error in turn reflects a weak first stage for the expansion subsample, and (not shown) including the baseline controls strengthens the first stage and yields a rejection of equality at least at the 5% level.

TABLE 6
ROBUSTNESS

SPECIFICATION	RECESSION-TO-RECOVERY CYCLE			CSAs	OBSERVATIONS
	β	Standard Error	p (Equality)		
1. Baseline	.87**	.25	.004	220	1,305
2. Bartik growth vigintiles	.98**	.30	.404	220	1,305
3. Control for age shares	1.10**	.23	.000	220	1,305
4. Predicted growth by period	.94**	.26	.008	220	1,305
5. Predicted growth by region	.99**	.30	.043	220	1,305
6. Drop if area has large industry	.86**	.29	.002	203	1,107
7. Drop if high manufacturing share	1.25**	.45	.007	193	960
8. Drop if high construction share	1.11**	.27	.000	206	966
9. Drop if high resources share	1.04**	.35	.160	177	973
10. Drop if high health care share	.70*	.33	.031	198	987
11. Drop persistent industries	1.47**	.37	.004	220	1,305
12. Drop 1990–93	.86**	.31	.008	220	1,091
13. Drop 2008–14	1.10**	.32	.002	220	1,101
14. HP filter dating	.88**	.28	.040	220	1,291
15. NAIRU dating	1.15*	.58	.399	221	1,246
16. Expand window $\pm$ 3 months	.71*	.28	.067	218	1,125
17. Peak-to-peak reallocation	.82*	.41	.030	220	1,305

NOTE.—Each row of this table reports the coefficient and standard error for recession-to-recovery periods and the p -value of equality of coefficients in a recession-to-recovery cycle and an expansion from a separate regression described in col. 1. All rows additionally include the controls from col. 1 of table 5. The first row, labeled “Baseline,” reproduces col. 1 of table 5. Row 2 controls for episode-specific indicator variables for belonging to each of 20 quantiles of predicted growth. Row 3 controls for the share of the population in 5-year age bins at the start of the cycle. Rows 4 and 5 allow the coefficients on predicted growth to vary by cycle and by region, respectively. Row 6 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. Rows 7–10 exclude observations in the cycle’s top quartile of employment share in the industry indicated. Row 11 excludes from the construct of predicted reallocation any industry that either expands or contracts nationally in every cycle in our data. Row 12 drops the 1990–93 cycle. Row 13 drops the 2008–14 cycle. In row 14, a recovery ends when the cyclical component of an HP filter of national employment (smoothing parameter: 129,600) equals zero. In row 15, a recovery ends when the national unemployment rate first equals the CBO’s estimate of the NAIRU. In row 16, the recession-to-recovery window is extended by 3 months on each side. Row 17 constructs predicted reallocation on a peak-to-peak basis.

* Significant at the 5% level.

** Significant at the 1% level.

to the national economy—the median CSA/MSA has total employment in 2013 of 113,000, or less than 0.1% of national private-sector employment—makes such concentration scarce. Indeed, the restriction removes 17 CSA/MSAs from the sample (primarily the largest ones) but has almost no effect on the estimated coefficients.

Rows 7–10 assess sensitivity to excluding areas with large employment shares in manufacturing, construction, resource extraction, and health

care. Goldsmith-Pinkham, Sorkin, and Swift (2018) show the equivalence of a shift-share IV estimator to a Rotemberg-weighted sum of IV coefficients using area shares in each individual industry and period as the excluded instrument in separate just-identified regressions. The Rotemberg weights characterize the influence of each industry on the IV coefficient. To assess the sensitivity to reducing the influence of high-weight industries, we first compute the Rotemberg weights for a close variant of our baseline regression that has the shift-share structure.²⁰ Of the five largest (in absolute value) weights for a recession-to-recovery period, two correspond to manufacturing industries, two to natural-resource extraction, and one to construction. We then remove from the sample areas in the top quartile of beginning-of-cycle employment share in each of these sectors. Excluding these areas substantially reduces the Rotemberg weight associated with the particular industry and directly addresses the concern raised by Goldsmith-Pinkham, Sorkin, and Swift (2018) of industries with high Rotemberg weights concentrating in areas experiencing other shocks. The coefficients from these restricted samples remain close to the baseline.

Row 11 follows Notowidigdo (2011) and keeps all areas but removes from the construction of predicted reallocation any industry that either increases or decreases national employment in each cycle in our data.²¹ Thus, the variation in row 11 comes only from industries experiencing persistent but not permanent expansions or contractions. Row 12 excludes the 1990–93 cycle from the second sample and reports the coefficient

²⁰ The variant contains only the national growth rate inside the absolute value in eq. (5). The 2SLS recession-to-recovery coefficient in this specification is 1.05 (SE = 0.30), close to that of our baseline coefficient. The weight formula for industry i in period t is $(\hat{R}^b R^\perp)^{-1} |g_{i,t}| Z_{i,t} R^\perp$, where $\hat{R}^b$ is the $AP \times 1$ vector of Bartik reallocation in each area $a \in A$ and period $t = 1, 2, \dots, P$; $R^\perp$ is the $AP \times 1$ vector of actual reallocations orthogonalized with respect to covariates; $g_{i,t}$ is the growth rate of national employment in industry i in period t ; and $Z_{i,t}$ is an $AP \times 1$ vector consisting of zeros in all rows not corresponding to period t and the location-industry initial employment shares in industry i for the rows corresponding to period t . Because construction of the Rotemberg weights requires computing the covariance of shares and the endogenous variable, we exclude the 1990–93 cycle from this exercise. Appendix table C.3 reports additional statistics related to the Rotemberg weights. We can alternatively motivate the robustness exercises in rows 7–11 as reducing the influence of industries that made the largest contributions to national reallocation over the past decades, consistent with the identification approach in Borusyak, Hull, and Jaravel (2018). In this vein, row 10 shows the robustness to excluding areas with large shares in health care.

²¹ More specifically, we exclude an NAICS (SIC) industry if national employment expands or contracts during each cycle in which NAICS (SIC) employment is reported. This procedure identifies 47 of the 165 SIC two-digit or NAICS three-digit industries in our data: 113, 114, 313–316, 322, 323, 325 (NAICS, shrinking); 112, 485, 488, 493, 541, 562, 611, 621–624, 712, 722 (NAICS, expanding); 09, 11, 12, 21–23, 31, 66 (SIC, shrinking); 02, 07, 08, 47, 58, 62, 64, 67, 72, 73, 75, 79–84 (SIC, expanding). These industries overlap substantially with the set of industries with high Rotemberg weights, with five of the 10 largest (in absolute value) Rotemberg weights in persistently shrinking or expanding industries.

from a conventional 2SLS regression. Row 13 shows that excluding the Great Recession has a small effect on the results.

Rows 14–17 explore robustness to the precise timing definition. Row 14 defines the end of the recovery as the first month following a peak in which the cyclical component of a Hodrick-Prescott (HP) filter (smoothing parameter: 129,600) of national private-sector employment turns positive. In row 15, we define the end of the recovery as the first month in which the national unemployment rate falls to or below the Congressional Budget Office (CBO) estimate of the nonaccelerating inflation rate of unemployment (NAIRU). Row 16 expands the recession-to-recovery window symmetrically by 3 months on each side. Finally, row 17 redefines the reallocation timing to measure reallocation between two national employment peaks.²² The basic pattern of a statistically significant positive coefficient during recession-to-recovery cycles and a smaller coefficient during expansions remains robust to these alterations.

V. Quantitative Model

The previous section demonstrated that reallocation causes an increase in unemployment if it occurs during a national recession-to-recovery cycle but not if it occurs during an expansion. The rise in the unemployment rate concentrates during the recession part of the cycle, with maximum impact around the national employment trough. We now study a model economy to better understand these patterns. The model illustrates what types of primitive shocks could give rise to our empirical setup and how frictions to labor mobility and downward wage rigidity allow the model to match the empirical patterns in the data.

A. Setup

Time is discrete. The economy consists of $\mathcal{A}$ islands, each of which has $\mathcal{I}$ industries. We assume no aggregate uncertainty and perfect consumption insurance within (but not across) islands, which implies an island-specific discount factor $m_{a,t,t+1}$.

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

²² For the 2008–14 cycle, we measure reallocation between the peak in January 2008 and December 2016.

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, Eichenbaum, and Trabandt (2016), 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 $n_{a,i,t} = Mv_{a,i,t}^{1-\alpha} x_{a,i,t}^\alpha$. Letting $\theta_{a,i,t} = v_{a,i,t}/x_{a,i,t}$ 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 $q_{a,i,t} = M\theta_{a,i,t}^{-\alpha}$. Thus, the value of a filled job to the firm, $J_{a,i,t}$, and the free-entry condition are

$$J_{a,i,t} = (p_{a,i,t} - w_{a,i,t}) + (1 - \delta_t)m_{a,t,t+1}J_{a,i,t+1}, \quad (10)$$

$$\kappa = q_{a,i,t}J_{a,i,t}, \quad (11)$$

where $p_{a,i,t}$ represents the real marginal product and $w_{a,i,t}$ represents the real wage.

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 semidirected search (Kline 2008; Artuç, Chaudhuri, and McLaren 2010; Kennan and Walker 2011; Dvorkin 2014; Pilossoph 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 $\Lambda_{a,i,t}$ at exogenous rate λ_t . An industry reallocation shock $\Lambda_{a,i,t}$ consists of an immediate job separation if previously employed, $\mathcal{I}$ time-invariant sector-specific taste parameters $\{\psi_{a,j}\}_{j=1}^{\mathcal{I}}$, and a draw of $\mathcal{I}$ idiosyncratic taste shocks $\{\varepsilon_{j,t}\}_{j=1}^{\mathcal{I}}$ from a distribution $F(\varepsilon)$. These taste parameters and shocks enter additively into the worker's value function for searching in each sector $j = 1, \dots, \mathcal{I}$. The value functions of an employed worker, $W_{a,i,t}$, and an unemployed worker, $U_{a,i,t}$, are then

$$\begin{aligned}
W_{a,i,t} = & w_{a,i,t} + m_{a,t,t+1} \{ [(1 - \delta_t) + (\delta_t - \lambda_t) f_{a,i,t+1}] W_{a,i,t+1} + (\delta_t - \lambda_t) (1 - f_{a,i,t+1}) U_{a,i,t+1} \} \\
& + m_{a,t,t+1} \lambda_{a,t} \left(\mathbf{E} \max_j \{ (1 - f_{a,j,t+1}) U_{a,j,t+1} + f_{a,j,t+1} W_{a,j,t+1} + \psi_{a,j} + \varepsilon_{j,t} \} \right),
\end{aligned} \tag{12}$$

$$\begin{aligned}
U_{a,i,t} = & z + m_{a,t,t+1} \{ (1 - \lambda_t) [f_{a,i,t+1} W_{a,i,t+1} + (1 - f_{a,i,t+1}) U_{a,i,t+1}] \} \\
& + m_{a,t,t+1} \lambda_{a,t} \left(\mathbf{E} \max_j \{ (1 - f_{a,j,t+1}) U_{a,j,t+1} + f_{a,j,t+1} W_{a,j,t+1} + \psi_{a,j} + \varepsilon_{j,t} \} \right),
\end{aligned} \tag{13}$$

where z represents the worker's flow opportunity cost of employment and $\mathbf{E}$ represents the expectations operator.

The parameter λ determines the share of unemployed reoptimizing their industry search market. Equivalently, λ has the interpretation of a stochastic death or retirement shock, with a new generation of workers of mass λ born each period and choosing afresh their industry of search. Holding the share reoptimizing below unity provides one important friction allowing reallocation shocks to affect employment. The $\psi_{a,j}$ parameters can be interpreted as permanent preferences to work in particular sectors. The ε_j shocks have the interpretation of transitory taste shocks that make some individuals prefer to work in certain sectors or of noise shocks that 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. Thus, consistent with Pilossoph (2014), net reallocation in our model occurs without requiring changes in the amount of gross flows. The level of λ and the volatility of the process generating ε_j together govern the magnitude of gross flows and the directness of search across industries.²³

²³ The assumption of time-dependent, stochastic reallocation shocks $\Lambda_{a,i,t}$ rather than a state-dependent reallocation decision and a fixed cost of moving makes the quantity of gross flows exogenous. In an aggregate steady state, the two approaches are isomorphic. The assumption of time-dependent shocks is more computationally tractable for a large number of industries. Quantitatively, the volatility of the preference shocks matters more for our results than the level of gross flows. Further, since we study the response to very long-lasting industry dispersion shocks, we do not think that allowing for "rest" unemployment as in Alvarez and Shimer (2011) would meaningfully affect the model's conclusions. We also abstract from geographical mobility. In app. sec. D.10, we extend the model to incorporate area reallocation. This extension yields quantitatively greater employment responses, due to the migration channel, but very similar effects of reallocation on local area unemployment.

We denote the transition probability from industry i to industry j conditional on an industry reallocation shock by $\pi_{a,ij,t}$. This probability does not depend on the worker's previous employment status or industry, $\pi_{a,ij,t} = \pi_{a,hj,t} = \pi_{a,j,t}$. We have three laws of motion for the evolution of job seekers, employment, and unemployment:

$$\begin{aligned} x_{a,i,t} &= \delta_{t-1} e_{a,i,t-1} + u_{a,i,t-1} - \lambda_{t-1} l_{a,i,t-1} + \pi_{a,i,t-1} \lambda_{t-1} l_{a,t-1}, \\ e_{a,i,t} &= (1 - \delta_{t-1}) e_{a,i,t-1} + f_{a,i,t} x_{a,i,t}, \\ u_{a,i,t} &= (1 - f_{a,i,t}) x_{a,i,t}. \end{aligned}$$

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

$$w_{a,i,t} = \max\{w_{a,i,t}^*, (1 - \chi^w) w_{a,i,t-1} / \Pi_{a,t}\}, \quad (14)$$

where $w_{a,i,t}^*$ represents the Nash bargained real wage, $\Pi_{a,t}$ represents gross producer price inflation, and χ^w is a parameter specifying the maximum-permitted decline in the nominal wage. Following Hall (2005) and Chodorow-Reich and Karabarbounis (2016), exogenous wage rigidity allows the model to generate realistic unemployment fluctuations without violating bilateral efficiency conditions or requiring counterfactual assumptions on the sources of wage rigidity. A large literature reports evidence of downward nominal wage rigidity in the data (Card and Hyslop 1997; Kahn 1997; Dickens et al. 2007; Goette, Sunde, and Bauer 2007; Daly and Hobijn 2014).²⁴

2. General Equilibrium

Output of industry i in area a is

$$Q_{a,i,t} = \eta_{i,t} e_{a,i,t}, \quad (15)$$

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

$$Q_{a,t} = \left[\sum_i \tau_{a,i,t}^{1/\xi} Q_{a,i,t}^{(\xi-1)/\xi} \right]^{\xi/(\xi-1)}, \quad (16)$$

²⁴ Still, this assumption is not without controversy. Pissarides (2009) shows in the context of a search model with exogenous separations that what matters for unemployment fluctuations is the wage rigidity of new hires. Daly and Hobijn (2014) and Gertler, Huckfeldt, and Trigari (2015) provide evidence of rigidity on this margin, including of downward wage rigidity.

giving rise to a downward-sloping industry-level demand curve $Q_{a,i,t} = \tau_{a,i,t} (P_{a,i,t}^Q / P_{a,t}^Q)^{-\zeta} Q_{a,t}$ and where $\zeta \geq 1$ and $P_{a,t}^Q = [\sum_i \tau_{a,i,t} (P_{a,i,t}^Q)^{1-\zeta}]^{1/(1-\zeta)}$. 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 (10) is the product of industry productivity and the real price of industry i 's good:

$$p_{a,i,t} = \eta_{i,t} P_{a,i,t}^Q. \quad (17)$$

With downward-sloping demand, the decline in output engendered by a decline in $\eta_{i,t}$ induces a rise in the real price $P_{a,i,t}^Q$, 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 $P_{a,i,t}^Q$, overall inflation, and the discount factor $m_{a,t,t+1}$. We assume that product prices are determined competitively. An Euler equation for each household determines consumption and the discount factor. 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 D.

B. Calibration

We calibrate a version of the model with two areas and 10 industries, $\mathcal{A} = 2$ and $\mathcal{I} = 10$, at monthly frequency. The two areas allow for one small (infinitesimal) area, which we treat as representative of a single local CSA/MSA, and one large area representative of the rest of the economy. Our choice of 10 industries represents a balance between computational feasibility and ensuring that national industry trends are representative of local industry trends.²⁵

We briefly describe the calibration of the labor market block of the model, shown in table 7. The parameters are the same in the small and in the large area. Appendix D contains further details and our procedure for finding the model steady state. We obtain a target for the steady-state job-finding rate f appropriate to a two-state labor market model of 0.5 by

²⁵ Recall our first-stage coefficient of roughly one in table 3. With $\mathcal{I} = 2$ and one location having a near-100% employment share in the expanding sector, clearly the employment share could not grow at the same rate locally as it does nationally. We found that $\mathcal{I} \geq 10$ largely eliminates such cases.

TABLE 7
CALIBRATED LABOR MARKET PARAMETERS

Parameter	Description	Value	Source
f	Job-finding rate	.5	Shimer 2012
q	Job-filling rate	.75	Davis, Faberman, and Haltiwanger 2013
δ	Separation rate	.066	Matched monthly CPS
λ	Industry reallocation	.043	Matched monthly CPS
ρ	Industry-reallocation noise	.95	Peak unemployment effect in recessions
D	Discount factor	.997	4% annual rate
β	Bargaining power	.6	
χ^w	Downward wage rigidity	.0035	Average monthly nominal wage growth
z	Opportunity cost	.55 p	Chodorow-Reich and Karabarbounis 2016

updating the procedure described in Shimer (2012) and for the job-filling rate q of 0.75 from Davis, Faberman, and Haltiwanger (2013). Using the longitudinally linked monthly Current Population Survey (CPS) files, we find a monthly job-separation rate of 0.066. This separation rate exceeds that implied by the procedure in Shimer (2012) because it includes job-to-job transitions and therefore is more appropriate for our labor market setting. We also use the longitudinally linked monthly CPS files to calculate that 60% of employment-to-unemployment-to-employment spells end with the worker employed in a different three-digit NAICS industry. This fraction together with the job-finding and job-separation rates determines the reallocation intensity λ . Given that our model has neither aggregate productivity growth nor trend inflation, we set the downward wage rigidity parameter χ^w to the 0.35% average monthly increase in nominal hourly earnings of production and nonsupervisory employees. This value allows nominal wages to fall by 0.35% each month relative to trend, corresponding to zero nominal wage growth. We parameterize the idiosyncratic shocks $F(\varepsilon)$ as type I extreme value $(-\rho\tilde{\gamma}, \rho)$, where $\tilde{\gamma}$ is Euler’s constant. The parameter ρ governs the directness of search and does not have an easily observed counterpart—it determines the ratio of net reallocation to gross reallocation given an increase in dispersion of industry labor demand—and the existing literature provides a wide range of estimates. We therefore infer ρ together with the shocks, as we describe in the following section. This yields $\rho = 0.95$. We parameterize the fixed reallocation utilities $\{\psi_{aj}\}_{j=1}^J$, such that the steady-state allocation of labor is efficient.

C. Quantitative Exercise

We conduct a just-identified indirect inference to recover ρ , $\{(\eta_{i,t+T})/\eta_{i,t}\}$, and μ_ν , such that the model matches the following empirical targets over a




FIG. 4.—Productivity paths and local predicted growth and reallocation. *A*, Perfect-foresight productivity paths η_i of each of the $i = 1, \dots, 10$ sectors. The productivity paths are mean preserving in logs. *B*, Cross section of predicted growth and predicted reallocation in our simulated local areas over a recession-to-recovery cycle. A color version of this figure is available online.

full recession-to-recovery cycle: (1) the average employment share changes,²⁶ (2) the average unemployment increase and duration of the cycle, and (3) the peak cross-sectional effect of reallocation on unemployment (fig. 3).

We restrict the $\mathcal{I} = 10$ productivity paths to be log linear and mean preserving, so that they can be (heuristically) identified from 10 employment share changes. Figure 4A plots the productivity paths with industries in decreasing order of productivity change. The dispersion begins in period -24 and reaches full spread after 135 months. Period 0 will correspond to the national employment peak before the recession. The remaining duration (111 months) corresponds to the average duration of a full peak-to-peak cycle. We solve the model under perfect foresight, so the productivity paths are known as of $t = -24$. Thus, our calibration features predictable industry-specific trends before, during, and after the recession.

We create a demand-driven recession with a temporary increase in the wedge μ_t between the central bank's policy interest rate and the interest rate available to households. To avoid conflating news of the recession with news of reallocation, agents learn of the wedge shock 25 months after the reallocation shock, at $t = 1$. The wedge equals $\mu_t = 0.0057$ for $1 \leq t \leq 36$, where its size and duration are heuristically identified by the average unemployment rate increase (3.5 percentage points) and recession-to-recovery duration (53 months). The national unemployment peak occurs after 24 months, close to the 30 months in the data.

In appendix section D.9, we verify that industry employment growth in the large area over a full recession-to-recovery cycle is highly correlated with the idiosyncratic industry shocks. We also show that by the end of

²⁶ We sort the 93 NAICS 3 sectors into 10 quantiles on the basis of their employment share changes and then calculate the average employment share change for each quantile.

the recession-to-recovery cycle, the industry employment distribution has reached the new steady-state level implied by the idiosyncratic productivity shocks. Thus, our timing assumption ensures that the Bartik instrument reflects the desired amount of reallocation and not temporary frictions to employment or mobility.

We simulate the model 200 times for different initial employment shares in the small area.²⁷ We construct the initial employment shares from 10 independent draws from a uniform distribution, $U[0.06, 1]$, scaled to sum to one. The positive lower bound on the distribution rules out zero sectoral employment shares as well as employment shares above 65%. Figure 4*B* plots a cross section of the implied predicted growth rates and the predicted reallocation rates for each local area.

The above describes the shocks during the recession-to-recovery scenario. In the expansion scenario, we turn off the interest rate wedge shock. We solve the model by reversing the shocks after 250 periods and working backward from the initial steady-state equilibrium.²⁸

Figure 5*A* plots the residualized change in local unemployment from $t = 1$ to $t = 24$ (the national unemployment peak) against the residualized instrumented reallocation for the local area. Both variables are residualized with respect to predicted growth over the respective cycle and a constant, which corresponds to the second stage of our empirical estimates. The figure highlights a strong positive relationship between (residualized) unemployment and (residualized) instrumented reallocation in a recession, whereas there is no economically significant relationship in an expansion.

Figure 5*B* plots the marginal effect of reallocation on unemployment for the expansion and recession-to-recovery scenarios. These are the IV coefficients from running regressions of the change in unemployment on realized reallocation instrumented with predicted reallocation, controlling for predicted growth and a constant. The marginal effects for $t = 24$ (the national unemployment peak) correspond to the best-fit lines through the points in figure 5*A*. These marginal effects of reallocation on unemployment have their data analogue in table 5 and the plot in figure 3. In the model, as in the data, the marginal effect of reallocation on unemployment is significantly larger and more persistent in recession-to-recovery cycles than in expansions. The maximum marginal effect during the recession-to-recovery peaks at 2.73 after 34 months, consistent with the empirical impulse response function in figure 3. In contrast, the maximum marginal effect during the expansion is small (0.02), falling within the range of point estimates reported in table 5 (after multiplying the latter by 2.5). While the model fits the peak effect in recessions by

²⁷ Since the small area is infinitesimal, the solution for the large area is identical in each simulation.

²⁸ Allowing for a longer amount of time before reversal has little effect on our results.




FIG. 5.—Model impulse response function (IRF) and marginal effect. *A*, Cross section of the residualized increase in unemployment from period 0 to period 24 against the residualized instrumented reallocation, separately for the recession-to-recovery scenario and the expansion scenario. *B*, Marginal effects of realized reallocation based on an IV regression at different horizons. A color version of this figure is available online.

construction, the model is corroborated by the dynamics of the marginal effects and the asymmetry of the expansion results.

D. Mechanism

The model replicates the finding that reallocation generates a quantitatively large increase in local unemployment if it occurs during a recession but not if it occurs during an expansion. Compression of the wage distribution during recessions provides a key mechanism underlying this difference.²⁹ 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 unmarked lines in figure 6*A*. Due to imperfect labor mobility, tightness diverges as shown in figure 6*B*, causing a decrease in 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, as illustrated by the dashed lines with triangles in figure 6*A*. The wage constraint

²⁹ Our focus on wage rigidity does not mean that we reject other possible mechanisms. Recent work has 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., Ahn and Hamilton 2014; Hall and Schulhofer-Wohl 2015). Dvorkin (2014) and Pilossoph (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.




FIG. 6.—Model impulse response functions of real wages and labor market tightness. A, Impulse response functions of real wages relative to the initial steady state in the large area of the currency union. B, Corresponding impulse response functions for labor market tightness. For ease of exposition, we plot results only for three sectors. A color version of this figure is available online.

binds most tightly in the industries experiencing both secular decline and the cyclical recession, further compressing job values and disincentivizing vacancy creation, while rising wages in the expanding industry limit the increase in job values and therefore vacancy posting.³⁰ The decline in total vacancies causes unemployment to increase.³¹ In addition, a constant-returns-to-scale matching function means that the sharp divergence in labor market tightness and hence job-finding rates across sectors generates additional mismatch unemployment beyond that accounted for by the decline in aggregate vacancies (Sahin et al. 2014; Barnichon and Figura 2015).³²

In table E.1 (tables C.1–C.3, D.1, and E.1 are available online), we test for the asymmetry of wage compression during recessions and expansions using national hourly wages by industry from the CPS and QCEW

³⁰ Throughout these dynamics, the job surplus in each sector remains positive, so that the downward nominal wage constraints respect bilateral efficiency. This follows from our calibration of the worker's opportunity cost such that the steady-state surplus of a job is relatively large.

³¹ The decline in vacancies in our model contradicts the argument in Abraham and Katz (1986) that the decline in vacancies during recessions necessarily limits the potential role for reallocation shocks. Their logic rests on a downward-sloping Beveridge curve (a negative relationship between unemployment and vacancies) absent sectoral shifts, while reallocation shocks that shift the Beveridge curve induce positive comovement of vacancies and unemployment. Hosios (1994) first challenged the robustness of this logic.

³² Constant returns to scale make job-finding rates concave in tightness so that minimizing unemployment holding the number of vacancies fixed requires equalizing job-finding rates across sectors. The magnitude of mismatch unemployment in our model calibrated to features of the Great Recession (employment share changes, depth, duration, and inflation rate) is 1.15 percentage points, within the range reported in table 1 of Sahin et al. (2014).




FIG. 7.—Marginal effect of reallocation in alternative model specifications. *A*, Marginal effect of reallocation on unemployment in the baseline model and when period-by-period Nash bargaining replaces the downward nominal wage rigidity assumption. *B*, Contrast of the baseline model with a parameterization that sets the area and industry reallocation noise parameters to 20% of their baseline values, which weakens reallocation frictions and makes labor search more directed. In each case, we parameterize the reallocation and recession shocks to hit the same targets as in our baseline. A color version of this figure is available online.

employment share changes. Consistent with the model's mechanism, we find that industries with rising employment shares have rising wage differentials during expansions but that there is no economically or statistically significant relationship between the change in the wage premium during a recession and industry share growth.

To quantitatively illustrate the model's key frictions, figure 7A plots the baseline marginal coefficient impulse response function (darker-shaded solid line) along with counterfactual marginal coefficients obtained by selectively changing features of the model. The lighter-shaded dash-dotted line shows the marginal effects in recession-to-recovery cycles with wages determined by Nash bargaining in each period. Since this experiment does not feature any nominal rigidity, we induce a recession with a negative productivity shock, calibrated to hit the same recession targets as our baseline. Without downward wage rigidity, the effects are much smaller and of similar magnitude across business cycle states.³³

Figure 7B illustrates the importance of reallocation frictions. The dashed lines report the impulse response functions of the marginal coefficients when the variance of the reallocation taste shocks ε_j is reduced to 20% of its baseline values. This change makes the industry and location

³³ We kept the baseline ρ in the experiment with Nash bargaining, since raising ρ to hit the peak marginal unemployment effect resulted in much wider dispersion in the productivity paths, causing negative job surplus in the contracting sector. Figure D.3 displays the comparison with a higher $\rho = 4$, such that the surplus remains positive in all sectors. In this case, the marginal effects of reallocation on unemployment under Nash bargaining are too large in expansions, are too small in recessions, and counterfactually increase with the horizon.

choice much more directed. A greater directedness of search per se has only a minor effect on the employment response to a given set of productivity processes. However, it implies that the same employment share changes across sectors require much less dispersion in productivity. This limits the dispersion of labor market tightness and the extent to which wage constraints bind asymmetrically across sectors, which results in marginal coefficients of nearly zero in both expansions and recessions. Conversely, figure D.4 shows that an increase in labor market fluidity (Davis and Haltiwanger 2014) mitigates the adverse employment consequences of reallocation shocks.

VI. Conclusion

Reallocation of workers across industries matters asymmetrically over the business cycle. To establish this fact, we develop a methodology to isolate the secular component of industry dispersion. 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. Areas subjected to more reallocation have worse local recessions if national employment is depressed but similar outcomes 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 that generate sticky wages, reallocation does not much affect labor market 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 magnitudes similar to those found in our empirical exercises. Third, compression of wage differentials during recessions can explain the asymmetric response of aggregate employment to reallocation during recession-to-recovery cycles and expansions.

Our analysis has implications for policy not explored in this paper. The idiosyncratic industry shocks that underlie secular industry reallocation include real shocks, such as dispersion in productivity levels or consumer taste trends. Nonetheless, the ease with which reallocation occurs sharply depends 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.

Appendix A

Standard Errors in TS2SLS

This appendix derives the TS2SLS asymptotic standard error when the first- and second-stage residuals are not necessarily independent. The notation completely departs from the notation in the main text to (roughly) follow Murphy and Topel (1985).

Notation.—We assume two possibly interdependent samples of size N_1 and N_2 drawn from the same population, where the subscripts 1 and 2 define the samples. We are interested in the following two-stage model of the observed data:

$$\text{Structural equation: } Y_j = F_j\gamma + X_j\beta + \eta_j, \quad j = 2, \quad (\text{A1})$$

$$\text{First stage: } F_j = W_j\pi_w + X_j\pi_x + e_j, \quad j = 1, \quad (\text{A2})$$

$$\text{Second stage: } Y_j = H_j\pi\gamma + X_j\beta + u_j, \quad j = 2, \quad (\text{A3})$$

where $Y_j \sim N_j \times 1$ denotes a vector of observations of the outcome variable in sample j ; $X_j \sim N_j \times K$ is a matrix of included instruments; $F_j \sim N_j \times M$ is a matrix of endogenous variables; $W_j \sim N_j \times L$ is a matrix of excluded instruments, where $L \geq M$; $H_j = (X_j \ W_j) \sim N_j \times (K + L)$ is the matrix of first-stage regressors; $\eta_j \sim N_j \times 1$ is the vector of structural residuals; $e_j \sim N_j \times M$ is the matrix of first-stage residuals; $u_j \sim N_j \times 1$ is the vector of second-stage residuals; β , γ , π_x , and π_w are conforming matrices of parameters; and $\pi = (\pi_x \ \pi_w)' \sim (K + L) \times M$. Let $\hat{\pi} = (H_1'H_1)^{-1}H_1'F_1$ denote the fitted coefficients from the first-stage regression, and let $Z_2 = (X_2 \ H_2\hat{\pi}) \sim N_2 \times (K + M)$ denote the matrix of regressors in the second stage.

Assumptions.

$$E[Z_j'Z_j] = Q \sim (K + M) \times (K + M), \quad (\text{A4})$$

$$E[H_j'H_j] = R \sim (K + L) \times (K + L), \quad (\text{A5})$$

$$E[Z_j'H_j] = S \sim (K + M) \times (K + L), \quad (\text{A6})$$

$$E[Z_j'u_ju_j'Z_j] = \Omega_{zu} \sim (K + M) \times (K + M), \quad (\text{A7})$$

$$E[H_j'e_j\gamma\gamma'e_j'H_j] = \Omega_{H(e\gamma)} \sim (K + L) \times (K + L), \quad (\text{A8})$$

$$E[Z_j' e_j \gamma \gamma' e_j' Z_j] = \Omega_{Z(\epsilon_j)} \sim (K + M) \times (K + M), \quad (\text{A9})$$

$$E[Z_j' u_j \gamma' e_j H_j] = \Omega_{Zu(\epsilon_j)H} \sim (K + M) \times (K + M), \quad (\text{A10})$$

$$E[Z_j' u_j \gamma' e_j Z_j] = \Omega_{Zu(\epsilon_j)Z} \sim (K + M) \times (K + M), \quad (\text{A11})$$

$$\lim_{N_1, N_2 \rightarrow \infty} N_2/N_1 = \alpha > 0. \quad (\text{A12})$$

Note that in equations (A4)–(A11), the moment conditions hold for any sample, $j = 1, 2$.

Derivation.—Rewriting equation (A3) as $Y_2 = X_2\beta + H_2\hat{\pi}\gamma - H_2(\hat{\pi} - \pi)\gamma + u_2$, one can derive an expression for the sampling error in the least squares estimator $(\hat{\beta} \ \hat{\gamma})'$:

$$\begin{aligned} \sqrt{N_2} \begin{pmatrix} \hat{\beta} - \beta \\ \hat{\gamma} - \gamma \end{pmatrix} &= \left(\frac{1}{N_2} Z_2' Z_2 \right)^{-1} \left(\frac{1}{\sqrt{N_2}} Z_2' u_2 \right) - \left(\frac{1}{N_2} Z_2' Z_2 \right)^{-1} \frac{1}{\sqrt{N_2}} Z_2' H_2 (\hat{\pi} - \pi) \gamma \\ &= \left(\frac{1}{N_2} Z_2' Z_2 \right)^{-1} \left(\frac{1}{\sqrt{N_2}} Z_2' u_2 \right) - \sqrt{\frac{N_2}{N_1}} \left(\frac{1}{N_2} Z_2' Z_2 \right)^{-1} \\ &\quad \times \left(\frac{1}{N_2} Z_2' H_2 \right) \left(\frac{1}{N_1} H_1' H_1 \right)^{-1} \left(\frac{1}{\sqrt{N_1}} H_1' e_1 \right) \gamma, \end{aligned}$$

where the last lines substitute the first-stage relationship $\sqrt{N_1}(\hat{\pi} - \pi) = ((1/N_1)H_1'H_1)^{-1}((1/\sqrt{N_1})H_1'e_1)$.

Assuming sufficient regularity conditions to invoke a law of large numbers to substitute using (A4)–(A6) and the central limit theorem to use (A7), (A8), and (A10), we have

$$\begin{aligned} \sqrt{N_2} \begin{pmatrix} \hat{\beta} - \beta \\ \hat{\gamma} - \gamma \end{pmatrix} &\rightarrow N(0, \Sigma), \\ \Sigma &= Q^{-1} \left(\Omega_{Zu} + \alpha SR^{-1} \Omega_{H(\epsilon_j)} R^{-1} S' - \sqrt{\alpha} \left(\Omega_{Zu(\epsilon_j)H} R^{-1} S' \right. \right. \\ &\quad \left. \left. + SR^{-1} \Omega'_{Zu(\epsilon_j)H} \right) \right) Q^{-1}. \end{aligned}$$

LEMMA 1.

$$H(H'H)^{-1}H'Z = Z.$$

Proof. Substitute $Z = (XH\hat{\pi})$, $H = (XW)$, and note that Z is in the subspace of H . Thus, the projection of Z onto H yields Z .

Using lemma 1, $SR^{-1}\Omega_{H(\epsilon_j)}R^{-1}S' = \Omega_{Z(\epsilon_j)}$, $\Omega_{Zu(\epsilon_j)H}R^{-1}S' = \Omega_{Zu(\epsilon_j)Z}$, and

$$\Sigma = Q^{-1} \left(\Omega_{Zu} + \alpha \Omega_{Z(\epsilon_j)} - \sqrt{\alpha} (\Omega_{Zu(\epsilon_j)Z} + \Omega'_{Zu(\epsilon_j)Z}) \right) Q^{-1}. \quad (\text{A13})$$

With independent samples and homoskedastic residuals, $\Omega_{Zu(e\gamma)Z} = 0$ and equation (A13) reduces to equation (16) in Inoue and Solon (2010).³⁴ When the two samples fully coincide, one can show using $u = e\gamma + \eta$ and $\alpha = 1$ that $\Sigma = Q^{-1}\Omega_{Z\eta}Q^{-1}$, the standard expression for the 2SLS variance matrix. A feasible estimator of Σ is

$$\hat{\Sigma} = (Z_2'Z_2)^{-1}(\Omega_{Z_2u_2} + \alpha\Omega_{Z_1(e_1\hat{\gamma})} - \sqrt{\alpha}(\Omega_{Z_{1\cap 2}u_{1\cap 2}(e_{1\cap 2}\hat{\gamma}_{1\cap 2})} + \Omega_{Z_{1\cap 2}u_{1\cap 2}(e_{1\cap 2}\hat{\gamma})_{1\cap 2}}))(Z_2'Z_2)^{-1},$$

where the subscript $1 \cap 2$ denotes the intersection of the samples. The functional form of the cross-product matrices $\Omega_{Z_2u_2}$, $\Omega_{Z_1(e_1\hat{\gamma})}$, and $\Omega_{Z_{1\cap 2}u_{1\cap 2}(e_{1\cap 2}\hat{\gamma})_{1\cap 2}}$ depends on the assumed covariance structure of the residuals (homoskedastic, heteroskedastic, clustered, etc.).

References

- Aaronson, Daniel, Ellen Rissman, and Daniel Sullivan. 2004. "Can Sectoral Reallocation Explain the Jobless Recovery?" *Fed. Reserve Bank Chicago Econ. Perspectives* 28 (2): 36–49.
- Abraham, Katharine, and Lawrence F. Katz. 1986. "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *J.P.E.* 94 (3): 507–22.
- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. 2018. "Shift-Share Designs: Theory and Inference." Working Paper no. 24944 (August), NBER, Cambridge, MA.
- Ahn, Hie Joo, and James D. Hamilton. 2014. "Heterogeneity and Unemployment Dynamics." Working paper, Fed. Reserve Board, Washington, DC.
- Alvarez, Fernando, and Robert Shimer. 2011. "Search and Rest Unemployment." *Econometrica* 79 (1): 75–122.
- Angrist, Joshua, and Alan Krueger. 1992. "The Effect of Age at School Entry on Educational Attainment: An Application of Instrumental Variables with Moments from Two Samples." *J. American Statis. Assoc.* 87 (418): 328–36.
- Artuç, Erhan, Shubham Chaudhuri, and John McLaren. 2010. "Trade Shocks and Labor Adjustment: A Structural Empirical Approach." *A.E.R.* 100 (3): 1008–45.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *A.E.R.* 103 (6): 2121–68.
- Barnichon, Régis, and Andrew Figura. 2015. "Labor Market Heterogeneity and the Aggregate Matching Function." *American Econ. J.: Macroeconomics* 7 (4): 222–49.
- Bartik, Timothy J. 1991. *Who Benefits from State and Local Development Policies?* Kalamazoo, MI: W. E. Upjohn Inst. Employment Res.
- Berger, David. 2014. "Countercyclical Restructuring and Jobless Recoveries." Working paper, Dept. Econ., Northwestern Univ.

³⁴ Khawand and Lin (2015) show that our TS2SLS estimator is equivalent to a weighted average of the 2SLS estimator using the overlapping sample and a split-sample 2SLS (SS2SLS) estimator that uses in the second stage only observations not in the first-stage sample. Equation (A13) is more general than their proposition 6 in that it allows for an arbitrary residual structure. Moreover, their proposition is incorrect because they assume independence of the 2SLS and SS2SLS estimators despite the shared first stage and because of several typos.

- Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2018. "Quasi-experimental Shift-Share Research Designs." Working Paper no. 24997 (September), NBER, Cambridge, MA.
- Brainard, S. Lael, and David M. Cutler. 1993. "Sectorial Shifts and Cyclical Unemployment Reconsidered." *Q.J.E.* 108 (1): 219–43.
- Card, David, and Dean Hyslop. 1997. "Does Inflation Grease the Wheels of the Labor Market?" In *Reducing Inflation: Motivation and Strategy*, edited by Christina D. Romer and David H. Romer, 71–122. Chicago: Univ. Chicago Press.
- Charles, Kerwin, Erik Hurst, and Matt Notowidigdo. 2014. "Masking and Aggregate Employment Changes: Housing Booms and Manufacturing Decline during the 2000s." Working Paper no. 18949, NBER, Cambridge, MA.
- Chodorow-Reich, Gabriel, and Loukas Karabarbounis. 2016. "The Cyclicity of the Opportunity Cost of Employment." *J.P.E.* 124 (6): 1563–618.
- Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt. 2016. "Unemployment and Business Cycles." *Econometrica* 84 (4): 1523–69.
- Daly, Mary C., and Bart Hobijn. 2014. "Downward Nominal Wage Rigidities Bend the Phillips Curve." *J. Money Credit and Banking* 46 (S2): 51–93.
- Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2012. "Labor Market Flows in the Cross Section and over Time." *J. Monetary Econ.* 59 (1): 1–18.
- . 2013. "The Establishment-Level Behavior of Vacancies and Hiring." *Q.J.E.* 128 (2): 581–622.
- Davis, Steven J., and John C. Haltiwanger. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Q.J.E.* 107 (3): 819–63.
- . 2014. "Labor Market Fluidity and Economic Performance." Working Paper no. 20479, NBER, Cambridge, MA.
- DeLong, J. Bradford. 2010. "Identifying Cyclical vs. Structural Unemployment: A Guide for Slate Writers." *Brad DeLong's Grasping Reality* (blog), August 24.
- DeLong, J. Bradford, and Lawrence H. Summers. 1988. "How Does Macroeconomic Policy Affect Output?" *Brookings Papers Econ. Activity (Microeconomics)* (2): 433–94.
- Dickens, William T., Lorenz Goette, Erica L. Groshen, et al. 2007. "How Wages Change: Micro Evidence from the International Wage Flexibility Project." *J. Econ. Perspectives* 21 (2): 195–214.
- Dvorkin, Maximiliano. 2014. "Sectoral Shocks, Reallocation and Unemployment in Competitive Labor Markets." Working Paper no. 1229, Soc. Econ. Dynamics, Minneapolis.
- Farmer, Tracy, and Michael Searson. 1995. "Use of Administrative Records in the Bureau of Labor Statistics' Covered Employment and Wage (ES-202) Program." Tech. report, Bur. Census Ann. Res. Conference 1995.
- Foster, Lucia, Cheryl Grim, and John Haltiwanger. 2016. "Reallocation in the Great Recession: Cleansing or Not?" *J. Labor Econ.* 34 (S1): S293–S331.
- Fujita, Shigeru, and Giuseppe Moscarini. 2013. "Recall and Unemployment." Working Paper no. 19640, NBER, Cambridge, MA.
- Garin, Julio, Mike Pries, and Eric Sims. 2013. "Reallocation and the Changing Nature of Economic Fluctuations." Working Paper no. 1258, Soc. Econ. Dynamics, Minneapolis.
- Gertler, Mark, Chris Huckfeldt, and Antonella Trigari. 2015. "Unemployment Fluctuations, Match Quality, and the Wage Cyclicity of New Hires." Working Paper no. 438, Soc. Econ. Dynamics, Minneapolis.
- Glaeser, Edward L. 2005. "Reinventing Boston: 1630–2003." *J. Econ. Geography* 5 (2): 119–53.

- Goette, Lorenz, Uwe Sunde, and Thomas Bauer. 2007. "Wage Rigidity: Measurement, Causes and Consequences." *Econ. J.* 117 (524): F499–F507.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2018. "Bartik Instruments: What, When, Why, and How." Working Paper no. 24408 (March), NBER, Cambridge, MA.
- Groshen, Erica, and Simon Potter. 2003. "Has Structural Change Contributed to a Jobless Recovery?" *Current Issues Econ. and Finance* 9 (8).
- Hall, Robert. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *A.E.R.* 95 (1): 50–65.
- Hall, Robert, and Sam Schulhofer-Wohl. 2015. "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers." Working Paper no. 20939, NBER, Cambridge, MA.
- Hosios, Arthur J. 1994. "Unemployment and Vacancies with Sectoral Shifts." *A.E.R.* 84 (1): 124–44.
- Inoue, Atsushi, and Gary Solon. 2010. "Two-Sample Instrumental Variables Estimators." *Rev. Econ. and Statis.* 92 (3): 557–61.
- Jackman, Richard, and Stephen Roper. 1987. "Structural Unemployment." *Oxford Bull. Econ. and Statis.* 49 (1): 9–36.
- Jaimovich, Nir, and Henry Siu. 2014. "The Trend Is the Cycle: Job Polarization and Jobless Recoveries." Working Paper no. 18334, NBER, Cambridge, MA.
- Kahn, Shulamit. 1997. "Evidence of Nominal Wage Stickiness from Microdata." *A.E.R.* 87 (5): 993–1008.
- Kennan, John, and James R. Walker. 2011. "The Effect of Expected Income on Individual Migration Decisions." *Econometrica* 79 (1): 211–51.
- Khawand, Christopher, and Wei Lin. 2015. "Finite Sample Properties and Empirical Applicability of Two-Sample Two-Stage Least Squares." Working paper, Dept. Econ., Michigan State Univ.
- Kline, Patrick. 2008. "Understanding Sectoral Labor Market Dynamics: An Equilibrium Analysis of the Oil and Gas Field Services Industry." Working paper, Dept. Econ., Univ. California, Berkeley.
- Koenders, Kathryn, and Richard Rogerson. 2005. "Organizational Dynamics over the Business Cycle: A View on Jobless Recoveries." *Fed. Reserve Bank St. Louis Review* 87 (4): 555–79.
- Krugman, Paul. 2014. "Jobs and Skills and Zombies." *New York Times*, editorial, March 30.
- Lilien, David M. 1982. "Sectoral Shifts and Cyclical Unemployment." *J.P.E.* 90 (4): 777–93.
- Mehrotra, Neil R., and Dmitriy Sergeyev. 2012. "Sectoral Shocks, the Beveridge Curve and Monetary Policy." Working paper, Dept. Econ., Columbia Univ.
- Molloy, Raven, Christopher L. Smith, Riccardo Trezzi, and Abigail Wozniak. 2016. "Understanding Declining Fluidity in the U.S. Labor Market." *Brookings Papers Econ. Activity (Microeconomics)* (47): 183–259.
- Murphy, Kevin M. and Robert H. Topel. 1985. "Estimation and Inference in Two-Step Econometric Models." *J. Bus. and Econ. Statis.* 3 (4): 370–79.
- . 1987. "The Evolution of Unemployment in the United States: 1968–1985." In *NBER Macroeconomics Annual 1987*, vol. 2, edited by Stanley Fischer, 11–68. Cambridge, MA: MIT Press.
- Notowidigdo, Matthew J. 2011. "The Incidence of Local Labor Demand Shocks." Working Paper no. 17167 (June), NBER, Cambridge, MA.
- Pacini, David, and Frank Windmeijer. 2016. "Robust Inference for the Two-Sample 2SLS Estimator." *Econ. Letters* 146:50–54.
- Pilossoph, Laura. 2014. "Sectoral Shocks and Move Unemployment." Working paper, Fed. Reserve Bank New York.

- Pissarides, Christopher A. 2009. "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica* 77 (5): 1339–69.
- Sahin, Aysegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. "Mismatch Unemployment." *A.E.R.* 104 (11): 3529–64.
- Schmitt-Grohé, Stephanie, and Martin Uribe. 2016. "Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment." *J.P.E.* 124 (5): 1466–514.
- Schumpeter, Joseph A. 1942. *Capitalism, Socialism, and Democracy*. New York: Harper.
- Shimer, Robert. 2007. "Mismatch." *A.E.R.* 97 (4): 1074–101.
- . 2012. "Reassessing the Ins and Outs of Unemployment." *Rev. Econ. Dynamics* 15:127–48.
- Staiger, Douglas, and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65 (3): 557–86.
- Walsh, Carl E. 2005. "Labor Market Search, Sticky Prices, and Interest Rate Policies." *Rev. Econ. Dynamics* 8 (4): 829–49.