Moving to Fluidity: Regional Growth and Labor Market Churn*

Eran B. Hoffmann — Hebrew University
Monika Piazzesi — Stanford & NBER
Martin Schneider — Stanford & NBER

October 2025

Abstract

This paper studies the connection between regional growth trends and labor market dynamics. New data on manufacturing worker flows for U.S. cities 1969-1981 show more new hires and more voluntary quits in growing cities, but more forced layoffs in shrinking cities. Recessions are special in growing cities in that hires and quits drop, whereas in shrinking cities layoffs rise. A quantitative business cycle model with migration and on-the-job search accounts for a large share of variation in growth and worker flows both over time and across space. Growing cities in the South and West had low job creation costs and only gradual in-migration, so labor markets were tight and encouraged search on-the-job. Unlike shrinking cities in the Rust Belt with slack markets, growing cities responded to aggregate job destruction shocks by a decline in labor market churn.

Keywords:

JEL codes:

*Email addresses: eran.hoffmann@mail.huji.ac.il, piazzesi@stanford.edu, schneidr@stanford.edu. An earlier version was called "Jobs at risk, regional growth, and labor market flows". We thank Tanya Baron, Katka Borovičková, Mike Elsby, Ed Glaeser, Erik Hurst, Lisa Kahn, Philipp Kircher, Ryan Michaels, Antonella Trigari, Eric Smith, Joseph Zeira, and audiences at Hebrew University, Midwest Macro 2019, Search and Matching Conference at BI Norway, 2019, EEA-ESSEM 2019 at Manchester, Frontiers of Macroeconomics at TAU, 2019, NBER conference on Cities, Labor markets, and the Global Economy 2019, and Micro-Macro Labor conference 2019 at the SF FRB, for helpful comments. This research was supported by a grant from the United States - Israel Binational Science Foundation (BSF), Jerusalem, Israel.

1 Introduction

Reallocation of workers happens at different frequencies. On the low end, regional growth trends that reflect movements of workers and jobs across locations unfold over many decades. For example, between 1970 and 2000, manufacturing employment in the Northeast and Midwest regions of the U.S. gradually fell by 23%, while in the South and West regions it increased by 30%. At the same time, continual reallocation of workers occurs in response to idiosyncratic shocks: data show large gross worker flows due to hiring, quits, and layoffs. Finally, the distribution of worker flows changes with the business cycle – for example, more layoffs occur in recessions. While large literatures study both regional change and labor market "churn" over the cycle, the two modes of reallocation are typically not studied jointly.

This paper shows that regional growth trends and gross worker flows are closely connected. Using new data on manufacturing worker flows for U.S. cities in the 1970s, we document that workers in the growing cities of the South and West quit their jobs more often and were laid off less often than workers in the shrinking cities of the Northeast and Midwest. Moreover, what made a recession special in growing cities was a sharp drop in hiring and quit rates, whereas in shrinking cities it was a spike in the layoff rate. A search-and-matching model with city-specific job creation costs and aggregate job destruction shocks quantitatively accounts for large shares of the evolution of the labor force across cities as well as the comovement of worker flows and unemployment over the business cycle.

The empirical part of the paper uses newly digitized archival records from the Bureau of Labor Statistics' Labor Turnover Survey (LTS), a precursor to the widely used Job Opening and Labor Turnover Survey (JOLTS). We combine LTS data with other sources to build a balanced panel of manufacturing employment growth and hiring, quit, and layoff rates for 64 cities during 1969-81, together with unemployment and vacancy rates in selected years. The new dataset thus allows a comparison of distinct regional labor markets for the same broad sector over a long sample that features both large cross-city worker reallocation and three recessions. Since regional heterogeneity in the modern JOLTS data is limited, the new data are of independent interest beyond their use in the present paper.

Figure 1 illustrates the stylized facts we document by comparing two example cities: Buffalo, NY, where manufacturing employment shrank by 27% over our sample and Salt Lake City, UT, where it grew by 69%. Salt Lake City had a much more dynamic labor market: on average, there were twice as many new hires and three times as many workers quit their job voluntarily. Since only half as many workers were laid off, the total separation rate (quits plus layoffs) is only 25% higher than in Buffalo. Finally, recessions look quite different across the two cities. In Salt Lake City, they are marked by a slowdown of worker churn: fewer quits and new hires. In Buffalo, in contrast, we see bursts in layoffs, but relatively little change in quits or hiring.

Our model attributes variation in worker flows across cities to a mismatch between to a mismatch between population size and the cost of job creation that only gradually changed with migration. In the late 1960s, low-cost cities in the South and West had

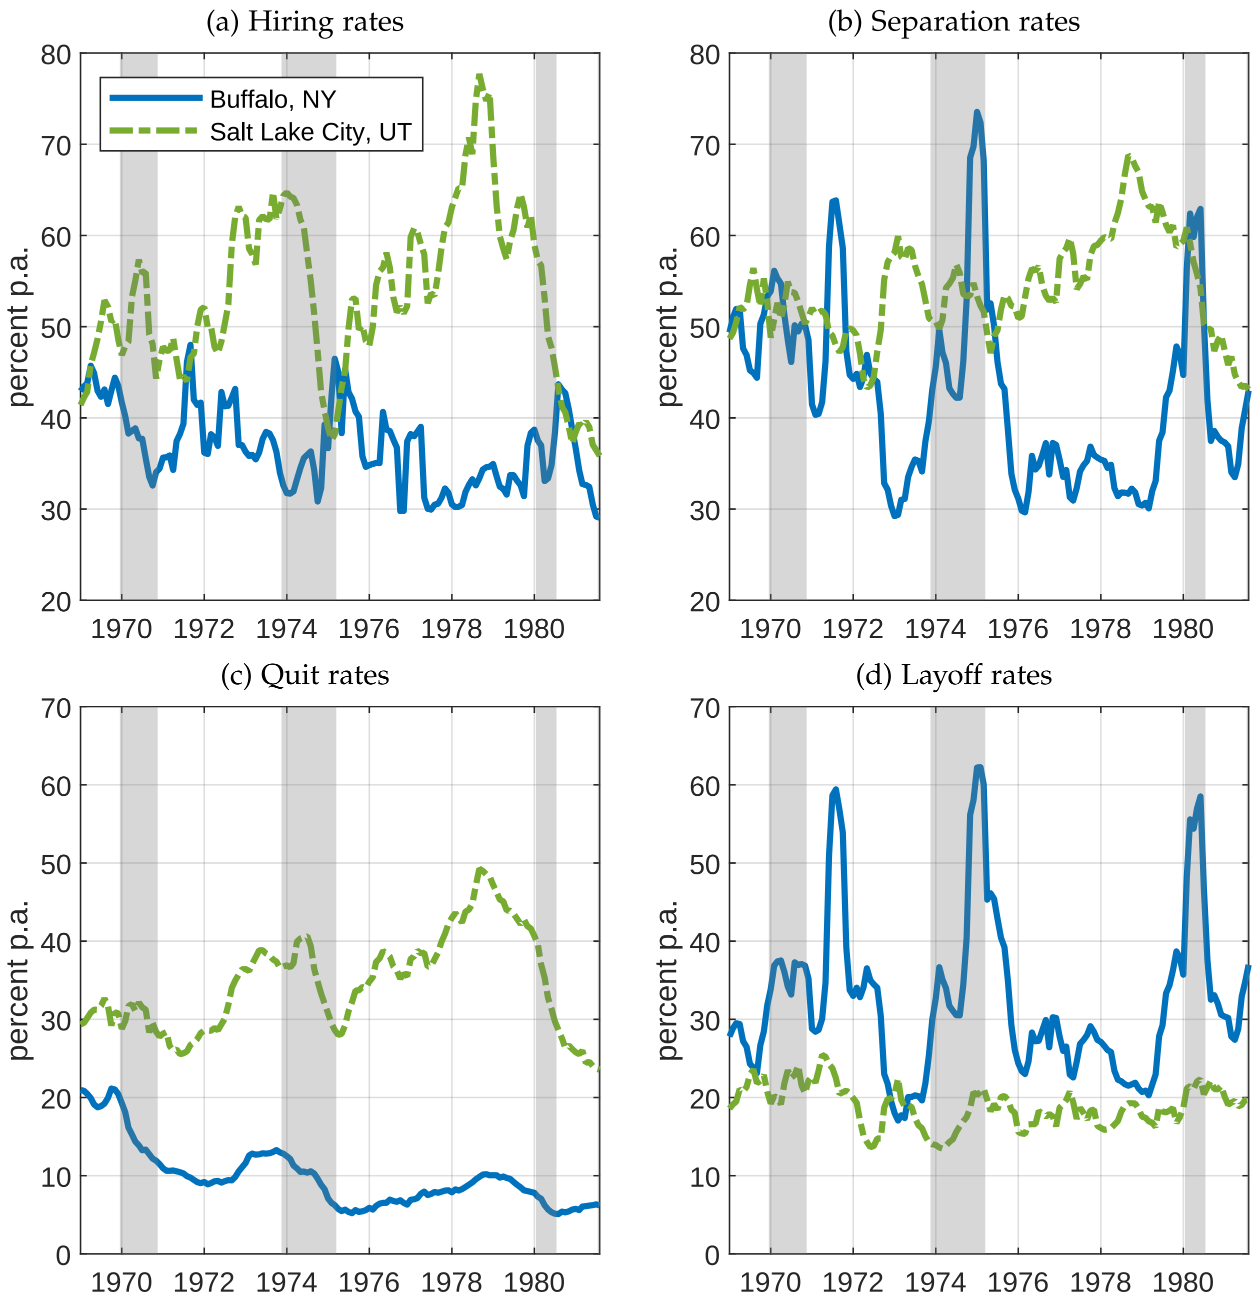

Figure 1: The labor markets of Buffalo, NY, and Salt Lake City, UT

Notes: The figure shows worker flows in manufacturing for Buffalo, NY (blue solid line) and Salt Lake City, UT (green dashed line) from the LTS. All flows are annualized and presented in percentage points. Panel (a) shows the hiring rate, panel (b) the separation rate, panel (c) the quit rate, and panel (d) the layoff rate. All series are seasonally adjusted as explained in the main text, and smoothed using an equally weighted and centered 5-month moving average. Grey bars indicate NBER recessions.

relatively few workers. Job creation made the labor market tight, in the sense of many vacancies per searcher. Workers thus found jobs easily, which not only attracted gradual in-migration, but also encouraged search on-the-job for higher quality matches. This is why quits and hiring were higher in high-growth cities such as Salt Lake City. Moreover, layoffs were lower since on-the-job search moved more workers from unstable matches with high layoff rates to more desirable stable matches. In high-cost cities such as Buffalo, in contrast, few jobs were created, the labor market was slack, and vacancies were hard to find. Workers held on to unstable matches and thus experienced high layoff rates. The overall takeaway is that low frequency movement of workers across cities shapes high frequency churn within cities.

Our quantitative exercise makes three points. First, we estimate city-specific costs of job creation for 64 cities. The quantitative success in the cross-section is that this one exogenous city characteristic generates the strong comovement we observe in all flow variables. Formally, the model almost exactly matches the first principal component in hiring, quits, and layoffs, which accounts for 62% of total variation. Growth is implied as the difference between hiring and separations. The estimated job creation costs—identified only from worker flows—are highly correlated (at .77) with the union participation rate in a city's manufacturing sector. The relationship captures not only differences across the Sun Belt and Rust Belt, but also variation within Census regions. We thus view a high job creation cost as a simple stand-in that captures low surplus due to labor conflict, consistent with the study of the Rust Belt by Alder et al. (2023).

Second, our model helps explain the cross-section of unemployment and vacancies. In cities with low job creation costs, layoff rates are lower and job-finding rates are higher, so the unemployment rate is lower. At the same time, more vacant jobs are available as it is more profitable to create and try to fill jobs. While our data on unemployment and vacancies is not as comprehensive as that on worker flows, we establish for selected years that unemployment indeed declines with growth. We also combine LTS vacancy data and Census unemployment data to derive a downward-sloping cross-sectional Beveridge curve for the year 1970. This evidence suggests that job creation costs, through the mechanism of the model, also drive a significant share of variation in unemployment and vacancies, although there is room for other complementary forces.

Finally, our model generates large fluctuations in worker flows and unemployment over the business cycle that differ across cities, as in the data. We assume that recessions are due to aggregate job destruction shocks. As a shock lowers vacancies and market tightness, search on-the-job becomes less attractive. UNCLEARAverage match quality declines as workers lose both high- and low-quality matches. In growing cities, the former effect is strong and leads to a sharp drop in quits and hiring. In shrinking cities, search on-the job is less common even in booms. A recession thus mostly lowers average match quality and increases layoffs. The quantitative success of our model for the business cycle is that job destruction shocks can account for (i) aggregate volatility of worker flows and unemployment as well as (ii) a large part of cities' relative exposure, as summarized by the first principal component in workers flows.

Our model describes the US economy as a finite collection of cities. Long-lived workers move infrequently and migrate between cities in response to economic conditions, captured by indirect utility. Firms move freely across cities to equate expected profits. Within each city, workers and firms interact in a labor market with search and matching a la Diamond-Mortensen-Pissarides (DMP). An equilibrium describes the evolution of the labor force across cities as well as the distribution of worker states (employed or unemployed) and firm states (vacant or matched) within each city. The fact that jobs can move faster than workers is what leads to long-term differences in market tightness when job creation costs are low in cities with initially small populations.

Labor markets within cities in our model differ from the DMP framework in two important ways. First, we distinguish a job from a match. A job is a long-lived capital good that is costly to create and exists until it is destroyed, as in Mercan and Schoefer (2020). It creates a flow of output while it is matched with a worker. When a worker chooses to leave a match or an unstable match is destroyed, the firm retains the vacant job. Job destruction, in contrast, eliminates the job altogether. Second, matches differ in the intensity of match destruction shocks, as in Nagypál (2005) or Jarosch (2023). In a tighter labor market with more on-the-job search, more unstable matches are upgraded to stable matches. The layoff rate is then lower as less match destruction occurs.

Since jobs are a capital good, our model generates high volatility of unemployment, vacancies, and tightness over the business cycle without assuming sticky wages, unrealistic fluctuations in productivity, or a high value of unemployment. A higher job destruction rate in a recession has two familiar effects: vacancies increase as more unemployed workers are available to fill them, and decrease with the present value of a job. In our model, jobs have long duration, so the present value is particularly sensitive to higher job destruction, and vacancies fall low in recessions. The economy moves into recession along a downward-sloping Beveridge curve. This is in contrast to the standard DMP models, where jobs are identified with short-duration matches. Match destruction shocks are typically not considered a driver of the business cycle in those models since they generate upward-sloping Beveridge curves.

We study the labor force in manufacturing, the sector for which we have detailed data from the LTS. This focus raises two natural questions related to labor force composition. First, how does manufacturing employment compare with the total population? We show that the two are highly correlated across cities in the 1970s. We thus think of our mechanism as capturing a key force behind overall migration, whereas employment in services and other nontradables complements manufacturing in a similar way everywhere. Second, do manufacturing worker flows simply reflect differences across cities in sectoral composition within manufacturing? We use a shift-share approach to show that our estimated job creation cost, and hence the first principal component of worker flows, does not correlate much with growth rate differentials due to sectoral composition.

We add to a large literature on how local labor markets change over time; for a survey, see Moretti (2011). Most studies focus on wages and employment. Early work interpreted the data through the lens of static spatial equilibrium models, as in Hall (1972) and Rosen (1979). Recent emphasis is on the dynamic relationship between wages and employment and the response of migration to persistent cross-sectional differences, for example, Topel (1986), Blanchard and Diamond (1990), Redding (2016), Amior and Manning (2018). In particular, Yoon (2017) and Alder et al. (2023) present models of the decline of the Rust Belt, which they attribute to changes in productivity and union activity, respectively. Our approach differs from this literature because the local conditions we study are statistics on gross worker flows, rather than wages or employment rates.

Our model builds on studies that connect migration and job search. Kennan and Walker (2011) develop a tractable choice model of worker location choice to estimate the response of migration to local conditions, a key parameter also in our model. Schmutz and Sidibe (2019) estimate such a model in French data and show that workers respond to opportunities for better job-to-job transitions, a key force in our model as well. Recent models of city wage differentials similarly emphasize the importance of on-the-job-search and match heterogeneity. They study its interaction with within-city firm heterogeneity (Heise and Porzio, 2022), knowledge diffusion within cities (Martellini, 2022), or firm sorting across cities (Lindenlaub et al., 2022). While we share the emphasis on search and migration, our focus is on equilibrium relationships between gross worker flows, both across space and over the business cycle.

There is relatively little evidence on gross worker flows by geography over time. Hall (1972) compared layoff rates in 12 large cities in the 1966 and shows that tighter markets see fewer layoffs. Davis and Haltiwanger (2014) show that worker reallocation (hiring plus separation) moves with the employment rate across U.S. states since 1998. Bilal (2023) studies unemployment and layoffs by region in French micro data that show strong comovement of unemployment and separation rates across cities. Kuhn, Manovskii, and Qiu (2021) study data from Germany and the UK as well as US JOLTS data for 18 MSAs since 2001 to show that labor markets with lower unemployment are tighter, and have higher separation rates and job-finding rates out of unemployment, consistent with our findings. We complement this literature by presenting a large US panel that decomposes separations into quits and layoffs and relates them to the business cycle as well as regional growth during a period of high migration rates.

We also contribute to an active literature on regional differences in unemployment. The typical approach is to embed a frictional labor market into a stationary spatial equilibrium framework where workers are indifferent across locations. For example, Beaudry, Green, and Sand (2012, 2014), and Kline and Moretti (2014) study the relationship between amenities, policy, and unemployment. Bilal (2023) estimates a multi-city search model with spatial sorting of firms and endogenous separation to explain a strong link between separations and unemployment. Kuhn et al. (2021) also consider a model with endogenous separation but add on-the-job search to generate a downward-sloping cross-sectional Beveridge curve. We differ from these papers in that we target data on gradual migration as well as quits and layoffs, which call for a nonstationary setup and different forces driving the cross-section and the business cycle.

The literature on *job flows* has studied reallocation and growth by industry and firm.[1] Foote (1998) notes that job creation and destruction are larger in growing sectors which also have a higher volatility of job creation than job destruction. He proposes an (S,s) employment adjustment model that qualitatively captures this relationship. Davis, Faber-

[1]Job flows are measured using micro data on firms, recording job creation by new or expanding establishments as well as job destruction by shrinking establishments. See Davis and Haltiwanger (2014) for a recent overview.

man, and Haltiwanger (2012) show that job reallocation rates are higher in retail and services than in manufacturing and higher at growing firms. Şahin, Song, Topa, and Violante (2014) emphasize sectoral differences in tightness. While our focus is on flows by region, our mechanism could be applicable for sectoral growth and churn as well. Studying this connection would require a model with a richer distinction between job and worker flows, see for example Kiyotaki and Lagos (2007), Burgess and Turon (2010), or Borovičková (2016).

Finally, our findings on worker flows in recessions relate to business cycle models with search. The typical approach studies labor reallocation at the national level. There has been some debate on the source of shocks: the relative importance of shocks to flow surplus and job destruction shocks depends on the relative cyclicality of job finding and separation rates (Shimer, 2012; Fujita and Ramey, 2009; Elsby, Michaels, and Solon, 2009; Yashiv, 2007). Our results show that these properties differ by region, and that a model with on-the-job search can account for differential exposure of regions to aggregate shocks. We thus add to a growing literature that explores the role of on-the-job search and variation in match quality to address the volatility of labor market flows over the business cycle (for example, Menzio and Shi, 2011; Fujita and Nakajima, 2016; Gertler, Huckfeldt, and Trigari, 2016; Lise and Robin, 2017; Moscarini and Postel-Vinay, 2018).

The paper is structured as follows. Section 2 presents our data and documents stylized facts. Section 3 introduces the model and explains the mechanism. Section 4 presents the quantitative implementation. Section 5 concludes.

2 Data and facts

This section describes our data and presents facts on the relationship between labor turnover, regional growth, and aggregate fluctuations.

2.1 Data sources

Separations. Our main data source is the Labor Turnover Survey (LTS), an establishment survey conducted by the Bureau of Labor Statistics (BLS) from 1931 to 1981. The LTS was a precursor to the currently run Job Opening and Labor Turnover Survey (JOLTS) that started in the year 2000. It sampled approximately 40,000 establishments each month, covering about 40% of employment in manufacturing. It was discontinued at the end of 1981 due to "severe budgetary cutbacks" (Utter 1982). Starting in 1958, data on turnover rates in manufacturing were also collected at the state and city level under a cooperative program with state statistical agencies. Key statistics were published monthly as part of the BLS "Employment and Earnings".[2]

We have hand-collected and digitized the LTS tables to form a panel of city-level flows. A typical observation in the raw data is a monthly labor flow—for example, the number of workers who were separated from a job—reported as a share of manufacturing employment in a geographic unit in percentage points with one decimal point

[2]Digitization is based on monthly issues of The United States Bureau of Labor Statistics, *Employment and Earnings* (January 1958-December 1981).

accuracy. We transform these monthly rates into annualized, continuously compounded rates. In particular, for any separation rate x, such as the quit rate, in percentage points, we compute the annualized rate $-100 \times 12 \times \log(1 - x/100)$. For the hiring rate, we calculate the annualized rate $100 \times 12 \times \log(1 + x/100)$. We remove seasonality from each flow, at the city level, using the X-13ARIMA-SEATS Seasonal Adjustment Program from the Census Bureau.[3]

We classify separations into the two standard categories used in the literature: voluntary separations initiated by the worker (“quits") or forced separations initiated by the employer (“layoffs"). Separations in the raw LTS data are instead classified into “quits", “layoffs", and “other". The LTS definition of layoffs is narrow in that it includes only forced separations from an employer “without prejudice to the worker". “Other" separations are defined as “terminations of employment because of discharge, permanent disability, death, retirement, transfers to another establishment of the company, and entrance into the Armed Forces expected to last more than 30 consecutive calendar days." Our baseline specification adds all “other" separations to layoffs, in order to match the total separation rate, and hence the net hiring rate. The stylized facts we present are robust to simply omitting the “other" category.

Employment and hiring. Starting in 1969, the Bureau of Economic Analysis (BEA) provides annual estimates of the number of full-time and part-time employees by industry and Metropolitan Statistical Area (MSA).[4] These estimates are based on quarterly tax filings from all business establishments as well as local statistics. We focus on full-time employment for all manufacturing (Table CAEMP25S, line code 400, industry D). For each Metropolitan Statistical Area (MSA), we compute the growth rate of manufacturing employment. We also use employment levels to weigh city-level flows and to initialize our model.

A drawback of LTS data on hiring rates is that the survey design oversamples large, established plants that tend to grow more slowly and create fewer jobs than newer establishments. Utter (1982) provides an overview of studies that compare LTS data to alternative sources for the 1970s, including Social Security Administration records. She concludes that the bias is considerably smaller for separation rates than for hiring. This is reassuring for our application since we do not have an alternative data course for separations. Consistent with these findings, we observe that BEA employment growth rates are systematically higher than growth rates implied by the LTS, defined as gross hiring less separation. Nonetheless, the two measures are highly correlated in the cross-section (correlation of .70), so our cross-sectional results are not sensitive to the choice of employment growth measure.

Given this data situation, we adopt two conventions for documenting facts and selecting calibration targets for our quantitative exercise. First, for all cross-sectional averages, we rely on annual BEA employment growth rates and define hiring rates as the sum of BEA growth rates and LTS separation rates. The constructed hiring rates are nearly identical to the raw LTS measures, with a correlation of .99. Second, for documenting

[3]We used MATLAB code to access the program, from Yvan Lengwiler, ’X-13 Toolbox for Matlab, Version 1.55’, Mathworks File Exchange, 2014-2023.

[4]Data are available at https://apps.bea.gov/regional/downloadzip.cfm.

facts about the business cycle, where monthly data is useful for accurately timing recessions, we use LTS data to compute differences in rates across booms and recessions. Any systematic bias in LTS levels should largely cancel out in these differences, as such a bias would act as a common shifter.

Vacancies. From June 1969 to December 1973 the BLS establishment survey underlying the LTS included a question on job vacancies in manufacturing. In collaboration with state agencies, the BLS compiled monthly vacancy rates for a sample of cities beginning in April 1970. Because the available data contain too few observations to construct a reliable panel, we focus instead on a single cross section of vacancy rates in June 1970 that we can compare to unemployment measures from the 1970 Census to derive a cross-sectional Beveridge curve.[5] Vacancy rates are published as the ratio of vacancies to employment in manufacturing, and reported in percentage points with one decimal point accuracy.

Unemployment. From 1977 to 1979, annual local area unemployment rates are available from the BLS and state agencies. The estimates are based on several sources, including the Annual Social and Economic Supplement (ASEC) of the Current Population Survey and statistics on unemployment insurance claims. To align with our vacancy data for 1970, we additionally construct MSA-level unemployment measures from the 1970 Census. In this case, we focus on individuals aged 15 to 65 who either report current employment in manufacturing or, if unemployed, report that their last job was in manufacturing.

Geographic units. The geographic unit for our panel is an MSA, which we refer to as a *city*. In the raw LTS data, geographic units are a mix of cities, counties, and "Standard MSAs" that differ over time. We begin with the list of MSAs reported in the BEA data and name-match them to the LTS by city and state. Partial matches are allowed when the central city is the same—for example, the city of Phoenix, AZ, is matched with the Phoenix–Mesa–Scottsdale MSA. When multiple counties or cities in the LTS correspond to a single MSA, we compute weighted averages of flow variables using annual county-level employment as weights. For instance, Dallas and Fort Worth are consistently treated as one MSA, even though they appear separately in early LTS years and jointly as a Standard MSA later on.

Our final sample consists of 64 cities, listed in Appendix A, observed over 151 months from January 1969 to July 1981. We include MSAs with at least 135 monthly observations of separations, more than 15,000 manufacturing workers at the beginning of the sample, and unemployment data available for at least one of the years 1977–1979. The most notable omissions are cities in California, a state that did not participate in the LTS during most of the sample period. We exclude Mobile, Alabama, and Johnstown, PA, as extreme outliers in all flow rates.[6] Finally, to construct aggregate labor flows at the national level, we compute monthly averages of city-level flows weighted by each city's

[5] See "Employment and Earnings", October 1970, table E-4, page 104.

[6] For example, the reported percent of laid off workers in Mobile is, in many months, in the double digits *per month*, or more than 100% turnover in yearly terms. Including these cities in the analysis would not significantly alter any of the main results.

manufacturing employment in that year.

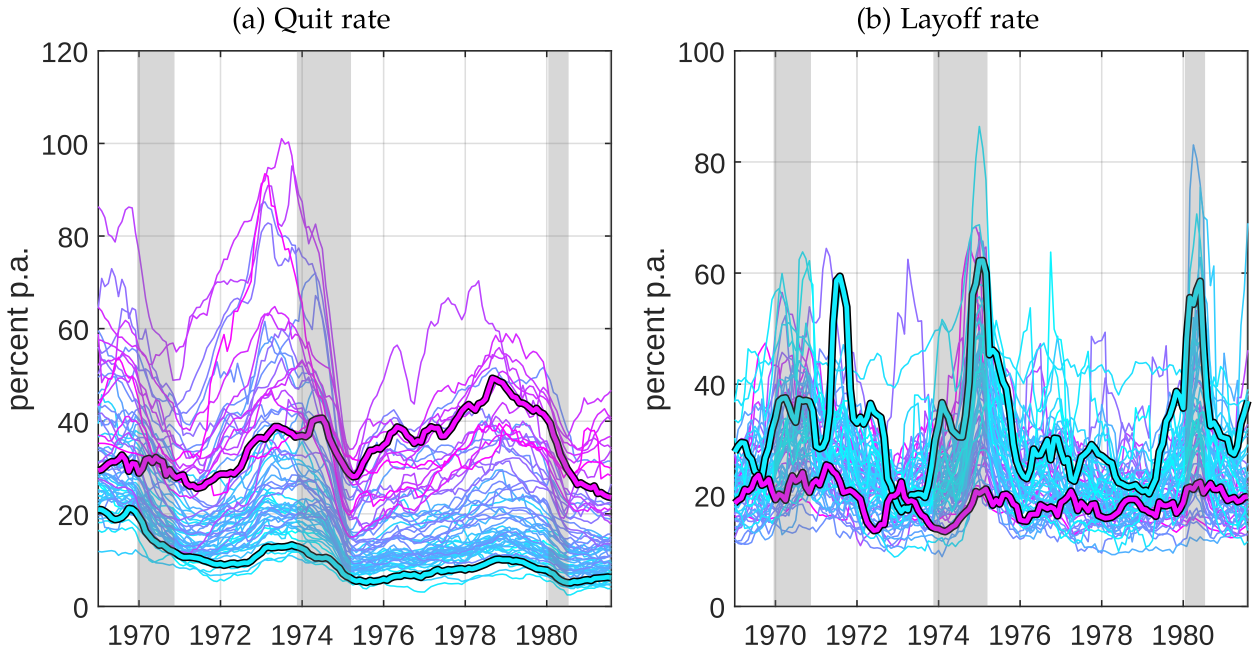

Figure 2: Quits and layoffs by city

Notes: This figure shows separation rates in manufacturing for each city in the sample. Panel (a) shows quit rates, while panel (b) layoff rates. All series are seasonally adjusted and detrended as explained in the main text. Line color indicates the mean employment growth rate. Light blue indicates lower growth rates, while light purple indicates higher growth rates. The thicker lines represent the separation rates for Buffalo, NY (low growth) and Salt Lake City, UT (high growth).

The city panel of worker flows. As a first look at the LTS data, Figure 2 displays quit rates in panel (a) and layoff rates in panel (b). Each line represents a city, with color indicating the mean growth rate of the city over the sample from light blue (lower growth) to light purple (higher growth). Even in a crowded figure of the entire panel, we can see the basic patterns exemplified by Buffalo and Salt Lake City in Figure 1. Faster-growing cities exhibit higher average quit rates and larger cyclical fluctuations between booms and recessions, shaded according to NBER dating. In contrast, their layoff rates tend to be lower and more stable. We now proceed to study the cross-sectional and time-series properties in more detail.

2.2 The cross section of worker flows

Separations, unemployment, and growth. Figure 3 summarizes our three key cross-sectional facts: *faster-growing cities exhibit more quits, fewer layoffs, and lower unemployment*. Each scatter plot relates the city-level average of quit rates, layoff rates, or unemployment rates to the city's average employment growth rate. Each dot represents a city, with colors and symbols identifying census regions. We also show nonparametric regression lines with 95% confidence intervals for reference. The patterns for quits and unemployment are very strong. For layoffs, the overall downward slope is weaker, in part because the shape of the curve changes around zero. This feature is also present to a lesser degree in the other two panels: the association with growth is stronger for shrinking cities.

The coloring underscores that the patterns are present both within and across regions. In particular, the cross-regional component represent movement from the Rust Belt to the Sun Belt during the 1970s: cities in the Northeast and Midwest were shrinking, whereas cities in the South and West were growing. The shift in population was sizeable: the average city in the South or West grew by 1.6% per year, or 21% over our 12-year sample, whereas the average city outside the South or West shrank by 1.1% per year or 13% overall. Average worker in the South or West quit their job twice as often as other workers, were laid off 5% less often and were 18% less likely to be unemployed.

Figure 4 compares hiring and separation in the cross-section: *when cities grow faster, more workers are hired and more workers separate from jobs.* Because separations equal quits plus layoffs, and hiring equals growth less separations, the figure does not really contain new information relative to Figure 3, but just looks at the same data from a different angle. The message is that gross hiring and separation rates look very similar: both are large in growing cities, consistent with many job-to-job transitions prompted by quits. In particular, the separation rate inherits the upward sloping relationship with growth from the quit rate even if the (downward sloping) layoff rate is added. A the same time, adding the (upward-sloping) growth rate to quits does not generate a big difference between hiring and quits.

Unemployment and layoffs. While Figure 3 already suggests a positive relationship between unemployment and layoffs, it is interesting to look at the cross-sectional relationship directly. Here we follow Bilal (2023) who proposes a decomposition based on a closed search market: if u unemployed workers find jobs at the rate f and $1-u$ employed workers are laid off at the rate λ, equating outflow and inflows to keep unemployment constant

$$\log(u/(1-u)) = \log \lambda - \log f.$$

Figure 5 plots the cross section of the LHS against $\log \lambda$. If city-specific differences in job destruction were the only force driving differences in unemployment, all dots would lie on the 45-degree line.

As expected, we find an upward-sloping relationship: unemployment moves with layoffs. The relationship is relatively weak, however: a linear regression delivers an R^2 of only 32%, leaving ample room for other variation. In particular, the typical Southern (Northeastern) city appears to have much lower (higher) unemployment than what one would expect based on variation in layoff rates alone. This finding is different from, for example, the French data studied by Bilal (2023), where layoffs and unemployment are tightly connected. In the 1970s US, cross-city differences in job finding rates and migration must have played a more prominent role.

The cross-sectional Beveridge curve. Figure 6 displays a scatter plot of unemployment against vacancy rates for the year 1970 for which both numbers are available for a cross-section of cities. The evidence here is not particularly strong, given the small number of cities. Nevertheless, a negative relationship is visible: *in cities with more unemployed workers, firms also post fewer vacancies*. The cross-sectional relationship between rates for different labor markets in one year is thus similar in shape to the familiar Beveridge curve that emerges when the rates are plotted at different dates for the same market, such as the US as a whole.

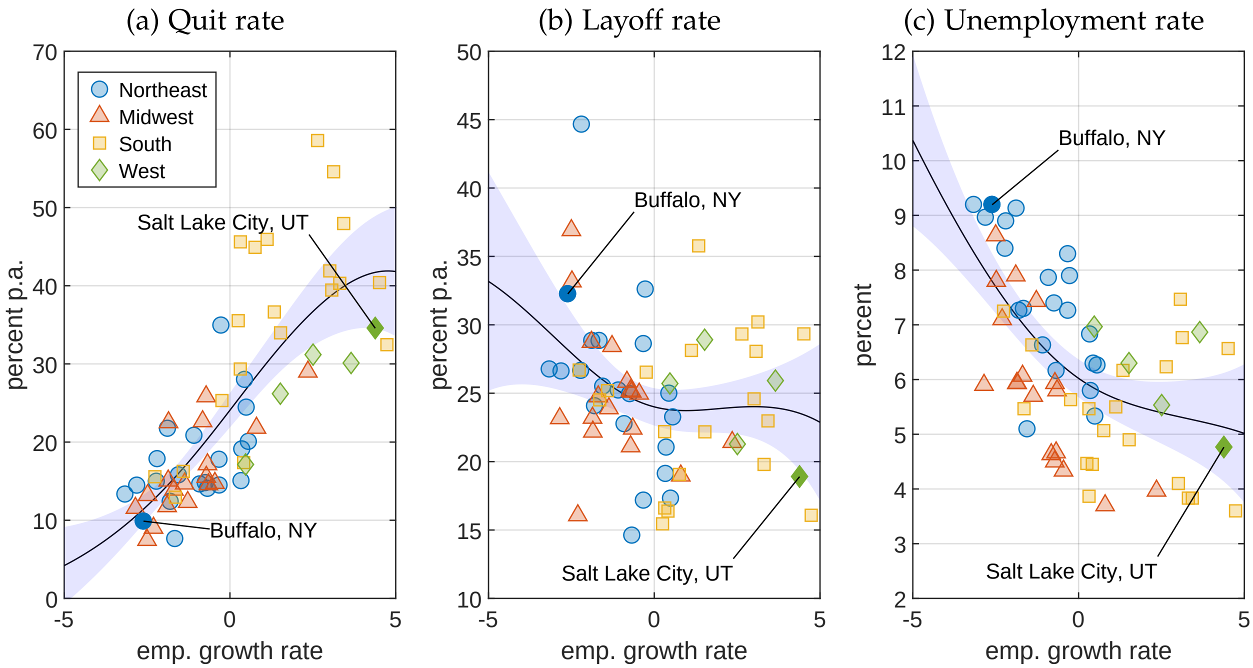

Figure 3: The cross-section of quit, layoff, and unemployment rates

Notes: The figure shows the cross-sectional joint distribution of mean employment growth in manufacturing and the mean quit rate (panel (a)), layoff rate (panel (b)), and unemployment rate (panel (c)). Each dot represents a city. Marker shape and color represent the census region of the city. Quit and layoff rates are in percentage points per year based on the LTS. Unemployment rate is the mean for the estimated unemployment rate in March of the years 1977-1979, based on CPS and state statistical agencies. The solid line is a local weighted regression (LOWESS) using a Gaussian kernel with bandwidth 2. Shaded area is 95% confidence interval based on bootstrapped standard errors with 1000 replications.

2.3 Worker flows over the business cycle

A stark pattern in Figure 2 is that quit rates fall and layoff rates rise together in all cities following aggregate downturns, indicated by NBER recessions shaded in gray. To obtain a parsimonious representation of cyclical comovement that focuses on manufacturing employment, we define a recession slightly differently from the NBER: Appendix B uses a clustering algorithm to classify months as boom or recession depending on the movement of aggregate layoff and hiring rates. A similar approach of detecting national recessions in labor market data to document city-level cyclical patterns was applied by Chodorow-Reich and Wieland (2020).

Figure 7 displays aggregate hiring, quit, and layoff rates together with our recession indicator as a purple line that is high in recession, and also shades NBER recessions as a benchmark. The algorithm clearly picks up three contiguous periods of unusual movement in worker flows. Those episodes typically occur slightly later than each of the three NBER recessions in our sample, for which dating relies relatively more on

Figure 4: The cross-section of gross hiring and separation rates

Notes: The figure shows the cross-sectional joint distribution of mean employment growth in manufacturing and the mean hiring rate (panel (a)) and separation rate (panel (b)). Each dot represents a city. Marker shape and color represent the census region of the city. Hiring and separation rates are in percentage points per year based on the LTS. The solid line is a local weighted regression (LOWESS) using a Gaussian kernel with bandwidth 2. Shaded area is 95% confidence interval based on bootstrapped standard errors with 1000 replications.

movements in aggregate GDP. The episodes are also similar in length. We thus maintain the language "boom" and "recession" when we refer to our classification.

Figures 8 presents facts on worker flows over the cycle. As before, the horizontal axis shows average employment growth over the entire sample, but we now plot the average difference between booms and recessions in each panel. First impressions are that hiring and quit rates are typically procyclical (dots are located above zero on the vertical axis), whereas layoff rates are countercyclical (dots are below zero). Moreover all nonparametric regression lines are upward sloping, that is, the difference between booms and recessions is larger (possibly less negative) in faster-growing cities.

Comparing how cities locate in the four quadrants reveals the stark difference between cycles in growing versus shrinking cities. In growing cities with dynamic labor markets, a recession mostly dampens this dynamism: quits and hiring decline. While there are some more layoffs, separations overall are basically unchanged. In shrinking cities, in contrast, hiring and quits change little from their lower average. A recession mostly brings many more layoffs and hence separations. Interestingly, the dynamics we observe in growing cities is similar to the conventional wisdom about the aggregate US economy. Shirnking cities, in contrast, behave very differently.

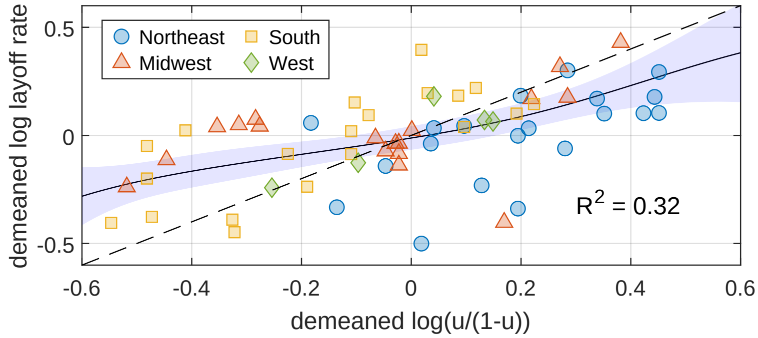

Figure 5: Unemployment and layoffs

Notes: The figure shows the cross-section of the mean layoff rate in booms and the unemployment rate in 1977-1979. Each dot represents a city. Marker shape and color represent the census region of the city. The layoff rate is the mean taken over all boom months of LTS data. The solid line is a local weighted regression (LOWESS) using a Gaussian kernel with bandwidth 2. The shaded area is a 95% confidence interval based on bootstrapped standard errors with 1000 replications.

3 Model

We describe the US economy as a finite collection of cities. Workers are long-lived, move infrequently, and migrate between cities in response to differences in indirect utility. Firms move freely across cities to equate expected profits. Within each city, workers and firms interact in a frictional labor market with costly search on the job, following Diamond, Mortensen, and Pissarides (DMP). In contrast to DMP, jobs are long-lived capital goods, and match quality is reflected in the duration of the match.

Starting from an initial distribution of population taken from the data, two exogenous forces drive differences in worker flows across cities and over time. First, cities differ in the cost of creating new jobs. Second, cities experience common business cycle fluctuations, generated by common changes in the job destruction rate. An equilibrium thus describes the evolution of the labor force across cities as well as worker states (employed or unemployed) and firm states (vacant or matched) within each city, along with wages and search effort. The model generates a sample of quit, layoff, and unemployment rates that we can directly compare to our data.

3.1 Setup

Time is continuous. Cities are indexed by $i = 1, \ldots, N$. A continuum of workers live in every city. Let $L_t(i)$ denote the population of workers in city i at date t and write $\bar{L}_t$ for the total US population. A single homogeneous consumption good is tradable across cities. Exogenous US-wide technical progress grows at the rate g and affects productivity, unemployment benefits, the cost of search, and job creation. Aggregate risk is captured by a Markov chain with two states b and r indicating booms and recessions, respectively. We denote by π_{jk} the intensity of jumping to state k when the current state is j.

Preferences and migration. Workers have linear utility over consumption and discount the future at a constant rate $r > g$. At any date, workers are either employed, earn-

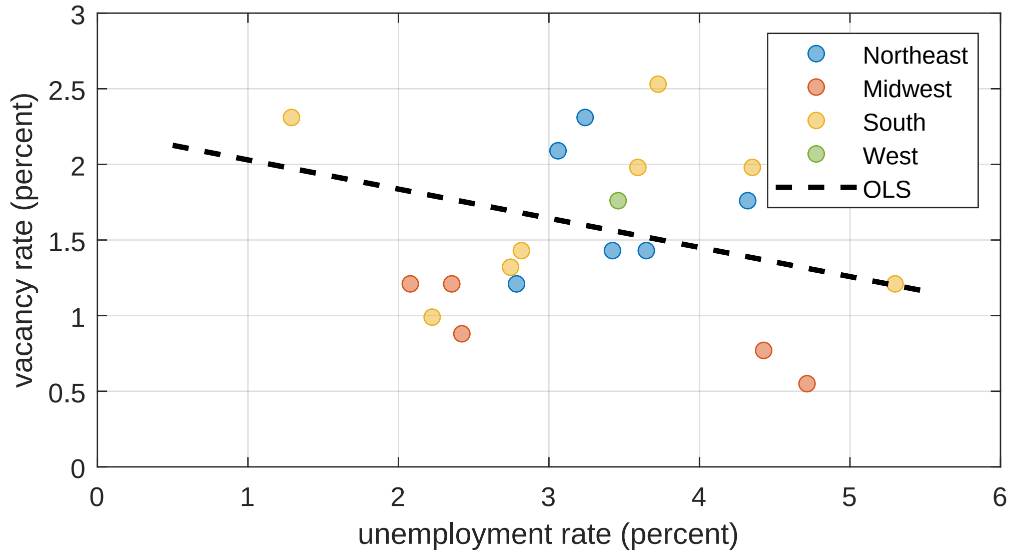

Figure 6: Cross sectional Beveridge curve in 1970

Notes: Vacancy rates are the mean of April-May 1970 vacancy rate reported by the BLS and multiplied by 2.11, as suggested by Abraham (1983). Unemployment rates are calculated based on the 1970 census (metro1 and metro2 1% samples), for workers 15-65 year old with self-reported current or last job in manufacturing. The dashed line is an equally weighted OLS regression line.

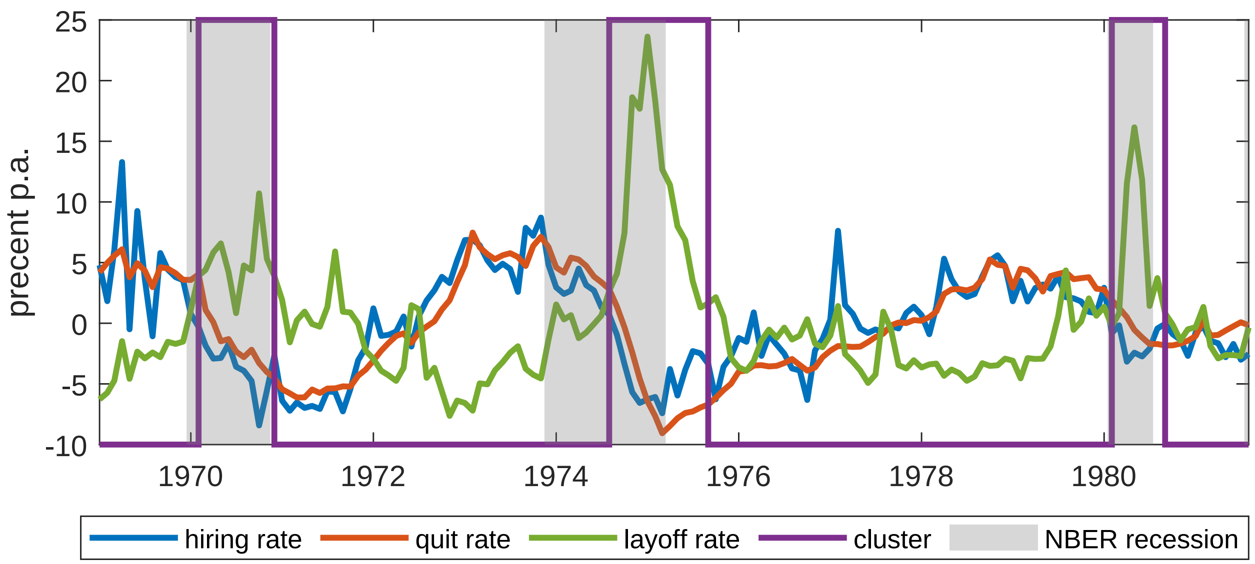

Figure 7: Aggregate labor-market flows and aggregate states

Notes: The figure shows the aggregate hiring, quit, and layoff rates together with the timing of NBER recessions and adjusted timing derived from a mixture of normals clustering algorithm as described in Appendix B. The clustered recessions overlap with the three NBER recessions but start later and last longer.

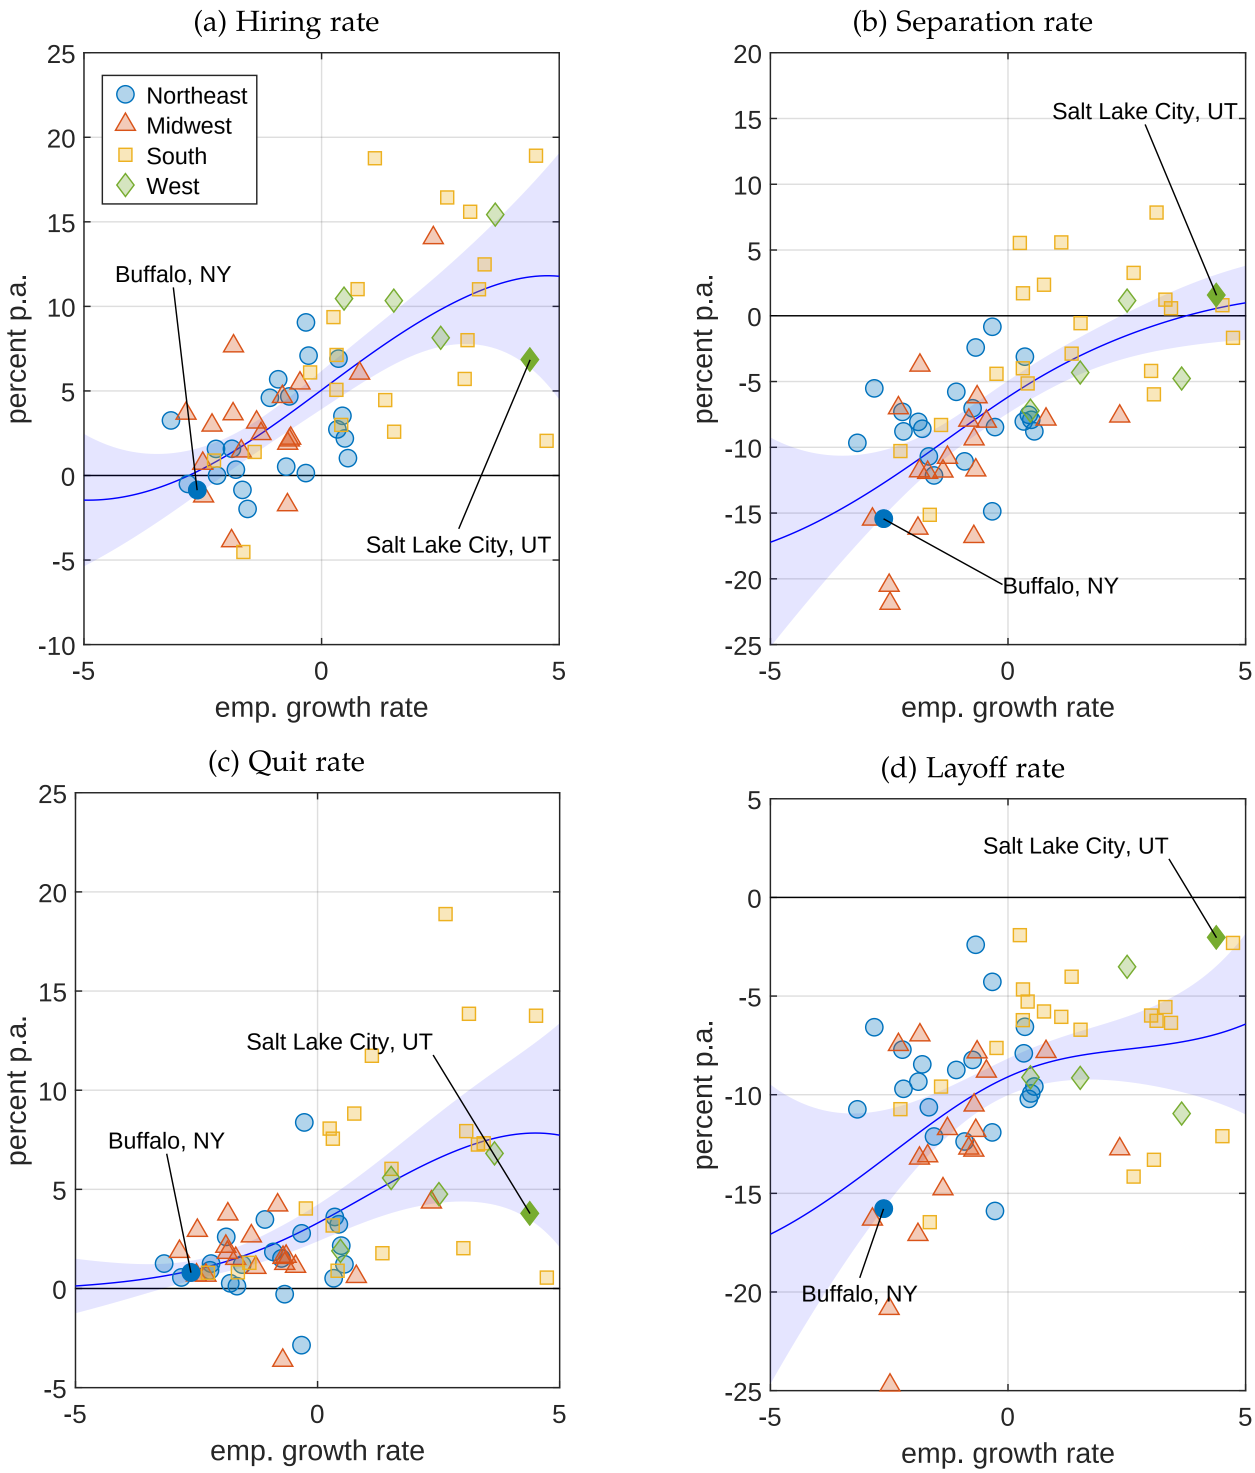

Figure 8: Workers flows in booms vs recessions

Notes: The figure shows the gap between the mean flows in booms and recessions in the cross-section of cities. Panel (a) shows the gap in the hiring rate, panel (b) the gap in the separation rate, panel (c) the gap in the quit rate, and panel (d) the gap in the layoff rate. Each dot represents a city. Marker shape and color represent the census region of the city. Flows are in percentage points per year based on the LTS data. The solid line is a local weighted regression (LOWESS) using a Gaussian kernel with bandwidth 2. The shaded area is a 95% confidence interval based on bootstrapped standard errors with 1000 replications.

ing a wage, or they are unemployed, receiving benefits $b_t = be^{gt}$. Both employed and unemployed workers can search for new jobs in the cities they live in. Unemployed workers supply one unit of search effort inelastically and at no cost. By contrast, on-the-job search by employed workers is costly: exerting effort e_t entails a consumption cost $c_t(e_t) = c(e_t)e^{gt}$, where the effort cost function $c(\cdot)$ is strictly increasing, concave, satisfies $c(0) = 0$, and has $c'(0) = 0$.

Workers leave a city only upon receiving an exogenous moving shock, which arrives at rate μ in every city. Movers exit the model and are replaced by new workers who are randomly assigned to a destination city where they start out unemployed. The assignment depends on the indirect utilities $U_t(i)$ that unemployed workers obtain in different cities. Rather than specifying the distribution explicitly, we directly assume that the growth rate of the labor force in city i is

$$\frac{\dot{L}_t(i)}{L_t(i)} = \nu + \mu \left(\frac{U_t(i)^\varepsilon}{\sum_{j=1}^{N} \frac{L_t(j)}{\bar{L}_t} U_t(j)^\varepsilon} - 1 \right). \tag{1}$$

where the parameter $\varepsilon > 0$ regulates migrants' sensitivity to local conditions.

Total population $\bar{L}_t$ grows at the constant rate ν. City i grows faster than the overall population growth rate ν if and only if it offers relatively more utility $U_t(i)$. Higher sensitivity ε means worker flows are more directed towards high utility cities. If $\varepsilon = 0$, workers do not respond to local conditions and are equally likely to land in any city, so there is no net migration: all cities grow at the same rate. There is still gross migration, however, regulated by the parameter μ. In the limit as $\varepsilon \to \infty$, all indirect utilities have to be equated in equilibrium. The typical case has both gross and net migration, regulated by μ and ε, respectively.

As a microfoundation for the migration equation, suppose that each city consists of many identical neighborhoods of equal population size, so larger cities contain more neighborhoods. Continuation utility for neighborhood k in city i is $\xi_{t,k}U_t(i)$, where the shocks $\xi_{t,k}$ are drawn independently across all neighborhoods in the country from a Fréchet distribution with shape parameter ε. If there are $K_t(i)$ neighborhoods in city i at date t, the net migration into city i is then

$$\dot{L}_t(i) = (\nu - \mu)L_t(i) + \mu \bar{L}_t K_t(i) \frac{U_t(i)^\varepsilon}{\sum_j K_t(j) U_t(j)^\varepsilon}$$

Here, the first term is the contribution of city i to common population growth net of gross out-migration. The second term is gross in-migration, given by the product of (i) all migrants in the country $\mu\bar{L}_t$, (ii) the number of neighbourhoods in city i, and (iii) the probability that any neighborhood in the city is a migrant's favorite neighborhood, which takes a simple form due to the Fréchet distribution of the preference shocks. Imposing that the number of neighborhoods is proportional to the size of its labor force results in equation (1).

Technology. Jobs are capital goods made from consumption goods. The cost of job creation is city-specific: one new job in city i requires $\sigma_t(i) = \sigma(i)e^{gt}$ units of consumption. When a job is matched with a worker, it produces output at the rate $y_t = e^{gt}$. Matches end when a worker quits because of a moving shock or after searching on the job, or when the worker is laid off because of an exogenous job or match destruction shock. There are two types of matches, high and low quality. The only difference is that low-quality matches break due to match destruction shocks that arrive at rate δ.

All matches are ex ante indistinguishable: match quality is determined by an iid draw and revealed only after the match has formed. The probability of a high-quality match is p. After a match ends, the job remains in place and can match with another worker. Jobs are destroyed by exogenous job destruction shocks. The rate of job destruction ρ_t depends on the state of the business cycle. In addition, we assume no adjustment costs to capital, so jobs can be transformed back into consumption goods. We can therefore think of job creation as reversible investment and job destruction shocks work on aggregate like depreciation of capital.

Search and matching. The labor market in every city is described by a constant-returns-to-scale matching function. If city i has $v_t(i)$ vacant jobs, that is, jobs not matched to a worker, and searching workers exert $s_t(i)$ total units of effort, then matches are formed at the rate $m(v_t(i), s_t(i))$, where the function m is increasing in both arguments and homogeneous of degree one. We assume that all jobs are equally likely to be matched. With constant returns, the rate at which vacant jobs in city i are filled can then be written as a decreasing function of the *tightness* of the labor market $\theta_t(i) = v_t(i)/s_t(i)$:

$$\frac{m(v_t(i), s_t(i))}{v_t(i)} = m(1, \theta_t(i)^{-1}) =: \phi(\theta_t(i)) \tag{2}$$

In a city with a tighter market it is harder to fill a vacant job so the job filling rate $\phi(\theta_t)$ is lower.

On the worker side, we assume that the likelihood of finding a job is proportional to search effort. The rate at which unemployed workers find jobs can be written as an increasing function of tightness:

$$\frac{m(v_t(i), s_t(i))}{s_t(i)} = m(\theta_t(i), 1) =: f(\theta_t(i)) \tag{3}$$

Here we are using the assumption that every unemployed worker exerts effort $e_t = 1$. The rate at which employed workers find jobs is then $e_t f(\theta_t(i))$ and hence scales with effort. In a tighter labor market, it is easier to find a job and the job finding rate $f(\theta_t)$ is higher. Other things equal, this increases the marginal benefit from searching on the job.

Firms and wages. Jobs are owned by competitive firms who value them by the present value of profits, that is, output y_t less wages for a matched job or zero for a vacant job. Firms discount the future at the same rate r as workers. Let $J_t^h(i)$, $J_t^l(i)$ and $V_t(i)$ denote the current value of a matched high or low quality job and a vacant job in city i, respectively. Since there is free entry and no adjustment costs, the value of a vacant job is always

$$V_t(i) = \sigma_t(i) \tag{4}$$

The surplus to the firm from being in a match of quality $d = h, l$ is $J_t^h(i) - V_t(i)$. Let $W_t^d(i)$ denote the utility of a worker at date t in a match of quality d in city i. The worker's surplus from being in the match is then $W_t^d(i) - U_t(i)$, the difference to the utility from unemployment in city i.

We assume that the wage in a match is determined every period by splitting the total mtach surplus from $S_t^d(i)$, defined as the sum of firm and worker surplus, in a fixed ratio, with share β going to the worker and $1 - \beta$ going to the firm:

$$(W_t^d(i) - U_t(i))/\beta = (J_t^d(i) - V_t(i))/(1 - \beta) = S_t^d(i) \tag{5}$$

This rule applies in all periods, including when the worker contemplates an outside job opportunity because on-the-job search has resulted in a match. A worker with an outside opportunity thus compares utility from the old match to expected utility from the new match before knowing the quality draw, taking as given the sharing rule. We note that since duration increases match value, it never makes sense to search on-the-job in a high quality match.

Equilibrium. We start the model at an initial (date 0) distribution of workers across cities as well as a distribution of matches, vacant jobs and unemployed workers within each city. An *equilibrium* consists of stochastic processes that describe, for every date $t > 0$ and city i, workers' and firms' decisions to accept matches, employed workers' search effort, firms' investment in jobs, and wages in every match such that (i) workers' and firms' decisions are optimal given expectations of future match rates and match surplus (ii) the wage in an individual match splits surplus according to (5) (iii) match rates (3) and (2) reflect aggregate vacancies and search effort, and (iv) population evolves according to (1).

We restrict attention to equilibria such that match surplus is always nonnegative - our quantification below sets parameters accordingly. Our model has two convenient properties that facilitate characterization of equilibrium. First, as usual in search models with quasilinear preferences, workers' and firms' decisions, as well as their continuation utilities, depend on individual histories only through their current match status. Wages in a city therefore depend only on the quality of the match: we write $w_t^h(i)$ and $w_t^l(i)$ for the wage in a high and low quality match, respectively.

A second key property is that equilibrium search effort as well as the tightness of the labor market in a city only depend on the exogenous city-specific parameter as well as on the exogenous state of the business cycle, but not on the distribution of workers within or across cities. We can therefore analyze the model as two blocks: we first determine how match rates and wages move over the business cycle, and then how jobs and workers flow between match states and as well as between cities. We derive this property in the next section from agents' decision problems.

3.2 Labor market dynamics within cities

We now study worker and firm decisions within a city. To ease notation, we drop the city index i for this subsection. The equations have the same structure in all cities, although of course the city specific exogenous variables imply different outcomes across cities. We also need notation for small changes in variables that are either smooth functions of time or follow Poisson processes. For any process x_t, we denote by $\mathcal{D}x_t$ its infinitesimal generator, that is, its the instantaneous expected rate of change $\mathcal{D}x_t = \lim_{\Delta\downarrow 0} E_t\left[(x_{t+\Delta} - x_t)/\Delta\right]$. If x_t is a smooth function of time, we have $\mathcal{D}x_t = \dot{x}_t$, the time derivative. For a Poisson process, the generator also includes the expected size of the jump multiplied by the arrival rate of the jump.

Worker and firm decision problems. Consider the evolution of worker utility. The only endogenous state variable for the worker is the individual employment state. Denote utility of an unemployed worker by U_t, and write W_t^h and W_t^l for the utilities of workers employed in high and low quality matches, respectively, all normalized by output y_t. It is also helpful to define the expected value of a match before quality is revealed, we write $W_t = pW_t^h + (1-p)W_t^l$. These values satisfy workers' Hamilton-Jacobi-Bellman (HJB) equations:

$$
\begin{aligned}
(r-g)U_t &= b + f(\theta_t)(W_t - U_t) - \mu U_t + \mathcal{D}U_t, \\
(r-g)W_t^h &= w_t^h - \rho_t(W_t^h - U_t) - \mu W_t^h + \mathcal{D}W_t^h, \\
(r-g)W_t^l &= w_t^l - (\rho_t + \delta)(W_t^l - U_t) - \mu W_t^l \\
&\quad + \max_{e_t}\left\{e_t f_t(W_t - W_t^l) - c(e_t)\right\} + \mathcal{D}W_t^l.
\end{aligned} \tag{6}
$$

An employed worker receives the (normalized) benefit b, finds jobs of unknown quality at rate f_t and moves at rate μ. Employed workers in either type of match receive wages and lose jobs due to either job or match destruction shocks. Workers employed in low quality matches further choose effort. Their private benefit to a worker from searching on the job net of search cost, in braces, is nonnegative because search on the job is voluntary. The first order condition for effort is $c_t'(e_t) = f_t(W_t - W_t^l)$: workers exert effort until the marginal cost of effort is equal to the marginal flow of private benefit to the worker.

Consider next the evolution of firm values. Let V_t, J_t^h, and J_t^l denote the present value of a vacant job, a high-quality matched job, and low-quality matched job, respectively, and let $J_t = pJ_t^h + (1-p)J_t^l$ be the expected value of a match before its quality is realized. These values solve firms' HJB equations:

$$
\begin{aligned}
(r-g)V_t &= \varphi(\theta_t)(J_t - V_t) - \rho_t V_t + \mathcal{D}V_t, \\
(r-g)J_t^h &= 1 - w_t^h - \mu(J_t^h - V_t) - \rho_t J_t^h + \mathcal{D}J_t^h, \\
(r-g)J_t^l &= 1 - w_t^l - (\mu + \delta + e_t f(\theta_t))(J_t^l - V_t) - \rho_t J_t^l + \mathcal{D}J_t^l.
\end{aligned} \tag{7}
$$

Jobs are destroyed at rate ρ_t whether or not they are matched to a worker. Match destruction shocks, worker moves or workers quits all turn a matched job into a vacant job.

To write a simple system for the business cycle dynamics in a city, we define the expected surplus of a new match as $S_t = pS_t^h + (1-p)S_t^l$. We then subtract HJB equations for unmatched values from those for matched values to characterize the joint evolution of surplus and the value of unemployment:

$$\begin{aligned}
(r-g+\rho_t+\mu)S_t^h &= 1-b_t-\left(\beta f(\theta_t)+(1-\beta)\varphi(\theta_t)\right)S_t+\mathcal{D}S_t^h, \\
(r-g+\delta+\rho_t+\mu+(1-\beta)e_t f(\theta_t))S_t^l &= 1-b_t-\left(\beta f(\theta_t)+(1-\beta)\varphi(\theta_t)\right)S_t+\mathcal{D}S_t^l \\
&\quad +e_t f(\theta_t)(S_t-S_t^l)-c(e_t), \\
(r-g+\mu)\,U_t &= b_t+\beta f(\theta_t)S_t+\mathcal{D}U_t, \qquad (8)
\end{aligned}$$

For high quality matches, the right hand side represents flow surplus equal to the benefit of production less the benefit of searching. To compute the value of a high quality match S_t^h, flow surplus is discounted at a rate that reflects destruction or migration shocks that terminate matches. For low quality matches, flow surplus also contains the option value of moving to a better job, less the cost of searching on the job; discounting then takes into account that the match can end when the worker quits.

Two additional equilibrium conditions reflect optimal effort and investment (or job creation). We obtain a first order condition for effort from the HJB equation for workers in low quality matches in (6):

$$c'(e_t) = \beta f(\theta_t)(S_t - S_t^l), \qquad (9)$$

Workers in low quality matches optimally equate the marginal cost of effort on the left hand side to its marginal benefit (on the right). Marginal benefit if higher the larger the worker's share in wage bargaining, the easier it is to find a job (that is, the tighter is the labor market, since f is increasing). It also gros with the difference between expected surplus and surplus in a low quality match, or equivalently with the difference in surplus between the two match qualities.

We obtain a job creation condition by substituting the free entry condition (4) into the HJB equation for vacant jobs in (7):

$$(r+\rho_t)\sigma_t - \mathcal{D}\sigma_t = (1-\beta)\varphi(\theta_t)S_t, \qquad (10)$$

Firms optimally equate the flow cost of job creation on the left hand side to the present value of profits from the jobs on the right. The flow cost of job creation amortizes the initial investment in the job σ_t, using a discount rate that takes job destruction ρ_t into account. It thus varies both across cities and over the business cycle: job creation is more costly in high σ_t cities, and it becomes more costly in recession when the duration of jobs shortens (ρ_t increases). We note that this countercyclical variation in job creation costs distinguishes our model from standard search models where the flow cost of job creation is a constant vacancy cost.

Equations (8)-(10) represent a system of 5 equations in 5 variables: search effort e_t, tightness θ_t as well as the 3 values S^h, S^l and U. Conjecture a solution such that all variables depend only on the exogenous state of the business cycle. For example, the value when unemployed takes only two values $U_t \in \{U^b, U^r\}$. It jumps if and only if the state changes, so the generator terms take a simple form. For example, the expected change in utility in the boom state b is $\mathcal{D}U_t = \pi_{br}(U^r - U^b)$. Expressions for other values and states similarly depend only on the intensity in the current state as well as values in both states. The system thus boils down to solving 10 equations in 10 unknowns.

We have shown that market tightness and search effort, together with surplus and all values, can be solved as functions of the business cycle states and parameters, independently of the initial distributions of the labor force across cities and workers and jobs in different states within cities. Intuitively, this follows from our standard assumptions of constant returns technologies for production and matching, linear utility, and free mobility of capital across locations. A scrambling of the distribution of workers, say, leads to shifts in the distributions of vacancies and matches, but ratios relative to the labor force are unchanged.

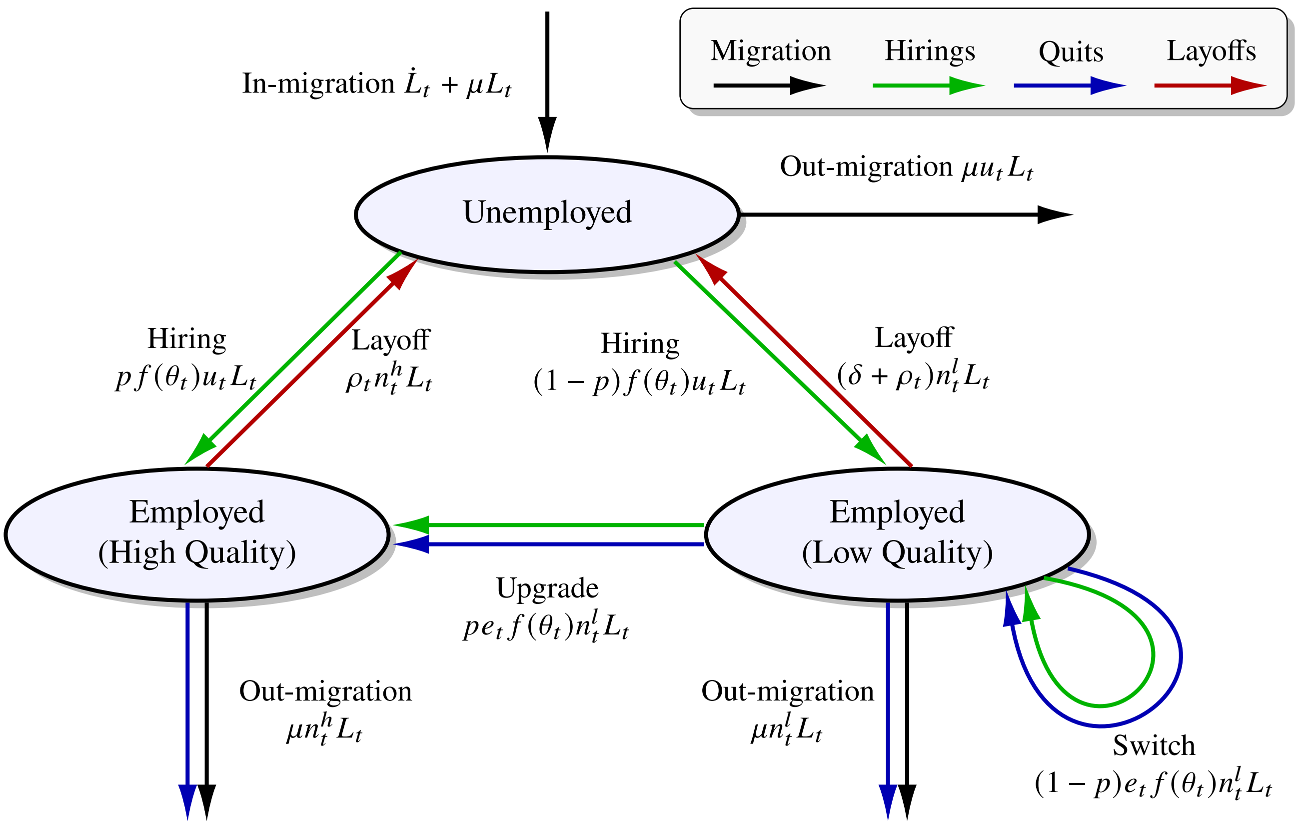

Figure 9: Worker flows within cities

3.3 Labor market flows

Figure 9 illustrates worker flows in the model. Within any city, workers are in one of three states: unemployed or employed in high or low quality matches. Let u_t, n_t^h and n_t^l denote their respective shares of the city's labor force L_t that add up to one, and let

$n_t = n_t^h + n_t^l$ denote the employment rate. Black arrows in the figure represent migration. Since all in-migrants start unemployed, gross in-migration is displayed at the very top. It includes $\dot{L}_t$ net in-migrants determined by the migration equation (1). Out-migration is due to migration shocks that hit all types of workers, so it adds up to the total μL_t.

We define *quits* as separations that occur when an employed worker receives a moving shock or finds a new job due to on-the-job search. The quit rate is the sum of the moving rate and the rate at which on-the-job searchers quit:

$$q_t = \mu + e_t f(\theta_t) \frac{n_t^l}{n_t}. \tag{11}$$

The second term is the job finding rate of OTJ searchers multiplied by their population share, or the share of workers in low quality matches. Quits are represented by blue arrows in Figure 9. Out-migration by employed workers is thus marked by two arrows to indicate the overlap of quits and migration. Since the match quality is not known in advance, only a share p of OTJ searchers who quit actually upgrade to a high quality match – a share $1 - p$ simply switches between low quality matches. Both flows are also marked in green representing new hires.

We define *layoffs* as separations that occur because of exogenous job or match destruction shocks. A laid off worker then flows to unemployment. The layoff rate is

$$\lambda_t = \rho_t + \delta \frac{n_t^l}{n_t} \tag{12}$$

Layoffs are represented by red arrows in the figure. They vary over time as well as across cities for two reasons. The first is variation in the job destruction rate ρ_t: more jobs are destroyed in recessions. In addition, match destruction shocks hit low quality matches, so layoff rates are higher when the endogenous share of low quality matches is larger. This share can again be driven by the business cycle or by city-specific parameters.

The evolution of the worker distribution in a city is described by

$$\begin{aligned}
\dot{u}_t + u_t \dot{L}_t / L_t &= \dot{L}_t / L_t + (1 - u_t)(\mu + \lambda_t) - u_t f(\theta_t), \\
\dot{n}_t^h + n_t^h \dot{L}_t / L_t &= p f(\theta_t) u_t + p e_t f(\theta_t) n_t^l - (\rho_t + \mu) n_t^h, \\
\dot{n}_t^l + n_t^l \dot{L}_t / L_t &= (1 - p) f(\theta_t) u_t - p e_t f(\theta_t) n_t^l - (\rho_t + \delta + \mu) n_t^l.
\end{aligned} \tag{13}$$

Here the left hand side of each equation is the number of workers of a given status normalized by total population, for example $(d(u_t L_t)/dt)/L_t = \dot{u}_t + \dot{L}_t / L_t$. The right hand side adds up over the different shown in Figure 9.

The number of unemployed workers grows because of migration and shrinks because of job finding. Both net and gross migration matter, because all migrants start out unemployed: in addition to the net inflow of new workers, captured by $\dot{L}/L$, purely gross migration also moves workers who were previously employed to unemployment. Numbers of matched workers of either quality grow because of new matches with unemployed workers and shrink because of separations due to destruction shocks or migration. In addition, the number of high (low) quality matches further increase (decrease) by poached workers who quit and actually get a higher quality match.

Summing up, the equilibrium can be characterized in two steps. We first solve for business cycle dynamics in every city. We obtain Markov chains for tightness, search effort and surplus. Job finding, quit and layoff rates are simple functions of tightness and search effort given by (3), (11) and (12), respectively. Second, we solve a system of equations for flows that consists of the job flow equations (13) as well as the migration equation (1) for every city, with city-specific parameters as well as job finding, quit and layoff rates from the first step. While the system of flows is large and interdependent, it is easy to solve because it does not involve expectations.

3.4 Comovement in labor market flows

In this section, we describe the core mechanism that allows our model to fit comovement of many variables in both the cross section of cities as well as over the business cycle. In a nutshell, growing cities and good times are special in that job creation makes the labor market tighter, so it is easier for workers to find jobs. A tighter market encourages both in-migration and effort to search on the job. When more workers quit unstable low-quality matches to find more stable ones, layoffs and unemployment are lower. These implications are neither qualitatively nor quantitatively obvious, and depend on several special features of our model, in particular the role of jobs as a capital good, heterogeneity in match stability and endogeneous search effort.

We provide intuition for the mechanism in two steps, in line with the solution strategy outlined above. We first focus on flows within a city given tightness and explain why changes in tightness generate the right comovement. We then explain why job creation costs leads to large changes in tightness across cities. Here we contrast our model to the standard DMP model where fluctuations in job destruction rates are typically not used as a driver of the cycle since they imply an upward-sloping Beveridge curve. Finally, we comment on the role of two assumptions, that match quality is identified with match duration, and that bargaining simply splits match surplus.

Stocks and flows in steady state. To understand variation in job flows, it is helpful to initially abstract from migration, growth and the business cycle, and study comparative statics for a stationary model of a single city. Much of the intuition from this exercise carries over to our quantitative results below. We thus set $g = \mu = \nu = 0$ and $\rho_t = \rho$ and solve for steady state values from the flow equations (13). For given constant tightness and search effort, the steady state share of low quality matches in a city is

$$\frac{n^l}{n} = (1-p)\frac{\rho}{\rho+\delta+pef(\theta)} < 1-p \tag{14}$$

It is always below $1-p$, the share of *new* low quality matches because workers leave low (but not high) quality matches even before the job is destroyed: low quality matches break at rate δ and workers quit to better matches at the rate $pef(\theta)$.

Tighter markets generate a lower share of low quality matches, or higher average match quality. The same is true for markets with higher search effort on the job. Intuitively, the easier it is to find a job for an OTJ searcher, the faster they leave low quality matches for higher quality ones. The share of low quality matches is also lower when fewer low quality matches are created to begin with (higher p) or when low quality matches break at a faster rate (higher δ). It is higher when jobs are destroyed at a higher rate: higher ρ means that match quality is drawn anew more often, pushing up the share towards $1-p$ as ρ becomes large.

Tightness and search effort generate comovement. It follows that tighter markets, as well as markets with more search effort, see fewer layoffs, more quits, and lower unemployment:

$$\lambda = \rho + \delta \frac{n^l}{n}, \qquad q = ef(\theta)\frac{n^l}{n}, \qquad u = \frac{\lambda}{\lambda + f(\theta)} \tag{15}$$

Layoffs fall with tightness and search effort because match quality captures risk: they reflect not only job destruction, but also destruction of low quality matches, which are less prominent in tighter markets. Quits rise with tightness and search effort even though there are fewer low quality matches, as the direct effect on the job finding rate dominates. Finally, the unemployment rate falls with tightness and search effort as fewer laid-off workers flow into unemployment. In addition more searchers flow out to new jobs in a tighter market.

Equations (15) are at the heart of our model predictions for comovement in both the cross section and the time series. Below we compare different exogenous fundamentals that drive tightness and search effort. In the cross section, a lower job creation cost σ increases both variables. In the time series, a lower job destruction rate ρ does the same. The job desctruction rate also appears separately in (14)-(15): it directly lowers layoffs and the share of low quality matches. While fewer low quality matches push quits down, this effect is quantitatively small so variation in job destructions also leads to the right comovement.

Vacancies and Beveridge curves. As a final moment, consider the vacancy rate

$$v = \frac{\lambda}{\varphi(\theta)} \tag{16}$$

Here, the effect of tightness is generally ambiguous, as two opposing forces are in play. On the one hand, tighter markets imply higher match quality, and hence lower layoff rates. As fewer jobs open up every period, the number of vacancies declines. On the other hand, a tighter labor market makes it harder to fill vacancies, so more vacancies remain open. Since unemployment always fall with tightness, it is thus possible for the Beveridge curve to slope either up or down, depending on parameters.

It is desirable for both time series and cross sectionl Beveridge curves to slope down: the cross sectional fact is Figure 6, whereas a large literature discusses the time series fact. We therefore would like unemployment and vacancy rates to move in opposite directions both in response to changes in tightness and search effort generated by job creation costs, and in response on changes in the job destruction rate. In either case, we need the negative effect on the vacancy rate through the layoff rate to be smaller then the positive effect through the job filling rate. Quantitatively, this works in our model if tightness moves a lot with the exogenous fundamental compared to both search effort and the share of low quality matches. We turn next to the response of tightness to shocks.

Job creation, job destruction and tightness. To clarify the link between investment in jobs and tightness, we continue with our stationary model of a single city. To allow comparison to the DMP model, we further abstract from on-the-job search by assuming that all matches are unstable, so $p = 0$, but instead introduce a flow cost κ of posting vacancies. The dynamics of surplus (8) then reduces to a single equation that expresses match surplus as the present value of flow benefits. Together with the job creation condition (10), we obtain two familiar schedules that relate surplus S and tightness θ, both constant in equilibrium:

$$\underbrace{\frac{(r+\rho)\sigma+\kappa}{(1-\beta)\varphi(\theta)}}_{\text{Job Creation}} = S = \underbrace{\frac{1-b-(r+\rho)\sigma}{r+\delta+\rho+\beta f(\theta)}}_{\text{Present Value}}. \tag{17}$$

The DMP model is the special case where $\kappa > 0$ and $\sigma = 0$, whereas the version of our model without search on-the-job is $\kappa = 0$ and $\sigma > 0$.

Figure 10 displays both schedules as solid lines in S-θ-plane, drawn identically for the DMP model in the left panel, and for our model in the right panel. Job creation schedules are downward sloping: as the market tightens, it is harder for firms to fill a job, so more match surplus is required for them recoup the flow cost of job creation. This logic is the same in DMP, where job creation requires only a vacancy cost, as in our model when it requires only a flow cost of investment. The schedules are identical as long as the vacancy cost is equal to the flow cost of investment.

Present value schedules are upward sloping: as the market becomes tighter, the benefit of being currently matched falls. As the job finding rate becomes large, we approach a Walrasian market where there is no match surplus since all workers can be rehired immediately. What is different in our model that the flow surplus from a match is net of $(r+\rho)\sigma$, the flow cost of investment. The reason is that this cost accrues whether or not the job is filled, in contrast to a vacancy cost that is saved when the job is filled. In order for the schedules to be identical, we thus need that the parameter b is lower in our model to obtain the same flow surplus.

Compare now the response of the two models to an increase in the job destruction rate, our exogenous business cycle shock. The key point here is that when jobs are a capital good, tightness falls (and unemployment rises) much more than in the standard DMP model. The comparative statics for tightness are displayed as dashed lines in Figure 10. In the DMP model (left panel), only the present value schedule shifts left as benefits are discounted at higher rate. The market slackens in the new equilibrium. Qualitatively, the DMP model thus generates the right comovement: lower tightness, lower layoffs and higher unemployment.

In our model (right panel), the same effect is present. However, it is now reinforced by an increase in the flow cost of investment, which shifts both schedules. First, the present value schedule shifts left further to reflect lower flow benefits. In addition, the job creation schedule also shifts left as investors are compensated for higher cost. As the labor market slackens, the unemployment rate $u = \lambda/(\lambda + f(\theta)$ rises not only because of higher λ but also because of a lower job finding rate. The resulting recession is thus deeper than in the DMP case, as the higher cost of investment discourages firms from creating jobs.

The decline in tightness in our model can be sufficiently strong to lower vacancies, so the economy moves into recession along a downward sloping Beveridge curve. In particular, this happens for typical calibrations where the DMP model predicts instead that vacancies increase with the job destruction rate. The vacancy rate takes the same form as in (17), so its response is

$$\frac{dv}{d\rho} = \frac{1}{\varphi(\theta)}\left(1 - v\varphi'(\theta)\frac{d\theta}{d\rho}\right) \tag{18}$$

We would like tightness to fall enough with the job destruction rate in order for the vacancy rate to fall.

To get a feel for the numbers, consider a typical calibration for the US. We first choose separation rates and a matching function such that steady state flow equations generate a lot of "churn", as in the data. Let $\delta = .14$ and $\rho = .1$ and target $u = .04$ and $v = .03$ by choosing $f(x) = 6.7\sqrt{x}$. It follows that $\theta = 0.72$, the job finding rate is $f(\theta) = 5.76$ and the job filling rate is $\varphi(\theta) = 8$. For a standard interest rate $r = .03$ and bargaining weight $\beta = .5$, we can now compare two parametrizations that generate identical schedules (17).

The DMP parametrization targets $b = .6$, which implies $\kappa = .48$. For our capital good parametrization, we pick $\sigma = 2.8$ to get the same flow of cost of job creation and adjust $b = .12$ to get the same flow benefit. With these numbers, we have $\varphi'(\theta) = 5.6$ and $d\theta/d\rho$ must be less than -6 for the derivative (18) to be negative. In the DMP calibration, we have instead $d\theta/d\rho = -0.54$, a much smaller number that implies an upward sloping Beveridge curve. Under our alternative calibration, in contrast, $d\theta/d\rho = -6.9$, so the Beveridge curve slopes down.

We can also use the simple model to clarify how variation in job creation costs generates cross-sectional differences in tightness. An increase in the cost σ has the same qualitative effects on the two schedules as drawn in the right panel of Figure 10. Like a higher job destruction rate, the cost σ increases the flow cost of investment and discourages job creation. Mechanically, the change in σ generates exactly the two shifts that distinguish our model from the DMP model, which also worked through the cost of investment. The small effect of ρ through discounting has no counterpart when job creation cost varies, consistent with the fact that the parameter σ does not exist in the DMP model.

Migration. The effect of tightness on migration and hence employment growth cannot be studied in the simple model of a single city. We can, however, obtain intuition on what drives the utility of unemployed workers, the key input to the migration equation (1). In steady state, the value of an unemployed worker from (8) simplifies to

$$rU = b + \beta f(\theta) S$$

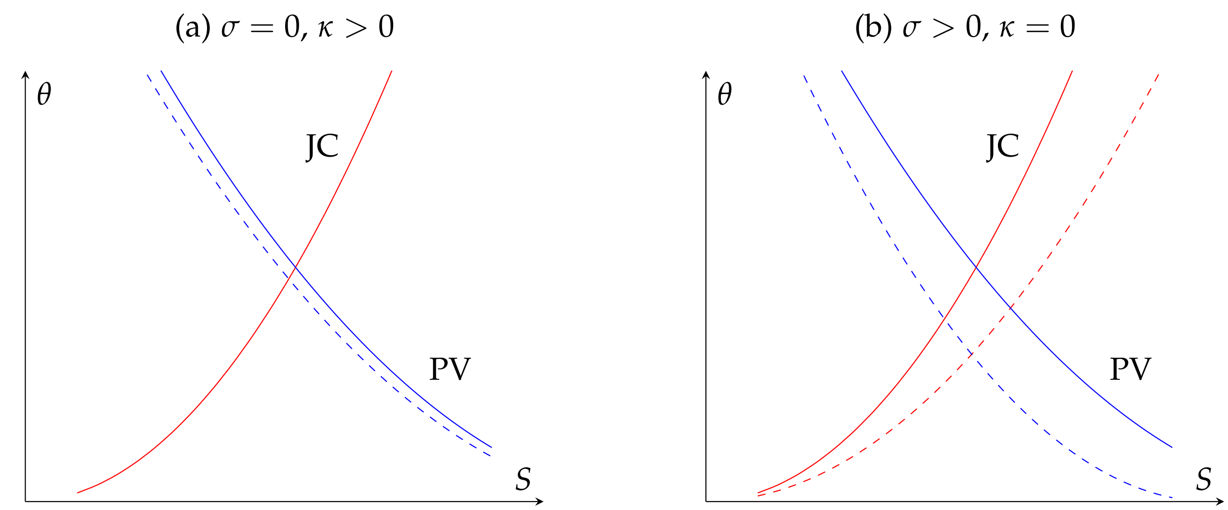

Figure 10: Comparative statics effect of a rise in job destruction rate ρ

Notes: The figure illustrates how an increase in the job destruction rate affects the labor market when job are free but job posting is costly (panel (a)) and when jobs are capital goods (panel (b)). In both panels, the horizontal axis is the surplus of a match, and the vertical axis is tightness. The solid job creation curve (JC) and present value curve (PV) are calibrated to match in two models. The dashed line represents the outcome of an increase in the job destruction rate. In case jobs are free to create, this change reduces the duration of the match and the present value but does not affect the job creation condition directly, resulting in a small decline in surplus and tightness. When jobs are capital goods, as in our model, the change decreases the present value more due to the outside option of the firm and also shifts the job creation curve down, resulting in a larger drop in tightness and relative stability of the match surplus.

As we have seen, surplus S increases together with tightness in response to shifts in either job destruction or job creation cost. Since the job finding rate is increasing in θ, a city becomes more attractive to unemployed workers when its labor market is tighter. Workers thus flow to the city, and we obtain the right comovement for employment growth in addition to quits, layoffs and unemployment.

4 Quantitative implications

In this section, we quantify the model to fit our panel of 64 cities. The goal is to provide a parsimonious account of worker flows over time and across cities based on the variation in labor market tightness, as described in Section 3.4. To this end, we assume that preferences and technology are the same across cities, with one exception: every city has its own constant job creation cost. Additionally, cities differ in their labor forces at the beginning of our sample. All cities then experience a common business cycle, driven only by fluctuations in the job destruction rate ρ_t, with the timing of the cycle taken from the data. We simulate the evolution of cities, as well as workers and jobs within cities, and compute the same moments as in Section 2.

4.1 Parametrization

We choose standard isoelastic functional forms for the matching technology and the cost of on-the-job search. In particular, the job finding rate is $f(\theta) = \bar{f}\theta^{\alpha}$, where $\alpha \in (0,1)$ is the elasticity of the matching function with respect to vacancies. Moreover, let $c(e) = \bar{c}e^{1+1/\eta}$, where η is the elasticity of search effort with respect to the payoff of search on the job. We then divide the parameters into two groups. One group is set up front, either to standard values from the literature or to match particular targets. Another group is chosen to fit a large set of target moments jointly.

Preset parameters. Preset parameters are listed in Table 1. We follow Shimer (2005) in selecting a discount rate $r = .048$ per year, the elasticity of the matching function is $\alpha = .72$, the bargaining weight of workers is $\beta = .72$, and the utility flow from unemployment is $b = .4$. The rate of productivity growth g is .023 per year, the rate of growth of labor productivity in manufacturing over the sample period.[7] The arrival rate of migration shocks μ, equal to the model-implied gross migration rate in the economy, is .06 per year, the higher rate of internal migration measured by Saks and Wozniak (2011).

The job destruction shock process is a continuous time Markov chain that switches between boom and recession states. It is parametrized by two values for boom and recession as well as two Poisson arrival rates, for a switch from boom to recession and vice versa. We set the arrival rates up front to match mean durations of business cycle phases from the business cycle classification introduced in Section 2. The mean duration of a boom is 40.6 months, and the mean duration of a recession is 10 months. We thus set the arrival rate of a regime switch to .295 for a boom and 1.20 for a recession.

Moment matching. Moment matching recovers 64 city-specific costs of job creation $\sigma(i)$ as well as nine parameters that are common to all cities, listed in Table 2. In addition to the rates of job destruction in booms and recessions, ρ^{b} and ρ^{r}, respectively, common parameters include the match destruction rate of low quality matches δ, the scale of the matching function, $\bar{f}$, the probability of a high quality match, p; the scale and elasticity parameters of the search cost, $\bar{c}$ and η, the natural rate of labor force growth, ν, and the elasticity of migration, ϵ.

We target 6 city-specific moments constructed from LTS and BEA data for 64 cities. For each city, we target the mean employment growth for the entire sample. We also target the mean layoff rates and quit rates, calculated separately for boom and recession months, and the mean unemployment rate between 1977 and 1979, the years for which we have data by city. We arrive at $64 \times 6 = 386$ city-specific moments. In addition, we target two numbers taken from other studies. One is the mean aggregate vacancy rate of 1.6% for the period between August 1969 and December 1973, as reported in Abraham (1983, Table 3, "Mean vacancy rate, JOLTS Program"). The second is an aggregate share of employed workers searching on-the-job over our sample of 22.4%. Here we rely on recent work by Faberman, Mueller, Şahin, and Topa (2022, Table 1, "Percent that actively searched for work, Employed").

[7] See "Historical labor productivity SIC measures for manufacturing sectors, 1949-2003", available on the website of the Bureau of Labor Statistics.

The calculation of simulated moments takes into account the transition dynamics of the model over the sample, as well as the precise timing of of the data used to form each empirical moments. We set the initial labor force of each city to its manufacturing employment in 1969. To determine an initial distribution of endogenous labor market states within cities, we solve (13) separately for every city for a constant employment growth rate and assuming that the economy is always in a boom, where the growth rate is calculated using the employment shares in January 1969. We then follow the two-step procedure outlined in Section 3: we solve for indirect utility, tightness, and search effort on the job for each city and for each business cycle state, and then solve the path of the worker distribution and flows in (13) over the entire sample.

With a model-generated worker distribution and flows, model moments can be constructed with the same dating as the data moments. In particular, we derive an empirical mean vacancy rate as the mean over the period from April 1969 to December 1973. Similarly, the calculation of the mean unemployment rate for each city uses values from March of 1977, 1978, and 1979 only, following the construction of these moments in the data sources (the unemployment rates reported by the BLS were based on the Current Population Survey).

Overall, the exercise is heavily overidentified. For a panel of N cities, the matching procedure targets 6N+2 moments using $N + 9$ degrees of freedom. In our case, $N = 64$ so we are fitting 388 moments with 73 parameters. City-specific moments are weighted by the inverse of their cross-sectional standard deviation. The two aggregate moments receive a weight of N divided by their respective mean values. Increasing or decreasing the weight on the outside moments by a factor of 10 does not significantly change the results qualitatively.

Identification. Identification comes jointly from trends and the business cycle, both on aggregate and in the cross-section of cities. Since cities differ in only one parameter, the quantified model can effectively capture one "nonlinear principal component" of the cross-section of city-specific moments. Hence, the city-specific parameters are chosen to place the cities along this single component, and the common parameters are chosen to fit (i) aggregate moments and (ii) the comovement of moments in the cross-section.

To see how different degrees of freedom are used to match the moments, consider briefly the model of a single city in steady state from Section 3.4 and think of it as a model of the US as a whole, informed by aggregate moments. Observations of the share of OTJ searchers n^l/n, the unemployment rate u, the vacancy rate v, the layoff rate λ, and the quit rate q pin down tightness θ, the scale of the matching function $\bar{f}$, and the search effort on the job e (based on Equations (14), (15), and (16)). Out of the parameters δ, ρ, and p, only two are needed to match the remaining moments, and so one degree of freedom remains. In addition, the optimality conditions (8)-(10) can only identify one of the remaining parameters: $\bar{c}$, η, and the job creation cost of the hypothetical city.

Our exercise thus must infer the common parameters in part from cross-sectional and business cycle variation. For example, the elasticity η determines how strongly search effort moves with tightness across cities and hence regulates the comovements of quits and unemployment, and also determines how much quits vary between boom and recession in the aggregate. At the same time, from equations (15), the probability of an unstable match p and the match destruction rate δ determine how much the layoff rate moves against the quit rate across cities. Similarly, comparing business cycle and cross-sectional variation in quits and layoffs helps the exercise gauge not only the size of fluctuations in destruction rates, but also the relative importance of job and match destruction. Finally, the migration elasticity ε regulates how strongly city growth responds to labor market conditions, and moves it away from the benchmark of aggregate population growth driven by ν.

Parameter values. The common parameter values are presented in Table 1. Consider first the two key sources of variation. The mean cost of creating a new job is 1.86 times the annual output of a worker, with a standard deviation is .37. The job destruction rate is .18 in booms, and jumps to .23 in a recession. The flow cost of investment in new jobs $(r+\rho)\sigma$, therefore, varies substantially for both reasons. From an investment perspective, moving into a recession in the average city increases cost by the same amount as moving to a city with a job creation cost that is higher by 1.1 standard deviations during a boom. While the change in the flow cost of investment is not the only way in which job destruction rates affect worker flows, as discussed in Section 3.4, it captures the variation in job creation at the heart of our mechanism.

The cost of search on the job to a worker who makes as much effort as an unemployed worker is 11% of match output. The elasticity of OTJ search is 1.55, that is, if the expected benefit of search increases by 1%, on-the-job searchers exert 1.55% more effort. The scale of the matching function is 7.45 and the probability of getting a stable match is 0.39. This means that it takes 3 moves, on average, until a worker ends up in a stable match. The elasticity of migration is 2.2, which means that a city that provides 1% more utility attracts 2.2% more in-migration. Finally, growth of the overall labor force is -0.006, reflecting the slow decline of national employment in manufacturing.

Table 1: Preset parameters

Parameter	Value	Source
r	.048	Shimer (2005)
α	.72	
β	.72	
b	.4	
g	.023	BLS
μ	.06	Saks and Wozniak (2011)
b→r	.295	boom/recession duration
r→b	1.20	

Model fit. The model closely fits the two aggregate moments. The mean vacancy rate in the quantified model is 1.74%, slightly above the target of 1.6%, and a share of 22.7% of workers are searching on the job, compared to 22.4% in the target moment. The fit for each of the six city-specific moments is illustrated in Figure 11. The model can capture the heterogeneity in growth rates and the relationship between growth rates to the mean quit and layoff rates in booms and recessions, as well as the unemployment

Table 2: Jointly-estimated common parameters

Parameter	Value	Interpretation
ρ^{b}	.182	job destruction rate during booms
ρ^{r}	.233	job destruction rate during recessions
$\bar{c}$	.103	scale of search on the job
η	1.55	elasticity of search on the job
$\bar{f}$	7.45	scale of matching function
p	.386	probability of high-quality match
δ	.257	rate of match destruction, unstable match
ϵ	2.20	elasticity of migration
ν	-.006	natural growth rate of labor force

rate. The model overstates, on average, the changes in mean quit rate between boom and recession, and understates the changes in mean layoff rate between boom and recession.

Since there is only one city-specific parameter, an interesting check on the model is whether it can indeed capture the first linear principal component of the 6 city-specific data moments, which explains 57% of the variation across cities. Figure 12 shows the fit of the first principal component scores from the data compared to the moment generated scores (using the PCA coefficients from the data). The model is able to match the first principal component remarkably well: the correlation between the model-generated and the data scores is 0.98. In other words, the model captures a significant share of the total variation in the data with only one free parameter per city.

4.2 Quantifying the mechanisms

We now use counterfactuals to quantify the contribution of different mechanisms discussed qualitatively above.

Job creation costs and Heterogeneous growth rates. The cross sectional differences in employment growth rate arise from differences in the cost of job creation. Cities with a low job creation cost attract more jobs, which makes the labor market tighter, and in turn attracts more workers to those cities. The new workers add to the labor force, making it even more attractive to firms and so on. This amplification mechanism means that small differences in job market outcomes on the worker side can have a large impact on the trajectory of a city.

To illustrate, consider again our leading example cities, Buffalo, NY and Salt Lake City, UT. Buffalo shrinks by 2.3% per year, whereas Salt Lake City grows at 4.3% per year. In the model, out-migration from both cities is 6% per year, so in-migration to Buffalo is only 3.7% per year, leading to net shrinkage. In contrast, in-migration to Salt Lake City is 10.3% per year. The model derives this gap from a $(10.3\% - 3.7\%)/\varepsilon = 3\%$ difference in the expected consumption equivalent utility. Our quantitative exercise thus infers that small differences in fundamentals can generate large divergence in growth when factors of production are mobile.

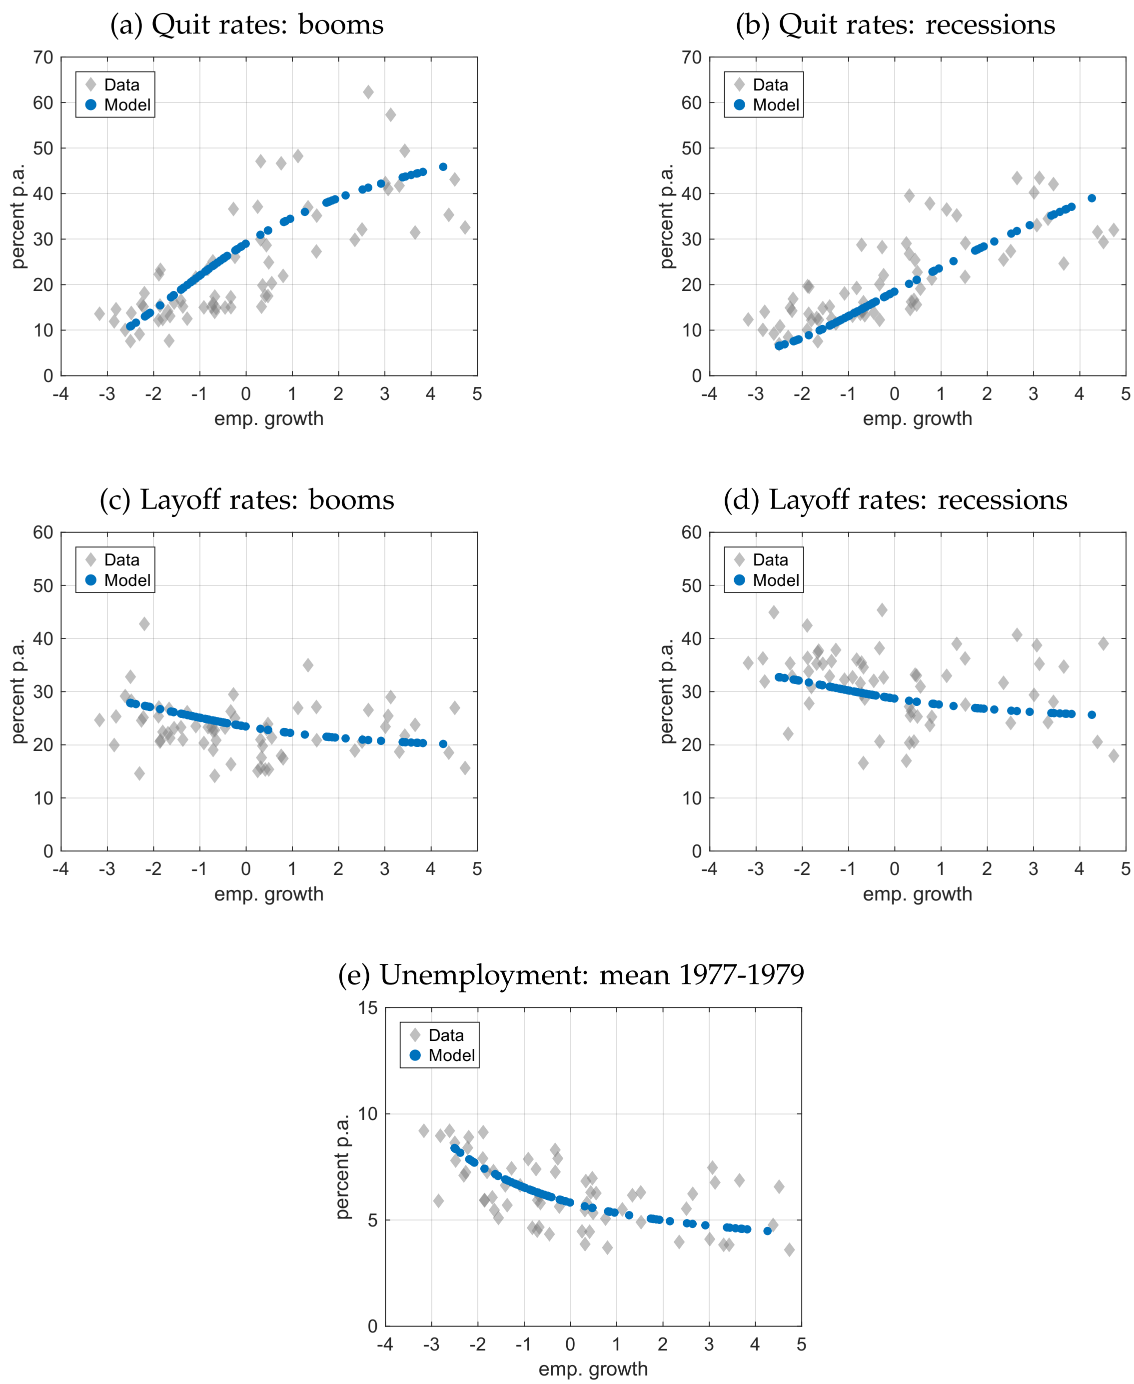

Figure 11: Fit to city-specific moments

Notes: The figure shows the model fit to the cross-sectional joint distribution of mean employment growth in manufacturing and the mean quit rate in booms (panel (a)), mean quit rate in recessions (panel (b)), mean layoff rate in booms (panel (c)), mean layoff rate in recessions (panel (d)) and mean unemployment rate in 1977-1979 (panel (e)). Each dot represents a city. Grey diamond dots are from the data and blue circle dots are model simulated.

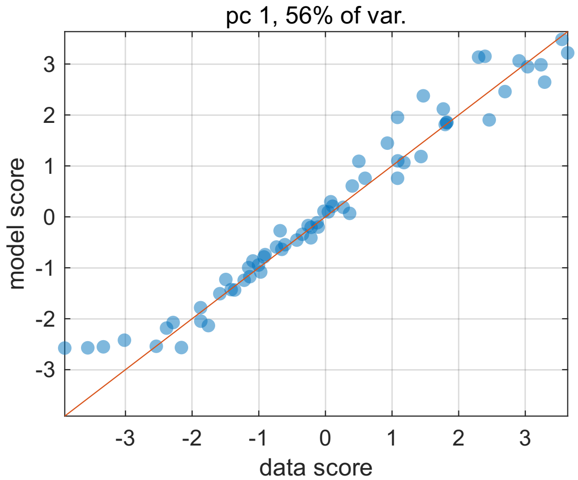

Figure 12: Model fit of first principal component

Notes: The figure shows the model fit to the first principal component (pc) in the data. Each dot represents a city. The horizontal axis is the principal-component score of the city in the data and the vertical axis is the principal-component of the city in the simulated model (using data weights). The red line shows the 45 degrees line.

Quit rates in the cross section. From the second equation in (15), three forces matter for the relationship between employment growth and the mean quit rate. First, a tighter market increases the job finding rate $f(\theta)$ for all searchers, including those who search on the job. Second, higher effort makes the search on the job more successful and further increases the quit rate. Finally, the quit rate is higher if a higher share of workers is in unstable matches and therefore search on the job for a better match.

Taking logs of the quit rate equation (net of migration quits), we obtain an additive decomposition of the three forces. Figure 13 shows how log quit rate in purple varies with mean employment growth, together with two counterfactual lines. The blue line isolates the share of employed workers searching on the job: it declines with growth rate as fewer workers are stuck in bad jobs. The green line adds to the blue line the job-finding rate due to the greater tightness. Cities that grow faster have more vacancies per searcher, and so it is easier to search in them. This effect is about the same size as the decline in the share searching on the job, so the line is relatively flat. The difference between the purple and green lines thus isolates the contribution of variation in search effort, which the model infers as being important for matching higher quit rates in faster-growing cities.

Layoff rates in the cross section. Layoff rates vary across cities with the distribution of match qualities: in cities where workers quit less frequently, they are more likely to remain in a low quality match. Figure 14 displays the contributions of job and match destruction to the overall layoff rate. The gap in layoff rates between high and low quality matches is significant. In booms, for example, workers in high quality matches are laid off roughly once every 7 years (compared to 5 years in a recession), while workers in low quality match are laid off every 2 years in booms (1.75 years in recession). The model predicts that workers in low growth cities are 3 times more likely to be in an unstable

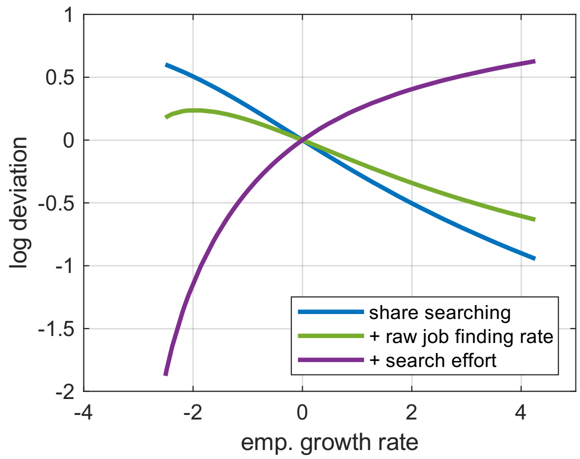

Figure 13: Decomposition of mean quit rates

Notes. The figure shows a decomposition of the mean log quit rate in the simulated model into its three components, relative to the mean quit rate of a simulated city with mean growth zero. The blue line shows the log deviation of the share of searching employed workers which is lower in faster growing cities because workers are better matched. The green line adds the log deviation of job find rates with constant effort. Job finding rates are higher in faster growing cities, but not enough to generate more quits. Finally, the purple line add the mean search effort. The estimated search effort needs to be significantly higher in faster growing cities to capture the cross-sectional relationship between growth and quits.

match.

Aggregate volatility of unemployment. The model is able to capture large aggregate fluctuations in labor market tightness between booms and recessions through changes in the rate of job destruction alone. Taking the total number of vacancies divided by the total number of unemployed in the simulated model and removing a trend (hp-filter with parameter 14400), the standard deviation of the logarithm of the vacancy to unemployed ratio is 0.356, and hence close to 0.382, the value originally calculated by Shimer (2005) for aggregate employment in the US.

Beveridge curves over time and across space. The model further replicates downward sloping Beveridge curves. Figure 15 displays model-implied Beveridge curve calculated separately for boom and recession months. Each dot represents a city. Blue dots are located at mean vacancy and unemployment rate over all boom periods, while orange triangles ate at means over all recession periods. Both cross-sectional Beveridge curves are downward sloping. Moreover, going into a recession the Beveridge curve shifts down and to the right.

Each individual city thus moves along a downward sloping Beveridge in the time series, but those Beveridge curves have quite different slopes. To illustrate, we connect the dots for our example cities Salt Lake City and Buffalo. A recession in Salt Lake City, highlighted near the top of the curves, comes with a modest increase unemployment and a considerable decline in the vacancy rate. When Buffalo goes into recession, in contrast, it experiences a smaller drop in the vacancy rate (in absolute terms) together with a large rise in the unemployment rate.

4.3 Interpreting city-specific job creation costs

A key result of our analysis is that a single city-specific parameter—the cost of job creation—accounts for the observed cross-sectional variation in employment growth, quits, layoffs, and unemployment rates. This parameter is identified structurally within our spatial labor market model, but it is not directly observable in the data. A natural question is therefore: where do these differences in job creation costs come from? Are they the result of differences in local economic fundamentals, such as agglomeration patterns across industries, or are they the result of institutional or regulatory frictions that raise the effective cost of starting a new plant or expanding an existing one? Similar questions have been posed in the literature on the long-run decline of the U.S. Rust Belt, where researchers have debated whether industry shocks, trade exposure, or labor market institutions were the main drivers (Alder, Lagakos, and Ohanian, 2023; Autor, Dorn, and Hanson, 2013, 2016).

We investigate these possibilities using two sources of observed variation, presented in Figure 16. First, we construct a Bartik (1991)-style predictor of local employment growth based on the 1970 industry composition of each city and national manufacturing-industries employment growth rates over 1969–1981.[8] This variable captures the amount of employment growth a city would have experienced if employment within each city-industry pair had grown at the national growth rate of that industry. If differences in job creation cost across cities stand in for differences in industry composition, we should see a strong negative correlation between the Bartik predictor and observed employment growth. Figure 16 Panel (a) shows that, in practice, the correlation between the Bartik predictor and estimated city-specific job creation cost is weak. Hence, differences in industry composition cannot account for regional differences in employment dynamics.

Other studies have emphasized the role of local labor market institutions, such as unions, in determining city-level outcomes. We calculate union coverage in manufacturing for the cities in the sample using CPS sample from October 1985 to December 1989 as a proxy for institutional frictions.[9] Figure 16 Panel (b) shows a scatter plot of job creation cost and city union coverage rates for 59 cities. In contrast to the industry composition, here we find a remarkably positive relationship between the estimated job creation costs and union coverage rates, with a cross-sectional correlation of .77. This result resonates with micro evidence that unionization increases establishment closure risk and reduces employment growth in manufacturing (Lee and Mas, 2012; Wang and Young, 2024), and with macro-level accounts linking labor market conflict to the slow growth of Rust Belt cities (Alder et al., 2023).

Taken together, these results suggest that institutional factors, rather than economic fundamentals, are the primary drivers of cross-city heterogeneity in job creation cost. Our interpretation is that cities with more unionized labor markets and potentially more stringent regulations faced higher barriers to expanding employment—whether

[8] We measure employment growth at the 2-digit manufacturing industries (21 industries) from NIPA Table 6.4B.

[9] The calculation follows the methods proposed by Hirsch and Macpherson (2003). October 1985 is the first month in which MSA identifiers are available for the majority of cities. We pool the data to December 1989 to get enough observations for a reliable estimate, assuming that union coverage is highly persistent.

through higher expected wages, slower permitting and siting of new facilities, or other frictions—leading to persistently lower growth rates. This interpretation aligns with recent evidence that local institutions shape job creation rates and with earlier work emphasizing that frictions affecting job creation and reallocation can have first-order effects on regional employment outcomes (Davis and Haltiwanger, 2014).

5 Conclusion

To conclude, we highlight key takeaways from our study. For the historical question we study—understanding the move of manufacturing from the Rust Belt to the South and West—our evidence and theory supports an emerging narrative that emphasizes institutional forces, rather than, say, industry composition. We show that spatial differences in job creation costs that are strongly correlated with union activity provide a natural explanation for the strong link between employment growth and labor market churn we document.

We further emphasize three points that we expect to be relevant beyond our concrete application. The first is that LTS data offers a detailed look at worker flows across space and time a period of US history that featured large population movements as well as cross-regional differences in growth rates together with deep recessions. Quantitative analysis of search models that make predictions for gross worker flows and their business cycle volatility can thus confront a much longer history than what is contained in the standard JOLTS data, as well as much more granular spatial heterogeneity. Variation in business cycle experiences across space further help identify where volatility can come from and how it is amplified.

Second, our modeling framework heavily uses the idea that jobs work like capital goods: they are costly to create and long-lived. Its key implication is that job creation is sensitive to parameters, much more so than in the standard DMP model where firms create matches of relatively short duration. Sensitivity underlies our results on both spatial and time series variation. On the one hand, jobs respond strongly to cost differences across cities, much more so than workers, who trade off utility gains against migration costs. This is what makes growing markets tight, the key feature at the heart of our mechanism. On the other hand, when jobs are destroyed at a higher rate in downturns, job creation responds strongly so vacancies decline. This is what allows our model to match the volatility of unemployment over the cycle.

Finally, we view our study as an example of how worker reallocation at different frequencies is closely connected, and modeling benefits from embracing this fact. In particular, the magnitude of job creation costs provides a reason both for different city-level employment trends and different city-level exposure to (common) business cycle shocks. Joint analysis of growth and the cycle is thus key to understand the mechanism, and fitting parameters based on both sets of facts sharpens estimates. This is in contrast much current practice in business cycle analysis that rither detrends the data in a first step or employs assumptions on preferences and technology that allow for limited interaction between trend and cycle. The tractability of labor search models makes a joint approach feasible, and our results confirm that it is fruitful.

References

Abraham, K. G. (1983). Structural/frictional vs. deficient demand unemployment: some new evidence. *The American Economic Review 73*(4), 708–724.

Alder, S. D., D. Lagakos, and L. Ohanian (2023). Labor market conflict and the decline of the rust belt. *Journal of Political Economy 131*(10), 2780–2824.

Amior, M. and A. Manning (2018, July). The persistence of local joblessness. *American Economic Review 108*(7), 1942–70.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review 103*(6), 2121–2168.

Autor, D. H., D. Dorn, and G. H. Hanson (2016). The china shock: Learning from labor market adjustment. *Annual Review of Economics 8*, 205–240.

Bartik, T. J. (1991). *Who benefits from state and local economic development policies?* WE Upjohn Institute for Employment Research.

Beaudry, P., D. A. Green, and B. Sand (2012). Does industrial composition matter for wages? a test of search and bargaining theory. *Econometrica 80*(3), 1063–1104.

Beaudry, P., D. A. Green, and B. Sand (2014). Spatial equilibrium with unemployment and wage bargaining: Theory and estimation. *Journal of Urban Economics 74*, 2–19.

Bilal, A. (2023). The geography of unemployment. *Quarterly Journal of Economics 138*(3), 1507–1576.

Blanchard, O. and P. Diamond (1990). The cyclical behavior of the gross flows of u.s. workers. *Brookings Papers on Economic Activity 21*(2), 85–156.

Borovičková, K. (2016). Job flows, worker flows and labor market policies.

Burgess, S. and H. Turon (2010). Worker flows, job flows and unemployment in a matching model. *European Economic Review 54*(3), 393–408.

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

Davis, S. J., J. R. Faberman, and J. Haltiwanger (2012). Labor market flows in the cross section and over time. *Journal of Monetary Economics 59*, 1–59.

Davis, S. J. and J. Haltiwanger (2014, September). Labor market fluidity and economic performance. Working Paper 20479, National Bureau of Economic Research.

Elsby, M. W., R. Michaels, and G. Solon (2009). The ins and outs of cyclical unemployment. *American Economic Journal: Macroeconomics 1*(1), 84–110.

Faberman, R. J., A. I. Mueller, A. Şahin, and G. Topa (2022). Job search behavior among the employed and non-employed. *Econometrica 90*(4), 1743–1779.

Foote, C. (1998). Trend employment growth and the bunching of job creation and destruction. *Quarterly Journal of Economics 113*, 809–834.

Fujita, S. and M. Nakajima (2016). Worker flows and job flows: A quantitative investigation. *Review of Economic Dynamics 22*, 1–20.

Fujita, S. and G. Ramey (2009). The cyclicality of separation and job finding rates. *International Economic Review 50*(2), 415–430.

Gertler, M., C. Huckfeldt, and A. Trigari (2016). Unemployment fluctuations, match quality, and the wage cyclicality of new hires. Technical report, National Bureau of Economic Research.

Hall, R. E. (1972). Turnover in the labor force. *Brookings Papers on Economic Activity 3*.

Heise, S. and T. Porzio (2022). Labor misallocation across firms and regions. Technical report, National Bureau of Economic Research.

Hirsch, B. T. and D. A. Macpherson (2003). Union membership and coverage database from the current population survey: Note. *ILR Review 56*(2), 349–354.

Jarosch, G. (2023). Searching for job security and the consequences of job loss. *Econometrica 91*(3), 903–942.

Kennan, J. and J. R. Walker (2011). The effect of expected income on individual migration decisions. *Econometrica 79*(1), 211–251.

Kiyotaki, N. and R. Lagos (2007). A model of job and worker flows. *Journal of political Economy 115*(5), 770–819.

Kline, P. and E. Moretti (2014). People, places and public policy: Some simple welfare economics of local economic development programs. *Annual Review of Economics 6*(1), 629–662.

Kuhn, M., I. Manovskii, and X. Qiu (2021). The geography of job creation and job destruction. Technical report, National Bureau of Economic Research.

Lee, D. S. and A. Mas (2012, 01). Long-run impacts of unions on firms: New evidence from financial markets, 1961–1999 *. *The Quarterly Journal of Economics 127*(1), 333–378.

Lindenlaub, I., R. Oh, and M. Peters (2022). Firm sorting and spatial inequality. Technical report, National Bureau of Economic Research.

Lise, J. and J.-M. Robin (2017). The macrodynamics of sorting between workers and firms. *American Economic Review 107*(4), 1104–35.

Martellini, P. (2022). Local labor markets and aggregate productivity. *Working Paper*.

Menzio, G. and S. Shi (2011). Efficient search on the job and the business cycle. *Journal of Political Economy 119*(3), 468–510.

Mercan, Y. and B. Schoefer (2020). Jobs and matches: Quits, replacement hiring, and vacancy chains and vacancy chains. *American Economic Review: Insights 2*(1), 101–124.

Moretti, E. (2011). Local labor markets. Handbook of Labor Economics.

Moscarini, G. and F. Postel-Vinay (2018). On the job search and business cycles. Working Paper, Yale.

Nagypál, É. (2005). On the extent of job-to-job transitions. *Unpublished Manuscript, Northwestern University*.

Redding, S. J. (2016). Goods trade, factor mobility and welfare. *Journal of International Economics 101*, 148–167.

Rosen, S. (1979). Wage-based indexes of urban quality of life. In P. Mieszkowski and M. Straszheim (Eds.), *Current Issues in Urban Economics*, pp. 74–104. Johns Hopkins University Press.

Şahin, A., J. Song, G. Topa, and G. L. Violante (2014). Mismatch unemployment. *American Economic Review 104*(11), 3529–3564.

Saks, R. E. and A. Wozniak (2011). Labor reallocation over the business cycle: New evidence from internal migration. *Journal of Labor Economics 29*(4), 697–739.

Schmutz, B. and M. Sidibe (2019). Frictional labour mobility. *Review of Economic Studies 86*(4), 1779–1826.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American economic review 95*(1), 25–49.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics 15*(2), 127 – 148.

Topel, R. H. (1986). Local labor markets. *Journal of Political Economy 94*(3), 111–143.

Utter, C. M. (1982, June). Labor turnover in manufacturing: the survey in retrospect. *Monthly Labor Review 105*(6), 30–36.

Wang, S. and S. Young (2024). Unionization, employer opposition, and establishment closure.

Yashiv, E. (2007). Us labor market dynamics revisited. *Scandinavian journal of Economics 109*(4), 779–806.

Yoon, C. (2017). Estimating a dynamic spatial equilibrium model to evaluate the welfare implications of regional adjustment processes: The decline of the rust belt. *International Economic Review 58*(2), 473–497.

A List of cities

Table 3: List of sample cities and initial manufacturing employment

MSA	Man. Emp. in 1969	MSA	Man. Emp. in 1969
Akron, OH	100,760	Jacksonville, FL	29,554
Albany, NY	63,959	Kansas City, MO	137,270
Allentown, PA	114,890	Lancaster, PA	56,985
Altoona, PA	16,374	Little Rock, AR	33,172
Atlanta, GA	190,310	Louisville, KY	130,840
Baltimore, MD	208,960	Memphis, TN	75,803
Binghamton, NY	43,139	Miami, FL	120,590
Birmingham, AL	79,393	Milwaukee, WI	215,590
Boston, MA	425,650	Minneapolis, MN	235,350
Buffalo, NY	180,330	New York, NY	1,879,700
Canton, OH	64,661	Oklahoma City, OK	38,817
Cedar Rapids, IA	30,157	Orlando, FL	25,238
Charlotte, NC	183,000	Philadelphia, PA	661,160
Chicago, IL	1,137,100	Phoenix, AZ	78,477
Cincinnati, OH	206,330	Pittsburgh, PA	323,250
Cleveland, OH	356,390	Portland, OR	100,430
Columbus, OH	126,020	Providence, RI	209,210
Dallas, TX	230,930	Reading, PA	60,516
Dayton, OH	134,880	Richmond, VA	68,369
Denver, CO	77,540	Rochester, NY	156,510
Des Moines, IA	28,273	St. Louis, MO	291,810
Detroit, MI	632,530	Salt Lake City, UT	28,997
Elmira, NY	15,892	Seattle, WA	186,810
Erie, PA	43,215	Syracuse, NY	67,903
Fort Smith, AR	16,890	Tampa, FL	56,745
Greensboro, NC	97,035	Toledo, OH	82,321
Greenville, SC	88,314	Trenton, NJ	41,394
Harrisburg, PA	40,660	Tulsa, OK	48,686
Hartford, CT	167,790	Utica, NY	44,178
Houston, TX	162,080	Wichita, KS	52,774
Indianapolis, IN	170,420	York, PA	54,590
Jackson, MS	23,828	Youngstown, OH	114,140

B Business cycle dynamics

In this appendix, we present our business cycle classification scheme in more detail. After describing our algorithm, we apply it to the recent JOLTS data to clarify its properties.

B.1 Clustering algorithm

Our goal is to sort all months in our sample into recession and boom, based on similar behavior of aggregate layoff and hiring rates. As a preliminary step to classification, we linearly detrend and smooth aggregate layoff and hiring rates using a 5-month equally weighted and centered moving average. We estimate, by maximum likelihood, the parameters of a mixture distribution of two multivariate normal random variables. This step implies, for each month, a conditional probability that it was drawn from either of the two normal distributions. We then assign each month to the most likely multivariate normal that generated it.

Panel (a) of Figure 17 demonstrates the classification. Each dot represents a month. The blue dots correspond to months classified as a "booms", and the red dots correspond to the months classified as "recessions". The contour lines trace the unconditional joint distribution. Panel (b) of Figure 17 shows the classification over time. The blue line is the indicator for a recession and the red line shows the estimated probability of a recession. The clustering algorithm has a remarkable resemblance to the peak-trough timing by the NBER Business Cycle Dating Committee (gray bars).

There are three identified "recessions" in the sample period that overlap with most of the months of the NBER timing. We thus use the term recessions without distinguishing it from the NBER term in the main analysis. The notable difference is that the clustering algorithm identifies the start and end of recession later than NBER timing. Importantly, the classification captured these elements without any other information, such as an initial guess or a requirement that recessions last at least 3 quarters.

Table 4 demonstrates the differences between our classification and that of the NBER. Panel (a) shows the mean rates by NBER classification, and panel (b) shows the mean rates by clustering classification. The number of months in recession is slightly lower based on clustering classification (30, compared to 34). And in both classifications, the layoff rate is significantly higher during recessions. Moreover, based on the NBER classification the hiring and quit rates are almost not fluctuating between boom and recession, and based on the clustering algorithm they move more significantly.

Table 4: Comparison of aggregate state dating

Panel a: mean rates and number of months based on NBER business cycle timing

	Hiring rate	Quit rate	Layoff rate	Num. months
Boom	44.5	20.8	25.8	118
Recession	42.2	20.9	33.3	34

Panel b: mean rates and number of months based on clustering algorithm

	Hiring rate	Quit rate	Layoff rate	Num. months
Boom	44.8	21.4	25.5	122
Recession	40.6	18.2	35.2	30

B.2 Application to JOLTS data

To cross-validate our method, we apply the same methodology to more recent data from the Job Opening and Labor Turnover Survey (JOLTS). JOLTS data are available for all establishments from 2000-2024. Figure 18 demonstrate the result. Similar to the outcome of the clustering algorithm in the LTS data, clustered recessions overlap with every NBER recession, and end at a later date. Here, too, our classification algorithm is able to capture larger business cycle variation in worker flows.

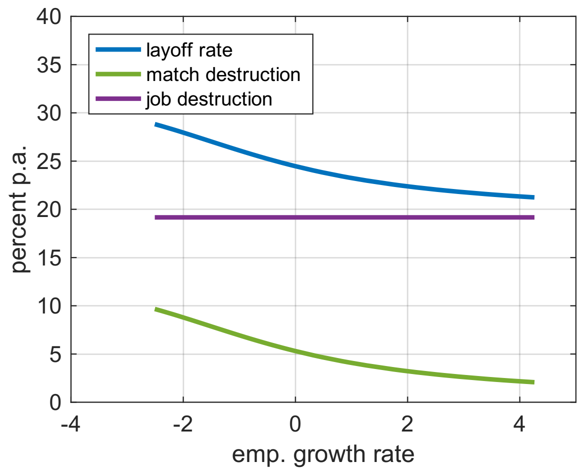

Figure 14: Decomposition of layoff rates in the cross section.

Notes. The figure shows a decomposition of the mean layoff rate (blue line) into match (green) and job (purple) destruction in the simulated model. The job destruction rate is common to all cities. Despite having the same match destruction rate conditional on match quality, faster growing cities have a smaller share of workers employed in unstable matches and so have a lower overall layoff rate.

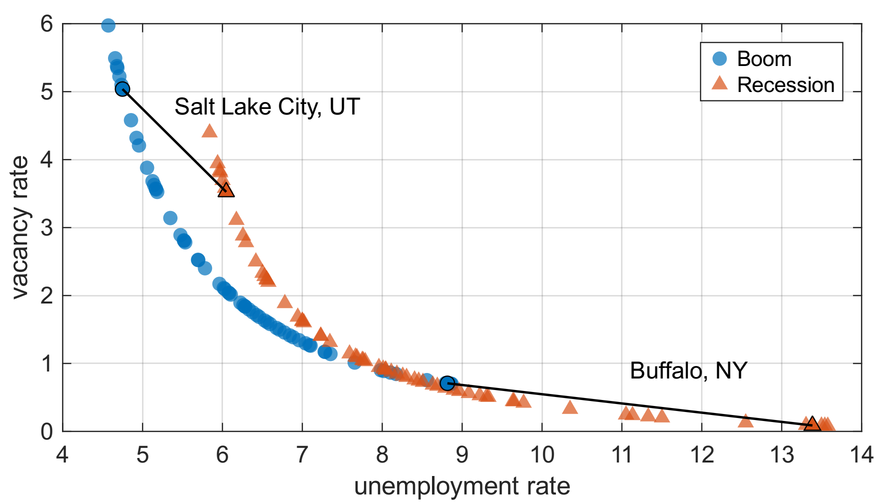

Figure 15: Beverage Curve

Notes: The figure shows the simulated Beveridge curves in the cross-section of cities. Each city is represented by two dots: blue circle (boom) and red triangle (recession). The horizontal axis is the mean unemployment rate and the vertical axis is the mean vacancy rate. The model features a downward sloping Beveridge curve in booms and recessions. Cities move down and to the right in recession, but in different proportions. Faster growing cities like Salt Lake City, UT see a modest fall in vacancies and rise in unemployment, whereas shrinking cities like Buffalo, NY see a large increase in unemployment.

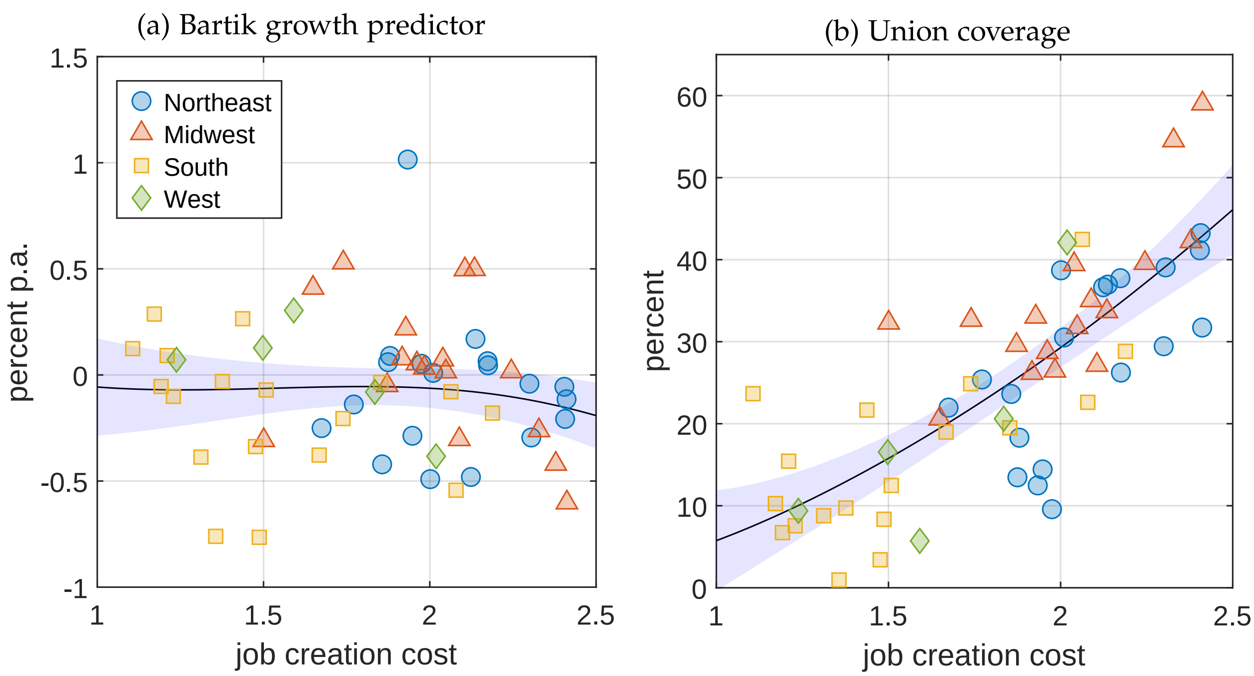

Figure 16: Job creation cost and covariates

Notes: The figure shows how the estimated city specific job creation cost varies with two city features: A Bartik predictor of city specific manufacturing employment growth (panel (a)), and a measure of union coverage in manufacturing (panel (b)). See text for details on the construction of the covariates.

(a) Hirings and layoffs

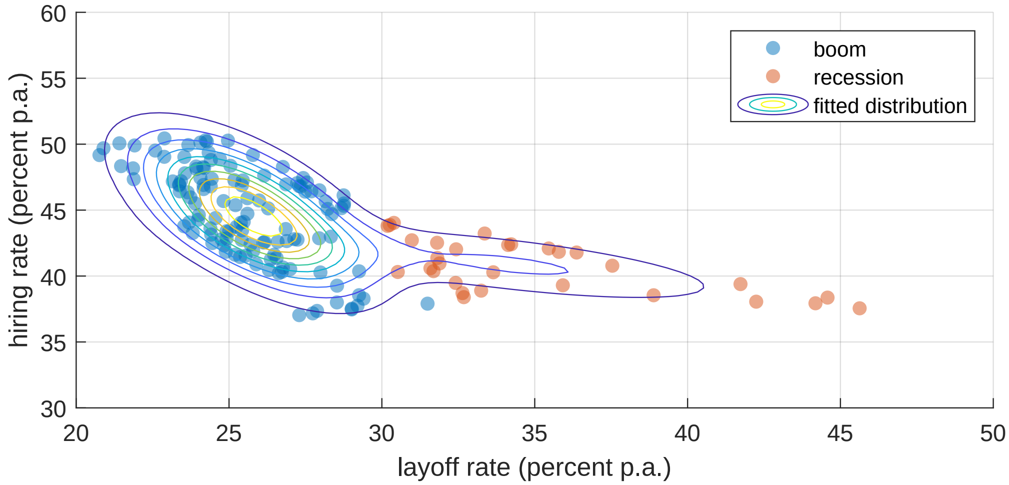

(b) Over time

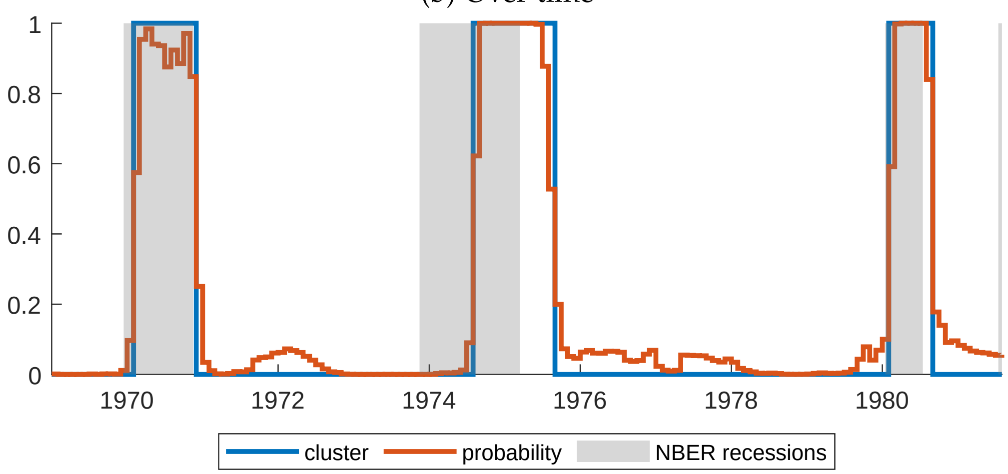

Figure 17: Classification of months to boom and recession

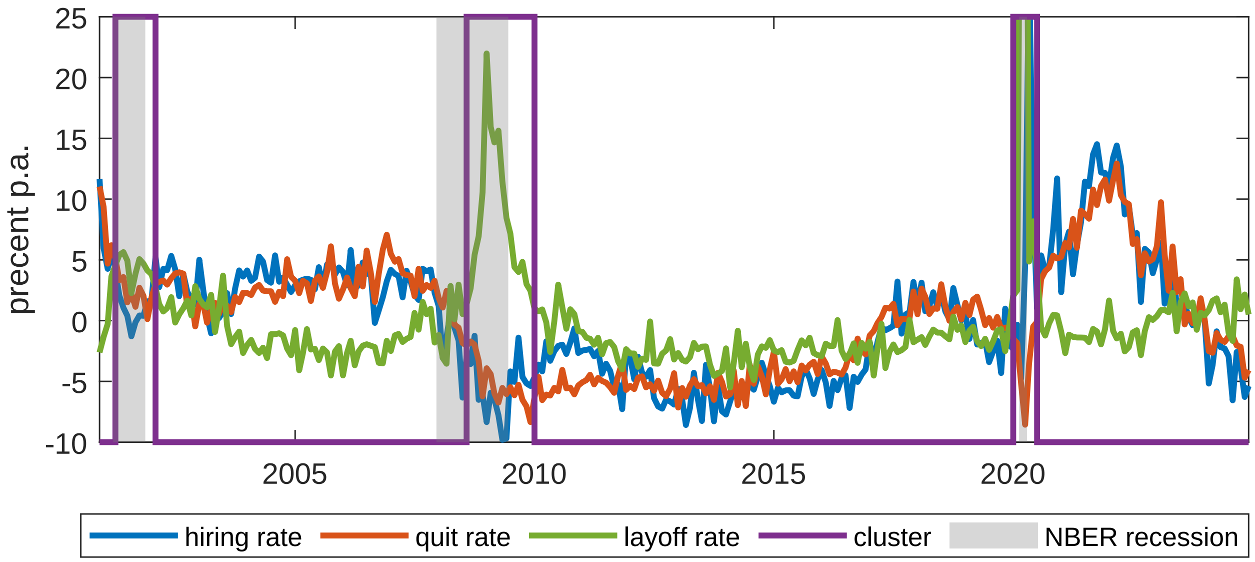

Figure 18: Aggregate labor-market flows and aggregate states, JOLTS

Notes: The figure shows the hiring rate, quit rate, and layoff rate in manufacturing together with NBER recessions and clustered recessions, based on JOLTS.

C Additional Figures and Tables

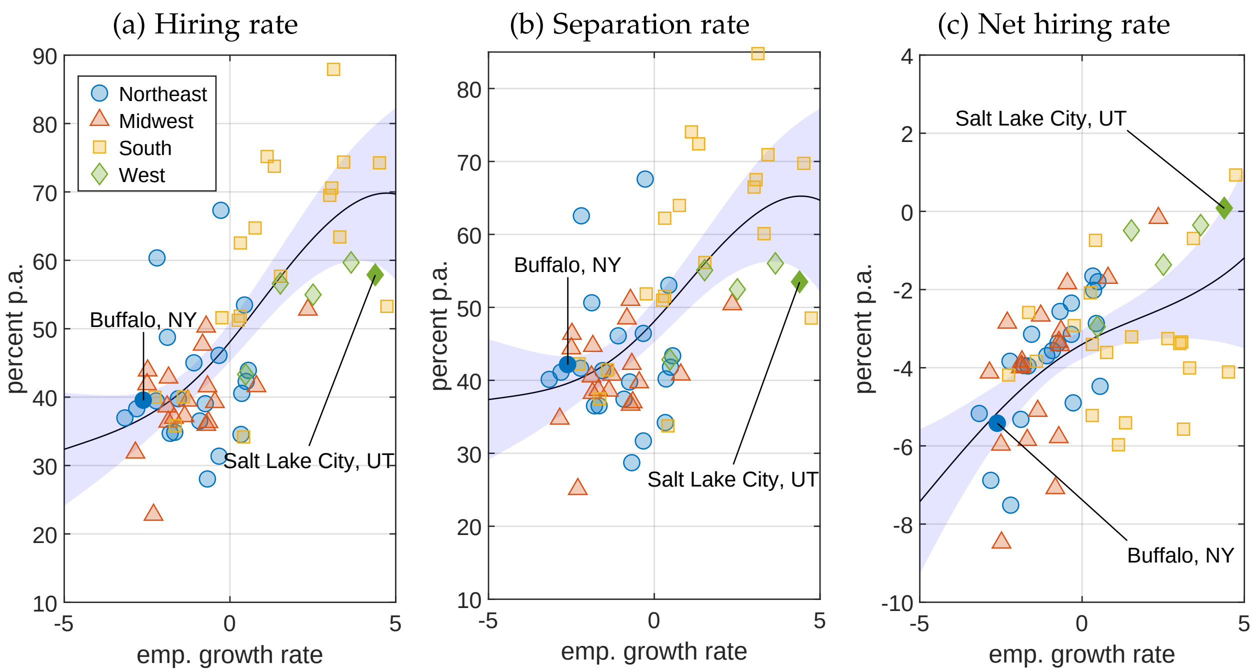

Figure 19: The cross-section of gross hiring, separation, and net hiring rates

Notes: The figure shows the cross-sectional joint distribution of mean employment growth in manufacturing and the mean hiring rate (panel (a)), separation rate (panel (b)), and the net hiring rate (panel (c)). Each dot represents a city. Marker shape and color represent the census region of the city. All rates are in percentage point per year. The mean hiring rate here is unadjusted, based only on LTS data. The solid line is a local weighted regression (LOWESS) using a Gaussian kernel with bandwidth 2. Shaded area is 95% confidence interval based on bootstrapped standard errors with 1000 replications.

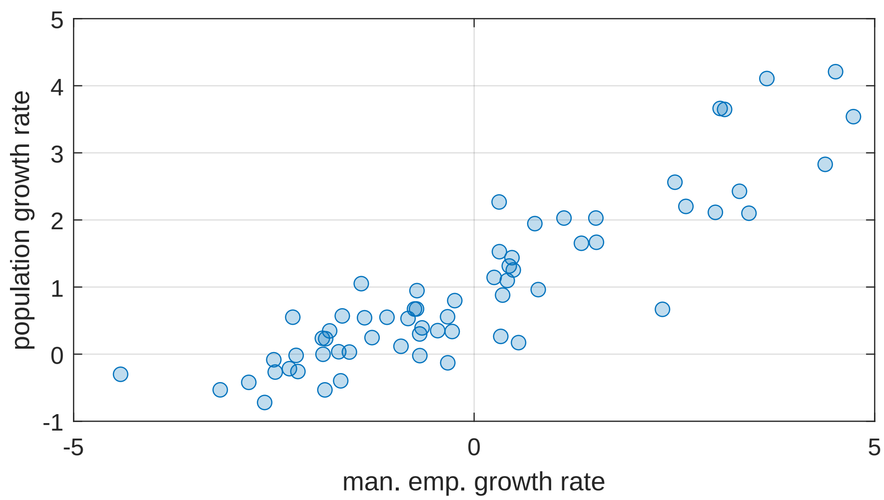

Figure 20: Employment growth and population growth

Notes: The figure shows the tight relationship between employment growth in manufacturing and overall population growth in the cross-section of cities. Every dot represents a city. The horizontal axis is the mean growth rate of employment in manufacturing, 1969 to 1981, in percent per year. The vertical axis is the mean population growth rate in the same period, in percent per year. Data are obtained from the BEA.