

Foote, Christopher L.; Ryan, Richard W.

Working Paper

Labor-market polarization over the business cycle

Public Policy Discussion Papers, No. 12-8

Provided in Cooperation with:

Federal Reserve Bank of Boston

Suggested Citation: Foote, Christopher L.; Ryan, Richard W. (2012) : Labor-market polarization over the business cycle, Public Policy Discussion Papers, No. 12-8, Federal Reserve Bank of Boston, Boston, MA

This Version is available at:

<https://hdl.handle.net/10419/99112>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Labor-Market Polarization Over the Business Cycle

Christopher L. Foote and Richard W. Ryan

Abstract:

During the last few decades, labor markets in advanced economies have become “polarized” as relative labor demand grows for high- and low-skill workers while it declines for middle-skill workers. This paper explores how polarization has interacted with the U.S. business cycle since the late 1970s. Consistent with previous work, the authors find that recessions are strongly synchronized across workers with different skills. Even high-skill workers favored by polarization suffer during recessions; this is particularly true during the last two downturns. Additionally, there is no evidence that polarization is driving the recent drop in the job-finding rate that has caused an adverse shift in the Beveridge curve. With this synchronization in mind, the authors then investigate the labor-market transitions of unemployed workers during recessions. When job-finding rates fall in recessions, middle-skill workers appear no more apt to leave the labor force or take low- or high-skill jobs than they are during booms. All in all, the results imply that current distress in the U.S. labor market extends far beyond middle-skill workers, and that recessions in general do not induce reallocation of middle-skill workers to jobs with better long-term outlooks.

JEL Classifications: E32, F62, J24, J63, J64

Christopher L. Foote is a senior economist and policy advisor in the research department at the Federal Reserve Bank of Boston. His e-mail address is chris.foote@bos.frb.org. Richard W. Ryan is a graduate student in the economics department at the University of Michigan. His e-mail address is Richard.W.Ryan@gmail.com.

Thanks are due to David Autor, David Dorn, and Melanie Wasserman for kindly sharing their occupational classification codes. The authors also received helpful suggestions from Chris Reicher and seminar participants at the Brookings Institution, the Swedish Riksbank, and a joint conference sponsored by the Kiel Institute for the World Economy and the Federal Reserve Bank of Richmond.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bostonfed.org/economic/ppdp/index.htm>.

This paper presents preliminary analysis and results intended to stimulate discussion and critical comment. The views expressed herein are those of the authors and do not indicate concurrence by the Federal Reserve Bank of Boston, or by the principals of the Board of Governors, or the Federal Reserve System.

This version: December 6, 2012

1 Introduction

Recent decades have seen distinct winners and losers emerge in the U.S. labor market. Consistent with a century-long trend, labor demand for high-skill workers has grown rapidly, as new technologies improve the outcomes of workers who have both the skills and the flexibility to use them (Goldin and Katz 2008). Moving down the skill distribution, workers with mid-level skills have not fared as well. Many of these workers are employed in routine jobs that can be replaced by automation or offshored to countries where wages are lower. Canonical examples of such jobs include assembly-line workers in manufacturing plants and workers in standardized office clerical jobs. Finally, low-skill jobs have proven relatively immune to replacement by automation or trade. While it may not take much formal education to clean a house or mow a lawn, a robot capable of doing most household chores has yet to be built, nor is it possible to ship a lawn to another country to be mowed. The hollowing out of job opportunities in the middle of the skill distribution has been termed the “polarization” of the labor market (Autor 2010; Acemoglu and Autor 2011; Autor and Dorn 2009; Autor, Katz, and Kearney 2008). In support of the idea that the decline in middle-skill jobs stems from the types of tasks that middle-skill workers perform and not from country-specific labor-market policies, researchers have found evidence of polarization in other advanced economies within Europe (Goos, Manning, and Salomons 2009).

Most of the empirical work on polarization has focused on the long-run relationship between labor-market polarization and wage inequality. Less research has explored how polarization might be related to the business cycle.¹ This paper attempts to fill part of that gap by measuring the degree of cyclical synchronization in the labor-market experiences of U.S. workers from different skill classes, and then asking how that synchronization affects the cyclical reallocation of workers across different skill types.

To see why a cyclical investigation of polarization is timely, consider Figure 1, which updates a figure in Autor (2010) and Acemoglu and Autor (2011). The bar chart depicts employment growth for 10 consistently defined occupations over four time periods: 1979–1989, 1989–1999, 1999–2007, and 2007–2011.² Appearing on the far left side of the chart are three high-skill occupations (managers, professionals, and technicians). On the far right side are three low-skill service occupations (protective services, food preparation and building and grounds cleaning, and personal care and personal services), which leaves four middle-skill occupations in the center (sales; office and administration; production, craft and repair; and operators, fabricators, and laborers). As Autor (2010) notes, the polarization of the labor

¹Some exceptions to this statement include Jaimovich and Siu (2012) and Faberman and Mazumder (2012), which we discuss below.

²The 10 occupations are created by grouping lower-level occupational codes available in the Current Population Survey (CPS). We explain this procedure more fully below. David Autor, David Dorn, and Melanie Wasserman kindly shared these codes with us.

market is clearly depicted in this graph; in the 1980s, 1990s, and early 2000s the occupations in the middle of the skill distribution experienced the weakest long-run employment growth.

With respect to current problems in the U.S. labor market, a particularly troubling feature of Figure 1 is shown by the gray bars, which indicate that middle-skill occupations had the worst employment performance during and immediately after the Great Recession (2007-2011). Poor middle-skill job growth over both the long and short horizons suggests that long-term structural forces exacerbated the massive job losses that took place during the Great Recession and have slowed employment growth in the ensuing recovery. The figure’s implications become less clear, however, if the employment figures are evaluated relative to trends—arguably employment in personal care and personal services fared worst relative to previous rates of employment growth.

Establishing a relationship between labor-market polarization and the business cycle would inform both theory and policy. Business cycle theorists have long investigated potential links between recessions and the reallocation of productive factors across alternative uses. Some papers contend that firms are more likely to reorganize production during cyclical downturns when the opportunity cost of foregoing current production in favor of reallocation is low. Other papers have suggested that allocational shocks help cause recessions in the first place.³ For current policymakers, a cyclical component to labor-market polarization could shed light on why recent U.S. recoveries have tended to feature slow employment growth, as suggested by Jaimovich and Siu (2012). In particular, polarization may explain why the degree of apparent mismatch between job vacancies and unemployed workers rose in the wake of the Great Recession. Figure 2 shows the well-documented adverse shift in the empirical relationship between job vacancies and unemployment—the Beveridge curve—which suggests that the labor market now produces fewer job matches for a given number of job vacancies.⁴ A possible reason for this shift is that the U.S. workers who lost jobs in the Great Recession were predominately middle-skill workers that firms do not want to hire, and concern that polarization is hindering the jobs recovery is now part of the policy debate. According to the public minutes of the September 2012 meeting of the Federal Reserve’s Federal Open Market Committee:

A few participants reiterated their view that the persistently high level of unem-

³The early seminal papers in the recessions-as-reallocations debate include Lilien (1982), Abraham and Katz (1986), Davis (1987), Davis and Haltiwanger (1990), Hall (1991), and Davis and Haltiwanger (1992). For more recent work, see Aghion and Saint-Paul (1998), Barlevy (2002), Caballero and Hammour (2005), and Lee and Mukoyama (2012).

⁴Because the Beveridge curve plots the number of unemployed workers against the number of job vacancies, the outward shift could have arisen from higher job separations among employed workers, not just low job-finding rates among the unemployed. But the preponderance of current research asserts that low finding rates generated the shift. See, for example, Barchenon et al. (2011), Barchenon and Figura (2012), and Sahin et al. (2012). Additionally, Davis, Faberman, and Haltiwanger (2010) provide evidence that the Great Recession and subsequent recovery saw a pronounced drop in “recruiting intensity” among firms.

ployment reflected the effect of structural factors, including mismatches across and within sectors between the skills of the unemployed and those demanded in sectors in which jobs were currently available. It was also suggested that there was an ongoing process of polarization in the labor market, with the share of job opportunities in middle-skill occupations continuing to decline while the shares of low and high skill occupations increased. Both of these views would suggest a lower level of potential output and thus reduced scope for combating unemployment with additional monetary policy stimulus (Federal Open Market Committee 2012, p. 7).

Yet it is not a simple exercise to establish a direct link from polarization to either recent movements in the Beveridge curve or to the recessions-as-reallocations theory. As we will see, a high fraction of middle-skill jobs are in manufacturing and construction, two sectors that are highly responsive to the business cycle. Large middle-skill losses during the Great Recession may therefore reflect the typical job losses that occur whenever economic growth declines, so it may be the case that the gray bars in Figure 1 overstate the extent to which middle-skill job losses reflect structural forces like polarization. Additionally, much of the recessions-as-reallocations literature implies that the incentives for workers to reallocate themselves go up in recessions. To our knowledge, convincing empirical evidence for this assertion has yet to emerge.

This paper uses individual-level data from the Current Population Survey (CPS) to ask some basic questions about the experiences of different skill classes over the last 30 years of the U.S. business cycle, paying particular attention to the most recent recession.⁵ To partially control for industry effects, in most of the analysis middle-skill workers are separated into three subclasses—middle-skill manufacturing, middle-skill construction, and middle-skill “other.” Adding high- and low-skill workers from all industries to the three middle-skill groups gives us a total of five industry–skill groups to analyze.⁶

The first set of empirical results concerns the strong *synchronization* in the business cycle experiences of workers in different industry–skill groups. Using some familiar statistical techniques, we estimate the common and idiosyncratic variation in the unemployment rates, job-finding rates, and job-separation rates for the different groups.⁷ Studies that analyze

⁵Following Autor (2010) and Acemoglu and Autor (2011), we use the occupation variable in the CPS to assign employed workers into different classes. For the unemployed, the CPS lists the individual’s most recent occupation, so we can classify jobless workers as well as employed ones. Because our paper does not investigate workers’ wages, we are free to use the entirety of the monthly files from the CPS (which do not have wages) not just the Merged Outgoing Rotation Groups (which include wage information).

⁶We perform robustness tests to ensure that our results are not driven by the inclusion of manufacturing and construction workers in the high-skill category.

⁷We define the job-finding rate as the rate at which unemployed workers find jobs and the job-separation rate as the rate at which employed workers enter the unemployment pool. These definitions do not account

disaggregated macroeconomic data with these statistical techniques often find that much of the variance in any component series is driven by common factors, and this paper is no exception.⁸ Skill-specific unemployment rates and job flows move together strongly over the business cycle, even though individual rates and flows often have different means and variances. With respect to the most recent recession, this synchronization is consistent with other work finding labor-market distress occurring across many different industries and varied demographic and educational groups (Elsby, Hobijn, and Sahin 2010; Dickens and Triest 2012). Our contribution is to show that comovement is also apparent after sorting workers on the basis of the long-run outlooks for their occupations. The results argue against a strong role for polarization in driving the currently slow U.S. recovery. In particular, we find that recent movements in job vacancies and the common component of job-finding rates do not support the view that labor-market polarization is responsible for the outward shift in the Beveridge curve depicted in Figure 2. And recent idiosyncratic movements in job flows provide no evidence that recessionary periods are becoming relatively easier for high-skill workers who have been favored by polarization trends. If anything, the recessions of 2001 and 2007–2009 were especially difficult for high-skill workers, given the experiences of similar types of workers in earlier downturns.

Synchronization in labor-market outcomes provides the context for the second set of empirical results focusing on the *reallocation* of unemployed workers in recessions and booms. Large numbers of middle-skill workers, especially in manufacturing, typically separate from employment in recessions. But the synchronization of business cycles discussed above means that these bursts of middle-skill separations take place at the same time that all other workers experience low job-finding rates. Thus, it is perhaps unsurprising that we find relatively little cyclical variation in the rate at which unemployed middle-skill workers transition out of unemployment to either high- or low-skill jobs. In fact, unemployed middle-skill workers appear reluctant or unable to move to other skill classes, as unemployed middle-skill workers who do find jobs accept middle-skill employment more than 75 percent of the time. While the fraction of middle-to-middle movements is trending down over time, this share does not display

for employment-to-employment flows or transitions into and out of the labor force. However, we broaden our analysis to encompass nonparticipation when we study the cyclical reallocation of workers to different skill classes, as described below. Also, because we want to know how worker skills affect job-finding probabilities, we calculate finding rates only for workers who held a previous job; that is, we exclude unemployed workers who are new entrants to the labor force. Finally, we use the term “job flows” when referring to job-finding and job-separation rates. In the academic literature, these flows are often labeled “worker flows.”

⁸Using techniques similar to ones below, Rissman (2009) and Reicher (2012) analyze common and idiosyncratic variation in employment data disaggregated by industry. The focus in this paper is on worker-level flows and unemployment rates disaggregated by skill. Tasci (2012) analyzes common and idiosyncratic variation in aggregate job flows to produce estimates of the long-term natural rate of unemployment, and Fleischman and Roberts (2011) extract a common business cycle from a variety of aggregate time series, including gross domestic product (GDP), gross domestic income, the labor-force participation rate, and the unemployment rate.

much cyclical variation.⁹ However, this is not the case for the share of unemployed workers who exit unemployment for labor-force nonparticipation.¹⁰ When the overall job-finding rate falls, the share of unemployed workers who leave unemployment for nonparticipation rises. But this pattern does not come about because unemployed workers are more likely to exit the labor force when finding rates decline. Rather, in recent recessions the explanation for the increased share of unemployment spells that end in nonparticipation owes more to a simple mechanical relationship: when job-finding rates fall, more unemployment spells end in labor-force nonparticipation because fewer unemployment spells end with transitions to employment.¹¹ Taken together, these findings argue against a straightforward theoretical link between recessions and polarization-based reallocation.

The rest of the paper is organized as follows. Section 2 discusses data issues confronting polarization studies that use high-frequency worker-level data. Section 3 presents the results on the synchronization of business cycles, and Section 4 investigates cyclical variation in the reallocation of unemployed middle-skill workers. Section 5 concludes by relating the paper’s empirical findings to the current debate on the cause of the slow U.S. recovery and the academic literature on the reallocation of labor in recessions.

2 Data

2.1 Occupational Classifications

One of the toughest tasks in the study of skill-based polarization is settling on the best way to classify workers. To a large extent we rely on previous work. Our data are drawn from the monthly basic files of the Current Population Survey (CPS), which categorize both employed and unemployed individuals into hundreds of detailed occupations. While the question used to elicit an individual’s occupation is constant throughout the sample, the classification system that the Bureau of Labor Statistics (BLS) uses to code occupations has changed near

⁹This lack of cyclical variation does not mean that workers do not upgrade their skills over the business cycle. But it suggests that cyclical upgrading tends to occur when workers move directly to new jobs without intervening spells of unemployment (Krause and Lubik 2006).

¹⁰A nonparticipating worker is neither employed nor looking for work. Sometimes this status is labeled “out-of-the-labor-force.”

¹¹Elsby, Hobijn, and Sahin (2012) argue that recessions release large numbers of workers with strong labor-force attachments into the unemployment pool. Consequently, the probability that an unemployed worker decides to leave the labor force declines in recessions due to a compositional effect. Our finding that the unconditional probability of moving from unemployment to nonparticipation does vary much over the cycle may reflect the specific demographic variables that we are holding constant as we investigate this correlation. It may also reflect the fact that we match workers in the CPS only on a month-to-month basis. Elsby, Hobijn, and Sahin (2012) link workers across 12 months and can therefore investigate the transition probabilities of unemployed workers who were previously employed 12 months ago. In any event, the authors find no evidence that unemployed workers are *more* likely to transition to nonparticipation in recessions, which is our conclusion too.

the start of each decade (Bowler et al. 2003, p. 18). Fortunately, Meyer and Osborne (2005) create a classification system that generates consistent occupational categories over time. Yet, while the Meyer-Osborne system has the advantage of consistency, it also has hundreds of entries—too many for a focused study of labor-market polarization.

Autor and Dorn (2009) and Autor (2010) aggregate the Meyer–Osborne occupations into the 10 coarser occupations, as shown in Figure 1. Figure 3 displays monthly employment levels for these 10 occupations from January 1976 (when micro-level CPS data become available) to April 2012. The vertical lines in each panel indicate a change in the BLS’s underlying occupational classifications. To the extent that both the Meyer–Osborne and the Autor–Dorn classification systems are consistent, there should be no discrete jumps in measured employment when these changes occur. Some consistency problems are apparent, however. The first row of graphs correspond to high-skill workers (managers, professionals, and technicians). After the first reclassification in 1983, many managers (far left graph of top row) appear to be reclassified as technicians (far right graph of top row). Data for the four middle-skill occupations are displayed in the middle row of Figure 3; the jump in the measured employment of sales workers (far left graph of middle row) suggests a reclassification issue after the 1983 change in occupational codes.¹² Finally, the measured employment of low-skill occupations, most notably personal care and personal services, also experienced big changes following reclassification dates.

To alleviate some of the reclassification problems, the employment levels from the 10 occupations are aggregated in the empirical work below. Figure 4 depicts monthly employment by skill after aggregating into high-, middle-, and low-skill categories, as in Autor and Dorn (2009), Autor (2010), and Acemoglu and Autor (2011). The effects of classification changes are less apparent, though some effect of the 1983 reclassification remains in high-skill employment, while a discrete increase in measured low-skill employment around the 2003 reclassification is also clear. The job-finding and job-separation rates we analyze are less susceptible to distortion from classification changes as long as these changes do not substantially alter the *type* of worker included in particular groups—even if the classification changes exert nontrivial effects on the *numbers* of workers in different groups. Figure 5 graphs the unemployment rates for the 10 individual occupations, providing some evidence that labor-market rates are not greatly affected by the reclassifications. Unemployment rates for the 10 occupations vary from month-to-month due to sampling error and the business cycle, but classification changes do not engender big jumps in these rates.¹³

Stepping back from classification effects, Figures 4 and 5 provide some visual evidence for

¹²Note that this increase also shows up in Figure 1, which indicates a big increase in sales employment in the 1979–1989 period.

¹³The unemployment rates in Figure 5 are adjusted for seasonal variation with a simple regression-based adjustment. We explain this adjustment more fully in Section 2.4.

a key assertion of this paper: while the labor-market outcomes of middle-skill workers vary more over the business cycle, recessions are bad for everyone. In Figure 5, the middle row shows that middle-skill workers are subject to large unemployment fluctuations. For example, workers in the production, craft, and repair group suffered jobless rates of approximately 15 percent during both the early 1980s and during the Great Recession. Unemployment rates for operators, fabricators, and laborers top out at more than 20 percent in the early 1980s and somewhat less than that recently. High-skill unemployment rates are lower; note the differences in the vertical scales in the top versus middle rows of Figure 5. But whenever middle-skill unemployment rises, so does high-skill unemployment. Moreover, for high-skill workers, the Great Recession stands out as a particularly difficult period. Unemployment rates for all three high-skill occupations rose to their highest levels since 1976 and have remained high since then.

2.2 The Role of Industry

In its most basic form, the polarization theory does not make specific predictions for the *cyclical* of middle-skill employment. Polarization predicts a negative long-run *trend* for middle-skill jobs as middle-skill tasks are automated or offshored.¹⁴ But disentangling cyclical and trend effects means dealing with a disproportionate share of middle-skill jobs in manufacturing and construction, two industries with highly cyclical demand. In Figure 6, Panel A graphs the industry shares of employment for the three broad skill classifications (high, middle, and low). The graph for high-skill jobs shows that at the start of the sample period, more than 80 percent of high-skill jobs were outside of manufacturing and construction; that share has risen even higher since then. The next graph shows that the share of middle-skill jobs outside of manufacturing and construction has also been rising, but starting from a smaller initial value, so that today more than 20 percent of middle-skill work remains in manufacturing and construction. The last graph in the panel shows that over the sample period virtually all low-skill work has been outside of manufacturing and construction. Panel B of Figure 6 shows why construction and manufacturing employment are so strongly represented in the middle-skill designation: most construction and manufacturing workers fall into the middle-skill group.

Up to now, the study of labor-market polarization has focused on occupation, rather than industry, but a study of polarization and cyclical employment should account for the special role that cyclical industries play in middle-skill employment. We examine this relationship in the simplest way possible by disaggregating middle-skill employment into manufacturing,

¹⁴As discussed below, Jaimovich and Siu (2012) provide a theory for why reductions in middle-skill employment should be concentrated during recessions when aggregate productivity is low, so that, as implied by the paper's title, "the trend [in middle-skill employment] is the cycle."

construction, and all other industries. Looking back at the employment levels in Figure 4 indicates that there is room to disaggregate middle-skill employment, as it has always been the largest of the three broad groups, despite polarization’s negative impact on middle-skill employment. Additionally, we choose to keep middle-skill employment in manufacturing and construction separate, mainly due to disparate long-run trends in those industries. As a share of total employment, manufacturing employment has been shrinking for most of the postwar era; it has been shrinking in absolute terms since the 2001 recession. Construction employment, on the other hand, has generally fluctuated from 4 to 6 percent of total employment since the 1950s.

Before investigating how similarly these five industry–skill groups behave over the business cycle, we point out one caveat regarding our chosen baseline classification. The top panel of Figure 6 shows that manufacturing and construction are heavily represented in middle-skill employment, but the figure also indicates that these two industries account for nontrivial shares of high-skill employment. We could separate manufacturing and construction employment from high-skill employment, as we do for middle-skill employment, but increasing the number of industry–skill groups would complicate the analysis; in particular, doing so could render the statistical models we estimate below less informative. We will therefore run robustness checks to make sure that any interesting results we find do not stem from the inclusion of some manufacturing and construction workers in the high-skill group.

2.3 Employment Growth, Educational Attainment, and Unemployment Rates for the Industry–Skill Groups

Figure 7 provides some detail on employment and educational attainment for our five industry–skill groups, with Panel A displaying employment levels. The lines for the high-skill and low-skill groups are the same as in Figure 4 because those designations are unchanged. The panel shows that the three disaggregated middle-skill groups have different long-run trends. The next two panels explore these employment trends more fully by indexing employment in each group to equal 100 in 1976:Q1. Panel B shows that from 1976 to the early 1990s, employment for the high-skill, low-skill, and middle-other groups grew at about the same rates. Shortly after the 1990–1991 recession, high-skill employment pulls away from the other two groups, leaving the middle-other and low-skill groups to grow together until the 2001 recession. At that point, an unfortunate CPS classification change causes measured low-skill employment to jump, making further trend comparisons difficult. Yet it is interesting that employment growth for the low and middle-other groups were so close for so long, given the expected negative effect of polarization on the middle-other group. The negative effect of polarization shows up much more strongly in the middle-manufacturing category, which is displayed with middle-construction in Panel C. The absolute decline in overall manufacturing

employment is consistent with the decline in absolute employment of middle-manufacturing workers over time. And, as pointed out by Jaimovich and Siu (2012), many of those losses are concentrated around recessions. Yet middle-construction workers have fared better. While middle-construction employment is cyclically sensitive, it grew to about 2.5 times its 1976:Q1 value by the eve of the Great Recession. Finally, Panel D measures educational attainment for the five groups. As we would expect, high-skill workers are most likely to be college graduates, though the college-educated are also found in middle- and low-skill work. Middle-skill workers outside of manufacturing and construction are also relatively well-educated, with more representation in the some-college or college categories than the other middle-skill or low-skill groups.

The employment data in Figure 7 provide useful background for our empirical work—particularly in regard to trend employment growth—but previously noted problems arising from changes in occupational classifications prompt us to emphasize instead the behavior of unemployment rates and job flows. Figure 8 graphs the quarterly unemployment rates for our five skill classifications. Later in the paper we will discuss data that have been adjusted for both seasonal effects and for the changing demographic composition of workers, but here we just plot the raw rates. As we saw in the unemployment rates among the 10 individual occupations in Figure 5, the unemployment rates of all workers move together but there are substantial differences in both levels and seasonal sensitivities across the five industry–skill groups.¹⁵ For example, our figure replicates the high and seasonal unemployment rates for construction workers found in published BLS data.¹⁶ The figure also shows that average unemployment for the middle-other group falls in the middle of the pack, and is lower than the rates for other middle-skill and low-skill workers and higher than the rate for the high-skilled group. And as we have seen, recent unemployment rates are high for all groups, even high-skill workers.

2.4 Adjusted Job-Finding and Job-Separation Rates

The unemployment rates in Figure 8 vary across groups because workers flow in and out of unemployment at different rates. As in previous research, job flows are calculated for those CPS respondents who have been matched across consecutive survey months. To ensure that the resulting flows reflect cyclical forces, the flows are adjusted in several ways. To account for changing demographics, we run year-specific, individual-level logit regressions for each of

¹⁵The unemployment rates in Figure 5 were adjusted for seasonal effects but not demographic composition.

¹⁶Of course, the middle-construction unemployment rate we calculate corresponds only to construction workers with mid-level skills, while the published construction unemployment rate is relevant for all construction workers. However, Panel B of Figure 6 shows that around 80 percent of construction employment is middle-skill. This fact implies a close correspondence between the middle-construction unemployment rate we calculate and the published rate for all construction workers.

the five industry–skill groups. For example, to adjust (say) job-finding rates, we use the logits to model the probability that an unemployed worker in a specific group will transition from unemployment to employment from one month to the next. The regressors in this model are demographic variables reflecting race, marital status, gender, education, and age. This approach generates separate sets of demographic coefficients that vary by year for each of the five groups. To calculate the adjusted job-finding probabilities for one of the groups, we use these year-specific logit coefficients to generate a series of expected job-finding probabilities for a “typical worker” in that group. This typical worker is defined as one with demographic characteristics equal to the within-group sample average. Adjustments for job-separation rates are figured similarly, using logits modeling the probability that an employed worker will transition to unemployment.

After adjusting the job flows for demographic composition, we then adjust these flows for time-aggregation.¹⁷ When a worker in the CPS loses a job, the worker has on average about two weeks to find a new job before the CPS’s next monthly survey week. If the worker does find a job, then the CPS will not record the job separation, causing the measured job-separation rate to be too low. A number of ways to account for time-aggregation bias have been suggested in the literature. We use the method concurrently proposed by Shimer (2011) and Elsby, Hobijn, and Sahin (2012). This approach uses the eigenvalues from the transition matrix of unadjusted flows to back out the implied continuous-time flow probabilities of moving from employment to unemployment or vice versa. Using these continuous-time probabilities, one can then calculate the probability of (say) moving from employment to unemployment from month t to $t + 1$, abstracting from the possibility that a new job can be found before the next survey date occurs.¹⁸ Finally, to seasonally adjust a time series, we simply regress the series on a constant and a set of quarterly dummies, omitting the dummy for the first quarter of the year. The seasonally adjusted series is the constant from this regression plus the residual.

In Figure 9, the top row presents the adjusted job-finding rates for the five industry–skill groups, with Panel A depicting the finding rates for high- and low-skill workers and Panel B depicting the corresponding rates for the middle-skill groups (to facilitate comparisons the vertical scales in the two panels are identical). There is substantial commonality across the five finding rates both in their means and their cyclical movements. Panel A shows that unemployed workers at the two ends of the skill spectrum (high and low) find jobs at nearly the same rates. Panel B shows that the job-finding rates for middle-manufacturing and middle-other workers are also nearly identical, even though Figure 7 showed stark differences in the

¹⁷For example, see the discussion of time-aggregation in Shimer (2005, pp. 32–33).

¹⁸The full time-aggregation correction described in Shimer (2011) and Elsby, Hobijn, and Sahin (2012) also allows transitions to and from labor-market nonparticipation, but we abstract from nonparticipation in our adjustments to the job-finding and job-separation rates.

long-run employment trends of these two groups. The finding rate for middle-construction workers is somewhat higher than that of other middle-skill groups, but the difference appears to be a near-constant level effect. The similar time-series behavior of job-finding rates across skill classes is consistent with Elsby, Hobijn, and Sahin (2010), who find that job-finding rates for workers with different levels of formal education move together closely over the business cycle, including the most recent one (that is, the Great Recession).

The bottom two panels of Figure 9 graph job-separation rates. Here, the differences across industry-skill groups are substantial. To begin with, high-skill workers separate into unemployment less often than other workers. Also, job separations for middle-manufacturing workers spike in recessions, though otherwise these separations are usually only slightly higher than separations for the middle-other group. The middle-skill panel also reflects the fluid nature of construction work. As with finding rates, separation rates for middle-construction workers are the highest of any group.

To the extent that our definition of skill correlates with formal education, the large difference in average job-separation rates across groups should not be surprising. An inverse relationship between the incidence of unemployment and formal education has been documented at least since Mincer (1991).¹⁹ Differences in unemployment incidence have been studied with respect to other demographic variables as well. In a recent paper, Hoynes, Miller, and Schaller (2012) note that the differences in unemployment incidence across demographic groups have proven remarkably stable over the past three decades, and that during the Great Recession these unemployment disparities correlated with demographic differences in industry and occupation.

Yet while the previous literature has emphasized differences in the unemployment incidence of various groups, it is also important to note the strong cyclical comovement in the separation rates of different groups. This comovement emerges more clearly when we graph job flows that have been standardized, so that each series has a mean of zero and a standard deviation of one. As shown in Figure 10, standardizing the separation rates shows in particular that even though high-skill workers suffer fewer separations during recessions than other workers, their separations cyclically rise and fall along with those of everyone else. Indeed, a salient feature of this graph is that in the last two recessions, the job-separation rate for high-skill workers is elevated relative to its sample history.

In the next section, we quantify the common and idiosyncratic variation in job-flows for

¹⁹Mincer (1991) found a negative relationship between the incidence of unemployment and years of education among male respondents in the Panel Survey of Income Dynamics. Further, differences in the duration of unemployment (roughly, the differences in job-finding rates) are much less important in driving the differences in unemployment rates across educational groups. Using a model with incomplete markets, Mukoyama and Sahin (2006) show that Mincer's finding on unemployment incidence implies that the costs of business cycles are relatively high for low-skill workers. Not only are low-skill workers unemployed more often, but their low levels of wealth make it harder for them to self-insure against fluctuations in consumption.

the five industry–skill groups. Doing so will allow us to investigate the importance of labor-market polarization as a potential explanation for parallels and contrasts in the experiences of differently skilled workers over the business cycle.

3 Synchronization of Business Cycles across Industry–Skill Groups

3.1 Principal Components Analysis

A common procedure to assess both common and idiosyncratic movements in a set of time series is principal components analysis (PCA). This nonparametric procedure models common variation in a system of variables as arising from a set of common factors, to which the individual series are related by so-called factor loadings. Formally, let y_{it} denote an unemployment rate or a job flow for industry–skill group $i \in \{1, \dots, 5\}$ in quarter t . The full principal components model for y_{it} is

$$y_{it} = \phi_i^1 F_t^1 + \phi_i^2 F_t^2 + \phi_i^3 F_t^3 + \phi_i^4 F_t^4 + \phi_i^5 F_t^5, \quad (1)$$

where the symbol F represents the factors that are common to all five groups and the symbol ϕ depicts the group-specific factor loadings. The F s are identified by imposing orthogonality among the factors, then ordering the factors sequentially by how much variance they can explain. In other words, F_t^1 and the individual ϕ_i^1 s are constructed to explain the maximum amount of total variation among the five individual rates. Then the factor F_t^2 and its associated factor loadings are constructed to explain the maximal amount of remaining variance, and so on for all five factors. Because there are five industry–skill classes, five factors will explain all of the variance in each system of labor-market data and thus will completely characterize each individual unemployment rate or job flow.²⁰ One limitation of PCA is that the individual component series being modeled should have variances of similar sizes, so it is common to perform PCA on standardized series, such those graphed in Figure 10.²¹ If this standardization is not performed and the variances of the component series differ greatly, then the first factor will be unduly influenced by the series with the largest variance. This limits the usefulness of the PCA as a characterization of common variance across all component series.

²⁰There is no error term in equation (1) because this decomposition includes all five factors. Also, while we have described the construction of the factors sequentially, in practice the factors are estimated simultaneously through an eigenvalue decomposition of the system’s covariance matrix. For a thorough yet accessible introduction to PCA, see Shlens (2009).

²¹We also seasonally adjust the component series before estimating the PCA. Regarding the individual variances, PCA works well when individual component variables have the same variances, but that variance does not have to equal one. Normalizing the component series to have unit variances is a common choice, however, and is equivalent to performing a PCA on the correlation matrix of the variables rather than the covariance matrix. The correlation–matrix choice is the default in the Stata software package, for example.

Selected results from three separate PCAs used to analyze unemployment rates, job-finding rates, and job-separation rates appear in Figure 11.²² The table at the top shows the cumulative share of total variance explained by adding additional factors (F^2 , F^3 , and so on) to each PCA. Because there are five industry–skill groups in each PCA, the row corresponding to F^5 shows that the share of total variance explained by using all five F s equals 1.00. More interesting are the results in the first row, which indicate that most of the variance in each system is explained by the first factor, F_t^1 , alone. Among standardized unemployment rates, the first factor and its associated factor loadings, the ϕ_i^1 s, explain 90.2 percent of the total variance in the system. Moving to the individual job flows, the corresponding share of total variance explained in the job-finding system is 86.4 percent, a number that confirms the pattern suggested visually in Figure 9: the job-finding rates of the different industry–skill groups move together closely over the cycle. But the table also shows that once job-separation rates are standardized, a single factor explains 66.9 percent of their total variance—a smaller but still substantial amount.

PCA can also be used to study idiosyncratic variation in component series. Consider a model that uses only one factor to model common variation: $y_{it} = \phi_i^1 F_t^1 + e_{it}$. Here, the contributions of the second through fifth factors have been folded into the error term e_{it} , which reflects idiosyncratic variation orthogonal to the common variation driven by F_t^1 . It turns out that idiosyncratic movements in finding and separation rates for high-skill workers provide some interesting insights on the potential role of labor-market polarization in recent business cycles. Panels B and C of Figure 11 depict standardized high-skill finding and separation rates (y_{it}) along with the one-factor predictions ($\phi_i^1 F_t^1$); the gap between these lines equals e_{it} . Panel B shows that the actual and predicted finding rates for high-skill workers line up closely, as we would expect given the high degree of total variance in finding rates explained by only one factor. But the predicted high-skill finding rate dips below the actual finding rate for an extended period in the early 2000s.²³ Additionally, Panel C shows that the predicted high-skill separation rate lies below the actual rate during the last two business cycles. In other words, relative to their experiences in previous business cycles, high-skill workers experienced job-finding rates that were lower than expected in the early 2000s, while they experienced separation rates that were higher than expected during the two most recent business cycles.²⁴ If recent business cycles were strongly affected by labor-market polarization, we would not necessarily expect this pattern.

²²Full results of the PCAs for are available from the authors upon request.

²³There is also a shorter-lived drop in the early 1990s.

²⁴The (unreported) full results of the PCA indicates that the pattern of idiosyncratic movements for high-skill workers is essentially the opposite of the pattern experienced by low-skill workers, who enjoy higher-than-normal finding rates during the early 2000s and lower-than-normal separation rates during the two most recent recessions. This pattern is also apparent in the idiosyncratic movement estimated by the dynamic factor models, which are discussed below.

3.2 Dynamic Factor Models: Specification and Estimation

PCA is often used to characterize comovement in different series, but the need to standardize the data beforehand means that the variances that a PCA attempts to estimate are in some sense artificial. We therefore estimate dynamic factor models (DFMs), which in our case do not require pre-standardization. Like a PCA, a DFM generates a latent factor (or factors) to drive common variation in the system. Our DFM has one common factor and is specified as follows:

$$F_t = \rho F_{t-1} + \nu_t \tag{2}$$

$$y_{it} = \alpha_i + \phi_i F_t + \gamma_{i\tau} q_\tau + e_{it}, \tag{3}$$

where τ indexes the second through fourth quarters. The unobserved common factor F_t is constrained to be an AR(1) process by equation 2, which is often called the state equation.²⁵ The individual rates are determined by the so-called observation equations, the functional form of which is shown in equation 3. The observation equations include the common factor, F_t , the group-specific constants, α_i , and seasonal coefficients, $\gamma_{i\tau}$. The component series therefore have their own means and seasonal cycles, though these series are partly determined by common variation. In our system, after accounting for common and seasonal variation, all remaining variation is absorbed by the error term e_{it} . Like the disturbance term in the state equation, ν_t , this error is assumed to be normally distributed.²⁶ This normality assumption allows joint estimation of the system as a state-space model via maximum likelihood and the Kalman filter.²⁷

Table 1 presents the parameter estimates from the DFMs for job-finding rates (column 1) and job-separation rates (column 2).²⁸ The top row of the table shows that the common factors in both systems are strongly autocorrelated, with estimated AR(1) coefficients above 0.90 in both cases.²⁹ For the finding-rate DFM in column 1, the coefficient estimates in the observation equations reflect the similarity in individual finding rates evident in Figure 9. Because the means of the component finding rates are similar, the estimated constant terms in the observation equations are also similar; all but the middle-construction constant

²⁵Independent variables can be entered as right-hand-side variables in the state equation. These additions are especially useful in forecasting applications.

²⁶The variance of ν_t is also normalized to equal one.

²⁷Maximum likelihood estimation is feasible when the number of observation equations is not too large. It is infeasible in large systems now common in empirical macroeconomics, but the abundance of information in these systems means that it is usually appropriate to estimate the factors by PCA, then treat them as data for further analysis (Stock and Watson 2011). For an application of this approach, see Stock and Watson (2012).

²⁸Coefficients on the quarterly dummies from the observation equations are omitted from the table.

²⁹Graphs of the F_t factors themselves can be found in Figure A1, which compares these series with the first factors from the PCAs.

are estimated to be close to 0.25. Additionally, the common factor F_t has a similar effect on the finding rates of all industry–skill groups, with the relevant coefficients clustered in a tight range of 0.010 to 0.014. Finally, the finding-rate DFM does a particularly good job of explaining finding rates for middle-skill workers. The R^2 of the observation equation for middle-other workers, calculated as $\frac{\text{Var}(\hat{y}_{it})}{\text{Var}(y_{it})}$, is 0.95, while those for middle-construction and middle-manufacturing workers are both 0.86. However, the R^2 s for both the high- and low-skill workers are only marginally lower (0.83).³⁰

Coefficient estimates from the separation DFM vary more across observation equations, consistent with the larger differences in nonstandardized separation rates that appear in Figure 9.³¹ The constant terms range from a low of 0.565 for high-skill workers to almost 10 times that for middle-construction workers. The varying scale of separations across groups is also reflected in the large dispersion among F_t coefficients, ranging from 0.016 for the high-skill group to 0.379 for middle-construction. Also, the separation DFM does a better job fitting the data for middle-construction workers than for other workers. The R^2 for the middle-construction equation is 0.88, though all other R^2 s in the model are close to 0.50 or larger. The relatively good fit for middle-construction may reflect a strong cyclical component for separations in that industry–skill group, relative to the importance of idiosyncratic variation.

3.3 Idiosyncratic Variation Implied by Dynamic Factor Models

We now show that idiosyncratic movements in individual job flows implied by the DFMs provide little evidence that the disparate impacts of recent business cycles have been influenced by labor-market polarization. Figure 12 depicts the smoothed DFM errors for the three middle-skill groups.³² Panel A shows the finding and separation errors for middle-skill workers outside of construction and manufacturing, which provide some basis for characterizing polarization’s impact on middle-skill workers after stripping away the effect of working in an industry that is highly responsive to the business cycle. As suggested by the high R^2 from the relevant observation equation, the left graph in this panel shows that middle-other errors in the finding-rate DFM are very small. There do not appear to be any extended periods when middle-other finding rates are substantially higher or lower than expected given the common variation in all finding rates. The same is true for the middle-other errors in the separation DFM, pictured in the right graph (note the difference in vertical scales). From the late 1990s through the mid-2000s, middle-other separation errors tend to be higher than errors in other years, but the difference is not large. Moreover, in the last few quarters of the

³⁰Calculating the R^2 s as $1 - \frac{\text{Var}(\hat{e}_{it})}{\text{Var}(y_{it})}$ gives essentially the same R^2 s as those in the table. Note that the R^2 s also reflect the contribution of the quarterly dummies, whose coefficients are omitted from Table 1.

³¹Coefficient estimates in the observation equations of the separation DFM have been multiplied by 100 for notational convenience.

³²Errors are smoothed by averaging over quarters $t - 3$ to $t + 2$ with equal weights.

sample—a period that includes the slow recovery from the Great Recession—the finding and separation rates for middle-skill workers outside of manufacturing and construction are well described by the common variation in job flows for all workers, as the middle-other errors in both DFMs are near zero.

Idiosyncratic errors for middle-manufacturing and middle-construction workers display more substantial variation. The right graph in Panel B of Figure 12 shows that separation errors for middle-manufacturing workers spike in all recessions, as recessionary job losses for this group are especially acute and thus not completely reflected by changes in the common factor F_t .³³ The pattern in separation errors for middle-manufacturing is also consistent with trend employment growth for this group. In the first half of the sample period, middle-manufacturing separation errors drop sharply after recessions, but the errors remain high between the 2001 and 2007–2009 recessions, a period of generally declining employment in this group.³⁴ Also consistent with this trend is the pattern of middle-manufacturing errors in the finding-rate DFM, which generally remain low in the 2000s as well.

Idiosyncratic movements in job flows for middle-construction workers differ from the middle-manufacturing pattern. Relative to their predicted values, middle-construction finding rates are too low during the 1990–1991 recession but too high during the recession of 2007–2009. This pattern may seem odd at first, since the Great Recession is associated with a housing market collapse. Yet the early 1990s also saw serious housing reversals in some regions of the country, particularly the Northeast and California.³⁵ Because the 1990–1991 recession was mild for the economy as a whole, middle-construction finding rates turned out to be lower than expected given the modest dip in overall job-finding during the early 1990s recession. When the Great Recession occurred nearly two decades later, construction employment collapsed again—but so does employment in every other category. Because of the widespread drop in finding rates, job-finding rates among middle-construction workers turn out to be higher than expected once common variation in these rates is accounted for.

More to the point of this paper, there is a general inverse relationship in both finding and separation errors for middle-skill workers in manufacturing and construction. The 2000s saw middle-manufacturing workers lose jobs at higher than normal rates, while similarly skilled construction workers experienced fewer than expected separations. The opposite is true for job-finding rates, which were relatively low for manufacturing workers and high for construction workers during the same period. This pattern is consistent with recent work by Charles, Hurst, and Notowidigdo (2012), who argue that during the 2000s the housing

³³A potential qualification to this statement is that the spike in middle-manufacturing separations near the 1990–91 recession is substantially smaller than during other recessionary periods.

³⁴See Panel C of Figure 7.

³⁵BLS data show that the four-quarter change in overall construction employment was –9.5 percent in 1991:Q3. This four-quarter rate of decline is larger than any four-quarter decline experienced during the recessions of the early 1980s.

boom masked the effects of ongoing distress in manufacturing on the labor-market outcomes for men without college educations. The manufacturing–construction dichotomy in Figure 12 also highlights the importance of accounting for industry effects when investigating the interaction between polarization and the business cycle, perhaps encouraging us to put more weight on the results for middle-skill workers outside of these two cyclical (and differently trending) groups.

Figure 13 graphs the smoothed DFM errors for high-skill and low-skill workers. The errors for high-skill workers are consistent with the findings from the corresponding PCAs. High-skill finding rates are abnormally low after the 2001 recession and separations are elevated after the two most recent recessions. These patterns are generally the reverse of the pattern for low-skill workers. Also, the small absolute size of the high-skill separation errors has more to do with the low average high-skill separation rate than to a particularly good model fit. Recall that Table 1 shows that the R^2 for the high-skill separation equation is 0.52, the second lowest of any group.

Summing up, idiosyncratic movements in job flows are hard to square with a claim that labor-market polarization is making U.S. recessions worse for middle-skill workers. It is true that middle-skill workers lost the most jobs in the Great Recession, but this pattern generally holds true in all recessions. Outside of the highly cyclical industries of construction and manufacturing, middle-skill job flows are well-explained by the common variation in flows for all workers, even during the Great Recession.³⁶

3.4 Common Finding-Rate Variation and the Vacancy–Unemployment (VU) Ratio

We next relate the estimate of the common factor, F_t , from the finding-rate DFM to the tightness of the U.S. labor market, as measured over time by the ratio of total job vacancies to total unemployment (the VU ratio). This exercise can serve as a check of our industry–skill classifications. The DFMs do not take into account the size of the individual industry–skill classes, so the F_t s they generate are influenced by the classification scheme we choose.³⁷ A close relationship between the rate of hiring in the economy and the VU ratio—a relationship sometimes called the hiring function—is clearly implied by theory and has been demonstrated

³⁶In some unreported work, we preformed some robustness checks to ensure that our results for high-skill workers do not result from the inclusion of manufacturing and construction workers in that category. We found that PCA errors from a model in which high-skill workers did not include manufacturing and construction workers were very similar to errors from a model that included those workers in the high-skill group. See Figure A2.

³⁷Assume that we divided workers into five groups—four whose labor-market outcomes were noncyclical and a fifth where these outcomes were highly cyclical. Cyclical variability would be relegated to idiosyncratic variation of the last group, even if the first four groups included very few workers.

in previous empirical work using aggregate data.³⁸ Consequently, if our classification of workers is reasonable, then we should find a close relationship between the estimated F_t and the VU ratio. A second reason to relate the finding-rate F_t to the VU ratio is to see whether the recent shift in the Beveridge curve, which has been linked to a deterioration in the hiring function, can also be linked to polarization.

To construct the VU ratio, we use the vacancy series produced in Barnichon (2010), which spans the post-1995 decline in newspaper help-wanted advertising (the basis for the traditional vacancy series that started in the early 1950s) as well as the introduction of the BLS's Job Opening and Labor Turnover Survey (JOLTS) in December 2000.³⁹ We regress the common finding factor, F_t , on the VU ratio and a constant:

$$F_t = \alpha + \beta VU_t + \varepsilon_t. \quad (4)$$

We then graph the fitted values and out-of-sample forecasts from this regression, along with the actual values of F_t used in equation (4). A close fit would indicate that our classification system is reasonable and that the VU ratio explains economywide hiring rates well. A divergence would suggest either that the classification system is inappropriate, or that the hiring function has broken down for reasons unrelated to polarization.

Figure 14 presents the results of this exercise. In each of the two panels, the solid line is the common factor from a finding-rate DFM while the dashed lines are fitted values and out-of-sample predictions generated by estimating equation 4 over two different sample periods. The green dashed line in each panel uses a sample that ends in 1985:Q4 (first vertical line) and the red line uses one that ends in 2007:Q3 (second vertical line). Panel A uses the F_t from the baseline model reported earlier, for which the five component finding rates are adjusted for both demographic composition and for time aggregation. The relationship between the VU ratio and F_t is remarkably stable until the onset of the most recent recession. Both regression-based predictions for F_t are close to one another, even though the two predictions are generated from samples of very different lengths. This closeness indicates that the relationship between F_t and the VU ratio does not change much between the end of the first sample (1985:Q4) and the end of the second sample (2007:Q3). When the Great Recession begins, however, both of regression-based predictions imply that hiring should be much higher than it actually is. We also investigate whether the post-2007 gap is related to

³⁸See the theoretical discussion in Pissarides (2000) and the empirical results in Shimer (2005).

³⁹Barnichon (2010) augments the traditional help-wanted data beginning in January 1995 with data from a concurrent survey of online help-wanted ads. We use the traditional newspaper-based vacancy series from January 1976 to December 1994, then use Barnichon's composite series until December 2000, at which point the official vacancy series from the JOLTS program becomes available. This method gives us a consistent series of vacancies, in levels, throughout our sample period. For the denominator of the VU ratio we use the reported number of unemployed persons from the BLS.

our demographic adjustments. Panel B repeats the exercise above with an F_t from a model without demographic adjustments to the component rates. If anything, using unadjusted flows improves the fit before 2007:Q3 but worsens it afterward.

The results in Figure 14 suggest that the finding-rate DFM is reasonably specified, but that the economic environment has changed since the Great Recession in ways that adversely affect workers of all skill types, not just those adversely affected by polarization. Confidence in the model’s specification is supported by the close relationship between job-finding rates and vacancies before the recent recovery. Specifically, the “average” finding rate indicated by the model-based F_t is closely related to the VU ratio before 2007:Q3. The recent discrepancy between job-finding rates and vacancies suggests that the hiring environment has undergone a pronounced adverse change, a finding that is consistent with the recent outward shift in the Beveridge curve. As noted earlier, previous research has linked this shift to a drop in the job-finding rate of unemployed workers, not an increase in job separations. Figure 14 suggests that this shift has little to do with polarization. Unemployed workers in all industry–skill categories are having trouble finding jobs in the recovery from the Great Recession. Indeed, the recent job-finding experiences of different types of workers are well-described by the (low) overall rate of job-finding. A look back at the left panels of Figures 12 and 13 indicates that the errors for each of the five industry–skill groups have been relatively close to zero in recent quarters.

The most important takeaways from the analysis of business-cycle synchronization can be simply stated. Consistent with the findings of other authors, there is strong comovement in the labor-market experiences of different types of workers, even in the recent recovery. We show that this comovement is present even after disaggregating workers by how their skills are likely to be affected by polarization. To be sure, this synchronization is not perfect—a single common factor does not explain all of the variance in skill-specific job flows. But the remaining idiosyncratic variation in job flows indicates that high-skill workers favored by polarization have also suffered during recent recessions. Moreover, the recent breakdown in the hiring function suggests that the adverse Beveridge curve shift has occurred alongside low rates of job-finding for workers throughout the unemployment pool. In the next section of the paper, we study the reallocation of unemployed workers with cyclical synchronization in mind. How do unemployed workers—middle-skill workers in particular—reallocate themselves in recessions, given the long-term outlooks that polarization implies?

4 Reallocation of Unemployed Workers Over the Business Cycle

4.1 Outflows from Unemployment: The Importance of Exits to Labor-Force Nonparticipation

The most basic question we can ask about unemployed workers is what type of new jobs they take, and whether the probability of taking jobs in different skill classes changes over the course of the business cycle. Figure 15 displays the composition of outflows from unemployment for each of the five industry–skill groups. These outflows are taken from the CPS data with no demographic adjustments. Panel A depicts the transitions of workers who leave unemployment for a new job, either high-skill, low-skill, or middle-skill.⁴⁰ As we would expect, unemployed high-skill workers tend to take high-skill jobs, and low-skill workers generally take low-skill jobs. But middle-skill workers in particular tend to stay in their same skill class. Middle-manufacturing workers leaving unemployment take middle-skill jobs more than 75 percent of the time, while the corresponding percentage for middle-construction workers is even higher.

There are some changes in skill-switching probabilities over time, however. After the 1990–1991 recession, unemployed high-skill workers finding work became progressively less likely to take middle-skill jobs. A similar trend begins for low-skill workers around the onset of the 2001 recession. For unemployed middle-skill workers getting jobs, the trend away from middle-skill reemployment is more modest and begins earlier. This trend is obviously consistent with polarization of job opportunities away from middle-skill employment. But there does not appear to be much business cycle variation in the rate at which workers exiting unemployment switch skill classes. The probability of skill-switching for middle-skill workers appears particularly immune to business cycle fluctuations. Combined with the high percentage of middle-skill workers who stay in middle-skill work, the lack of cyclical variation in skill-switching for the middle-skill unemployed suggests that it is unattractive or infeasible for unemployed middle-skill workers to find high-skill or low-skill jobs at any phase of the business cycle.

Panel B of Figure 15 broadens the set of possible transitions out of unemployment to include exits to nonparticipation. Now some cyclical variation emerges. The black lines show that in the second half of the sample, a progressively higher share of workers who exit unemployment do so by leaving the labor force. Part of this upward movement may reflect the workforce’s demographic composition. As it ages, a higher share of the unemployed may choose to retire early by moving from unemployment (U) to nonparticipation (N). Just as importantly, there is also some business cycle variation in the share of unemployment exits ending in nonparticipation. For most industry–skill groups, the share of exits to nonpar-

⁴⁰For convenience, we do not disaggregate new middle-skill jobs by industry.

ticipation rises during later recessions and recoveries, including the Great Recession and its aftermath.

The rising importance of job flows involving nonparticipation has been noted in previous work that does not address polarization directly.⁴¹ But the rising share of U-to-N exits, combined with the cyclical nature of this share in the latter part of the sample period, suggests that a study of long-run labor-market trends and cyclical reallocation should incorporate the potential exit for unemployed workers to nonparticipation. Unemployed middle-skill workers facing a declining demand for their skills may be more likely to respond to recessions by dropping out of the labor force, either to enjoy leisure or to augment their human capital. Participation decisions by high- and low-skill workers may also be affected by the decline in middle-skill opportunities. As Figure 15 shows, these workers also take middle-skill jobs when unemployed, so high- and low-skill workers may also find nonparticipation more attractive as middle-skill opportunities decline.

One must be careful, however, when linking the rising share of unemployment spells that end in labor-market nonparticipation to optimizing behavior on the part of unemployed workers. There is a simple mechanical explanation for cyclical movements in the probability of moving from unemployment to nonparticipation *conditional* on exiting unemployment in a given month. A worker can exit from unemployment in one of two ways: leaving the labor force or finding a job. Consider what happens when the job-finding rate falls: the probability of moving from U-to-N conditional on exiting unemployment is likely to rise—simply because the probability of exiting unemployment in the only other possible way has declined. To see how this matters in practice, consider the top panel of Figure 16, which depicts the economywide share of unemployment spells that end in nonparticipation. The green line uses the CPS microdata from this paper and the red line uses published BLS data on gross worker flows, available after January 1990.⁴² Both lines show the steady upward march and greater cyclical nature in this share after the mid-1980s. The bottom panel shows the unconditional probability of moving from U-to-N; that is, the ratio of U-to-N movements to the total number of unemployed workers in the previous month. As other authors have pointed out this probability falls during recessions. Indeed, even though the top panel shows that the share of unemployment spells that end in nonparticipation is near its post-1976 high, the bottom panel shows that the unconditional probability of moving from U-to-N is near a

⁴¹Elsby, Hobijn, and Sahin (2012) note as a general matter that the cyclical nature of these flows is important for understanding the time-series behavior of the unemployment rate. Kudlyak and Schwartzman (2012) point out that flows involving nonparticipation are especially important in accounting for movements in unemployment around time of the Great Recession.

⁴²We would not expect these two lines to be identical, because the CPS microdata we use does not include new entrants (see footnote 7). Our data also omit a small number of unemployed workers whose occupation cannot be classified or is missing. Nevertheless, the seasonal and cyclical agreement in the two series in both panels of the figure is gratifying.

20-year-low. These two facts are explained, of course, by the recent collapse in the overall job-finding rate.

4.2 Modeling Unconditional Transition Probabilities for Unemployed Workers

To determine whether unemployed workers—middle-skill workers in particular—have responded differently to recent business cycles relative to their earlier behavior, we need to measure the cyclical responses of their unconditional transition probabilities controlling for potentially confounding factors.⁴³ We estimate multinomial logit models that ask how these reallocation probabilities change in response to the economy’s overall job-finding rate, controlling for worker-level demographic characteristics (which may be relevant for trends as well as the cycle) as well as unemployment duration. Formally, consider an unemployed worker j from industry–skill group i who can either stay unemployed (U), exit to employment (E), or exit to nonparticipation (N). With staying in U is normalized as the baseline choice, the unconditional probabilities of transitioning from unemployment to either E or N are:

$$\Pr(E_{ij,t+1}|U_{ijt}) = \frac{\exp(\Gamma'_{iE}X_{ijt})}{1 + \exp(\Gamma'_{iE}X_{ijt}) + \exp(\Gamma'_{iN}X_{ijt})} \quad \text{and} \quad (5)$$

$$\Pr(N_{ij,t+1}|U_{ijt}) = \frac{\exp(\Gamma'_{iN}X_{ijt})}{1 + \exp(\Gamma'_{iE}X_{ijt}) + \exp(\Gamma'_{iN}X_{ijt})}, \quad (6)$$

where X_{ijt} is a vector of regressors and the Γ s are parameters. For notational convenience, this representation does not distinguish between exits to high-skill, middle-skill, or low-skill employment, but this breakdown is made in the empirical work below. To capture worker-level attributes, the X vector includes data on educational attainment, gender, marital status, and age. Reported unemployment duration is also entered as a set of dummy variables, with the omitted dummy corresponding to workers who have zero to one weeks of unemployment duration.⁴⁴ To capture the state of the business cycle, we include the the common factor F_t from the finding-rate DFM.⁴⁵ For all unemployed workers except the high-skilled, we

⁴³Most recently, Elsbj, Hobijn, and Sahin (2012) has argued that recessions change the demographic makeup of unemployment pool by causing large numbers of workers with strong labor-force attachments to become unemployed. They find that controlling for demographics substantially reduces the downward recessionary movements in U-to-N probabilities apparent in the bottom panel of Figure 16. Their paper uses CPS data matched over 12 months, so they are able to control for the unemployed worker’s labor-force status 12 months before.

⁴⁴Age is specified as a cubic polynomial in the worker’s true age minus 35 years, so that all three age terms equals zero when the worker is 35 years old. Dummies are entered in the regression for nonwhite, female, and married. The three included education categories are less-than-high-school, some college, and college graduate. The female dummy is interacted with the nonwhite and married dummy as well as the cubic in age – 35. The duration dummies correspond to 2, 3, 4, 5–8, 9–13, 14–17, 18–21, 22–26, 27–51, 52, 53–78, 79–98, 99, and > 99 weeks of duration. We also exclude workers who are more than 70 years old from the estimation sample and include quarterly dummies (the first quarter is omitted).

⁴⁵For notational convenience this factor is standardized to have a mean of zero and a variance of one.

report transition probabilities for a baseline worker who is 35 years old, male, unmarried, and white, and who has a high school diploma but no further education. For the high-skilled probabilities, we construct a baseline worker with a college degree. For each of the five types of unemployed workers, we run two multinomial logits, one on a sample from 1976:Q1–1985:Q4 and another on a sample from 1986:Q1–2012:Q1.

Before presenting the results, a caveat is in order. Flows between unemployment and nonparticipation are known to be measured with error. Recently, Elsby, Hobijn, and Sahin (2012) have studied this issue by matching monthly CPS records for the same individuals across several months. Just as earlier authors have found, recent work by Elsby, Hobijn, and Sahin (2012) notes that workers often move back and forth between nonparticipation and unemployment, suggesting that the true economic meaning of those two states is blurred for many in the labor market. When the authors “iron out” those CPS histories by removing potentially spurious transitions, they generate corrected series that are close to the corrections implied by Abowd and Zellner (1985). Yet the authors also find that these adjustments do not have large effects on the cyclical nature of U-to-N flows (though the adjustments do affect the cyclical nature of N-to-U flows). The main purpose of the multinomial logit regressions below is to see whether transition decisions change over time. If the nature of measurement error has remained constant, it is not immediately clear how this error would affect our results. That said, we do enter a post-1994 dummy in the X vector where appropriate. This dummy controls for any level shift in the probabilities after the CPS redesign, including the upward shift in the unconditional U-to-N probability that appears after the redesign in Figure 16.

4.3 Results of Multinomial Logit Models

Figure 17 displays the results of the logits for the middle-skill unemployed. The graphs on the left side of the figure correspond to the early sample period and those on the right are from the later sample. The blue bars in each graph are the relevant probabilities when the common finding-rate factor, F_t , is one standard deviation above its mean; the gray bars correspond to probabilities when F_t is one standard deviation below its mean. Consistent with Elsby, Hobijn, and Sahin (2012), the figure shows that controlling for observables matters when estimating the transitions of unemployed workers to nonparticipation. The lower panel of Figure 16 showed that the average U-to-N probability for all workers falls in recessions, but in this figure the blue and gray bars corresponding to the nonparticipation exit are essentially constant in both sample periods.⁴⁶

What transition margins are affected for middle-skill workers when the overall job-finding

⁴⁶See Figure A3 for graphs of the same probabilities implied by models without demographic or duration controls. In that figure, a drop in the finding rate appears to have a negative effect on U-to-N transitions, consistent with the raw time series for U-to-N movements in the lower panel of Figure 16.

rate declines? Not surprisingly, middle-skill workers are less likely to exit unemployment for middle-skill jobs when the finding rate falls. A drop in the overall finding rate has little effect on the probability that a baseline middle-skill worker will find a job in another skill class, as these probabilities are low to begin with.⁴⁷ Because the unconditional probabilities must sum to one, the lack of a cyclical response for most potential exits means that the decline in middle-skill re-employment probabilities caused by a drop in the overall job-finding rate is reflected nearly one-for-one by an increase in the probability that a middle-skill worker remains unemployed.

For comparison, we also estimated multinomial transition logits for unemployed high- and low-skill workers. The results were similar, in that lower finding rates did not increase the unconditional probability of moving from unemployment to nonparticipation.⁴⁸ As Figure 15 showed, unemployed workers from the poles of the skill distribution sometimes take middle-skill jobs as well as jobs in their own skill classes. The multinomial logits for these workers indicate that both of these potential exits from unemployment are reduced when overall finding rates fall, but changes in the prevalence of the nonparticipation margin are minor.

What happens to unemployed workers after they transition to nonparticipation? In this paper, we match workers only across consecutive months in the CPS, and nonparticipating CPS respondents are not asked questions about their most recent occupations. Consequently, we are unable to determine whether “middle-skill nonparticipators” are more likely to emerge from nonparticipation with higher-skill jobs, as we would expect if these workers chose to use their time out of the labor market to augment their skills. We did, however, investigate outflows from nonparticipation over the sample period disaggregated by gender. To be sure, there were cyclical movements in the fraction of nonparticipating workers flowing into unemployment, consistent with the idea stressed in Elsby, Hobijn, and Sahin (2012) that workers flow back and forth between these two states. But when we look only at movements from nonparticipation to employment, there is little evidence that these flows are tilted toward higher-skill jobs in the wake of recessions.⁴⁹

Taken together, the results of the multinomial logits suggest that the rising and increasingly cyclical share of exits from unemployment to nonparticipation does not result from middle-skill workers—or any other type of worker, for that matter—choosing to transition to nonparticipation as an optimal response to lower job-finding rates. Rather, the time-series behavior of reallocations of unemployed workers to nonparticipation seems driven by a combination of demographic influences and the simple mechanical effect of being unable to exit

⁴⁷The probabilities in Figure 17 correspond to a 35-year-old baseline worker with a high-school degree, and choosing another baseline worker would affect the likelihood of a middle-skill transition out of unemployment into a low-skill or high-skill job. Setting the baseline equal to an “average” worker would still generate few transitions to high- or low-skill jobs; recall the low average rates for these transitions from Figure 15.

⁴⁸See Figure A4.

⁴⁹See Figure A5.

unemployment by finding a job.

5 Conclusions

This paper was motivated by a potentially troubling pattern that emerged during the Great Recession and the subsequent slow recovery. Figure 1 showed that from 2009 to 2011, the worst employment growth was experienced by middle-skill workers, the same group that has been adversely affected by long-term polarization trends in the U.S. labor market. Coupled with the outward shift in the Beveridge curve, Figure 1 appeared to suggest that unemployment today has an important structural component. If so, we would expect to find some class of workers in high demand, but previous empirical work using large research datasets has failed to do so. And efforts to measure the degree of structural mismatch directly have concluded that it is not a large and persistent feature of today’s labor market.⁵⁰

The study of high-frequency individual-level data over many years helps reconcile these results. It is true that recent middle-skill job losses were the most severe, but this is a common pattern due in part to the disproportionate fraction of middle-skill jobs located in manufacturing and construction. Outside of those cyclical industries, a dynamic factor model shows that middle-skill job flows are almost exactly what we would expect them to be given the poor state of the overall U.S. labor market. The results also provide context for the small high-skill employment losses in Figure 1. While also small in an absolute sense, the job-separation rate for employed high-skill workers is now large relative to the rate that would be expected given the experiences of high-skill workers in other business cycles. Along these lines, these results buttress other research that finds workers of all types are having trouble finding jobs; this pattern remains consistent even after workers are segmented on the basis of how their previous occupations line up with labor-market polarization. Finally, bringing job vacancies into the analysis suggests that polarization is not behind the recent shift in the Beveridge curve. Common variation in job-finding rates across industry–skill groups closely followed the vacancy–unemployment ratio until recently, a finding that is consistent with the idea that job-matching efficiency has declined for many types of workers, not just those in the middle of the skill distribution.⁵¹

Historical context is also useful for relating labor-market polarization to the recent aca-

⁵⁰Dickens and Triest (2012), Sahin et al. (2012), and Lazear and Spletzer (2012) find that some measures of industry-level or occupational mismatch rose during the Great Recession but have since retreated.

⁵¹Faberman and Mazumder (2012) perform a skill-based breakdown of employment growth using annual CPS data from 2007 through 2011. Consistent with Figure 1, they find that job losses are most severe among occupations in the middle part of the skill distribution. However, the Conference Board’s monthly help-wanted data, which begins in May 2006, shows that middle-skill job vacancies have generally outpaced those from other groups during the recent recovery. The authors conclude that these results provide little support for the mismatch hypothesis, though the combination of high middle-skill job postings with low middle-skill employment growth may indicate some degree of mismatch *within* the middle-skill group.

demographic literature on recessions as reallocations. In particular, a provocative paper by Jaimovich and Siu (2012) contends that polarization is responsible for the jobless recoveries experienced after the 1990-1991, 2001, and 2007-2009 recessions. In its most basic form, the paper’s search-and-matching model of the labor market includes two types of occupations, high-skill and “routine,” the latter of which is equivalent to middle-skill jobs. To reflect polarization, the productivity of routine jobs is assumed to decline gradually relative to high-skill productivity. This trend encourages routine workers to switch into high-skill jobs, but doing so is difficult. Routine workers must first separate from their old jobs and then search in a so-called switching market. Once they find jobs in this market, former middle-skill workers augment their skills while working and soon become qualified for high-skill work. Unfortunately for potential switchers, jobs in the switching market are scarce. The authors assume that it is costly for firms to post vacancies in the switching market, so few openings are available there. The high posting cost therefore lengthens the unemployment spells of workers in the switching market, and thus approximates the myriad real-world factors likely to lead to long jobless spells among former middle-skill workers hoping to find new jobs.

A key goal of Jaimovich and Siu (2012) is to show how the reallocation of routine workers can be concentrated in recessions, even when the productivity gap between high-skill and routine matches grows smoothly over time. This concentration requires an aggregate, business-cycle shock that affects productivity in both high-skill and routine matches at the same time. When a negative aggregate shock causes a recession, routine workers are more likely to enter the switching market because the opportunity cost of leaving their old jobs has fallen. After the recession, the economy’s unemployment rate remains high for an extended period, as large numbers of former routine workers endure long unemployment spells in the switching market where jobs are hard to find. Polarization therefore contributes to jobless recoveries by creating a mass of poorly skilled workers hoping to improve their careers by leaving their jobs when productivity falls.

To our knowledge, the Jaimovich–Siu model is the first to link polarization with jobless recoveries in a plausible and internally consistent way. But the model’s specific predictions for the reallocation of middle-skill workers are in some tension with the reallocation patterns identified above. Unemployed middle-skill workers rarely reallocate to either high- or low-skill jobs, and this low reallocation rate does not vary much over the business cycle. In recent years, the unemployment spells of middle-skill workers are increasingly likely to end in nonparticipation—as are the spells of other workers—and Jaimovich and Siu (2012) notes that transitions between nonparticipation and unemployment could be related to the process of skill acquisition at the heart of their paper.⁵² But the multinomial logits showed that there

⁵²The authors write that “[m]ore broadly, a jobless recovery involves a slow transition into employment from any source of non-employment. Hence, we view the middle-skill worker’s move from the [middle-skill] market, to the [switching] market, to eventual employment also as a stand-in for temporary spells of labor

is no substantial increase in the unconditional probability that an unemployed worker of any type will leave the labor force when the aggregate job finding rate falls. This is true both in the early part of the sample, when jobless recoveries are not feature of the data, as well as in later years when jobless recoveries are the norm.⁵³

There are also some apparent discrepancies in the predictions for group-specific job flows. Their model implies that the only endogenous separations occur in middle-skill matches, as high-skill matches are insulated from productivity fluctuations by search costs and fully adjusting real wages. On average, this prediction is validated by the evidence presented above, because the high-skill separation rate is lower than the corresponding rates for other workers. But high-skill separations are cyclical nonetheless, and during the (jobless) recoveries from the 2001 and 2007-2009 recessions, the high-skill separation rate was high relative to its historical norm. As for job-finding rates, these flows are similar across industry-skill groups throughout the sample period, though they were particularly low for high-skill workers in the early 2000s. These findings do not necessarily rule out a role for polarization in generating jobless recoveries. But they suggest that the factors linking polarization to jobless recoveries extend beyond the job-matching frictions central to the Jaimovich-Siu model.

force non-participation that may arise from time spent relocating or retraining in order to switch occupations” (Jaimovich and Siu 2012, p. 29, insertions added).

⁵³Some other work on the reallocation of middle-skill workers focuses on the type of worker that reallocates, without an analysis of how the reallocation relates to the business cycle. Cortes (2012) uses annual and biannual data from the Panel Survey of Income Dynamics to show that middle-skill workers switching to high-skill positions tend to be better-quality workers, as predicted by a model with positive sorting on ability.

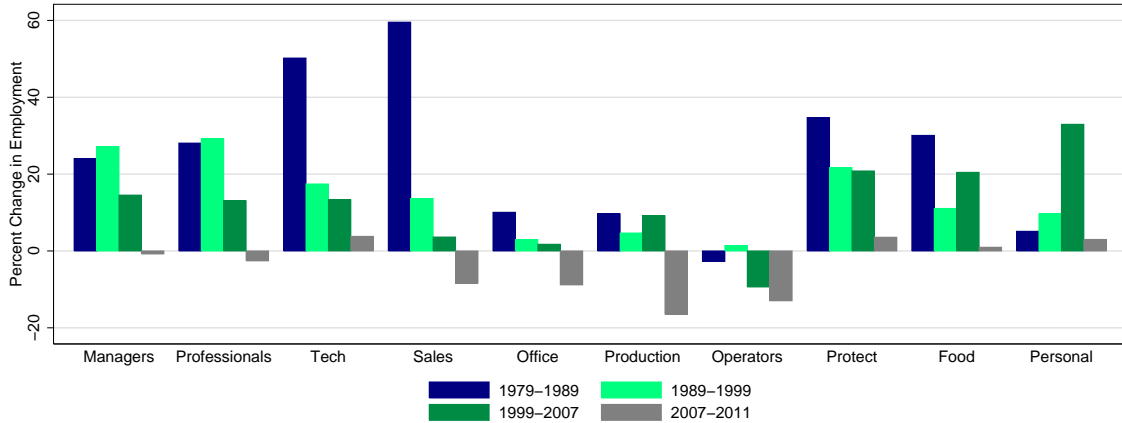


Figure 1. EMPLOYMENT GROWTH BY OCCUPATION: 1979–2011. Note: This figure is an updated version of a chart that appears as Figure 3 in Autor (2010, p. 9) and Figure 12 in Acemoglu and Autor (2011, p. 1073). Occupations depicted (from left to right) are managers; professionals; technicians; sales; office and administration; production, craft, and repair; operators, fabricators, and laborers; protective services; food preparation and building and grounds cleaning; and personal care and personal services. The original charts in Autor (2010) and Acemoglu and Autor (2011) use a subset of individual-level records from the Current Population Survey (CPS) that also includes wage information. This chart uses the complete set of monthly records from the CPS, but the resulting differences are minor. Source: Authors’ calculations using CPS microdata.

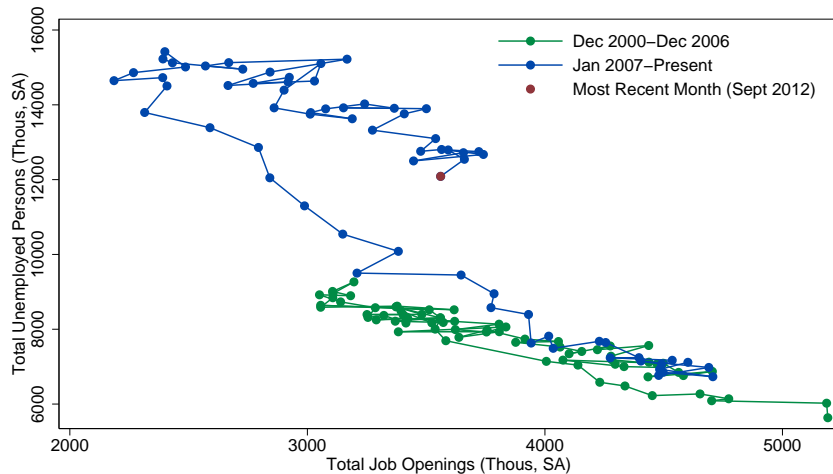


Figure 2. THE BEVERIDGE CURVE: DECEMBER 2000–SEPTEMBER 2012. Note: The recent outward shift in the Beveridge Curve has been linked to lower job-finding rates among unemployed workers. Source: Bureau of Labor Statistics.

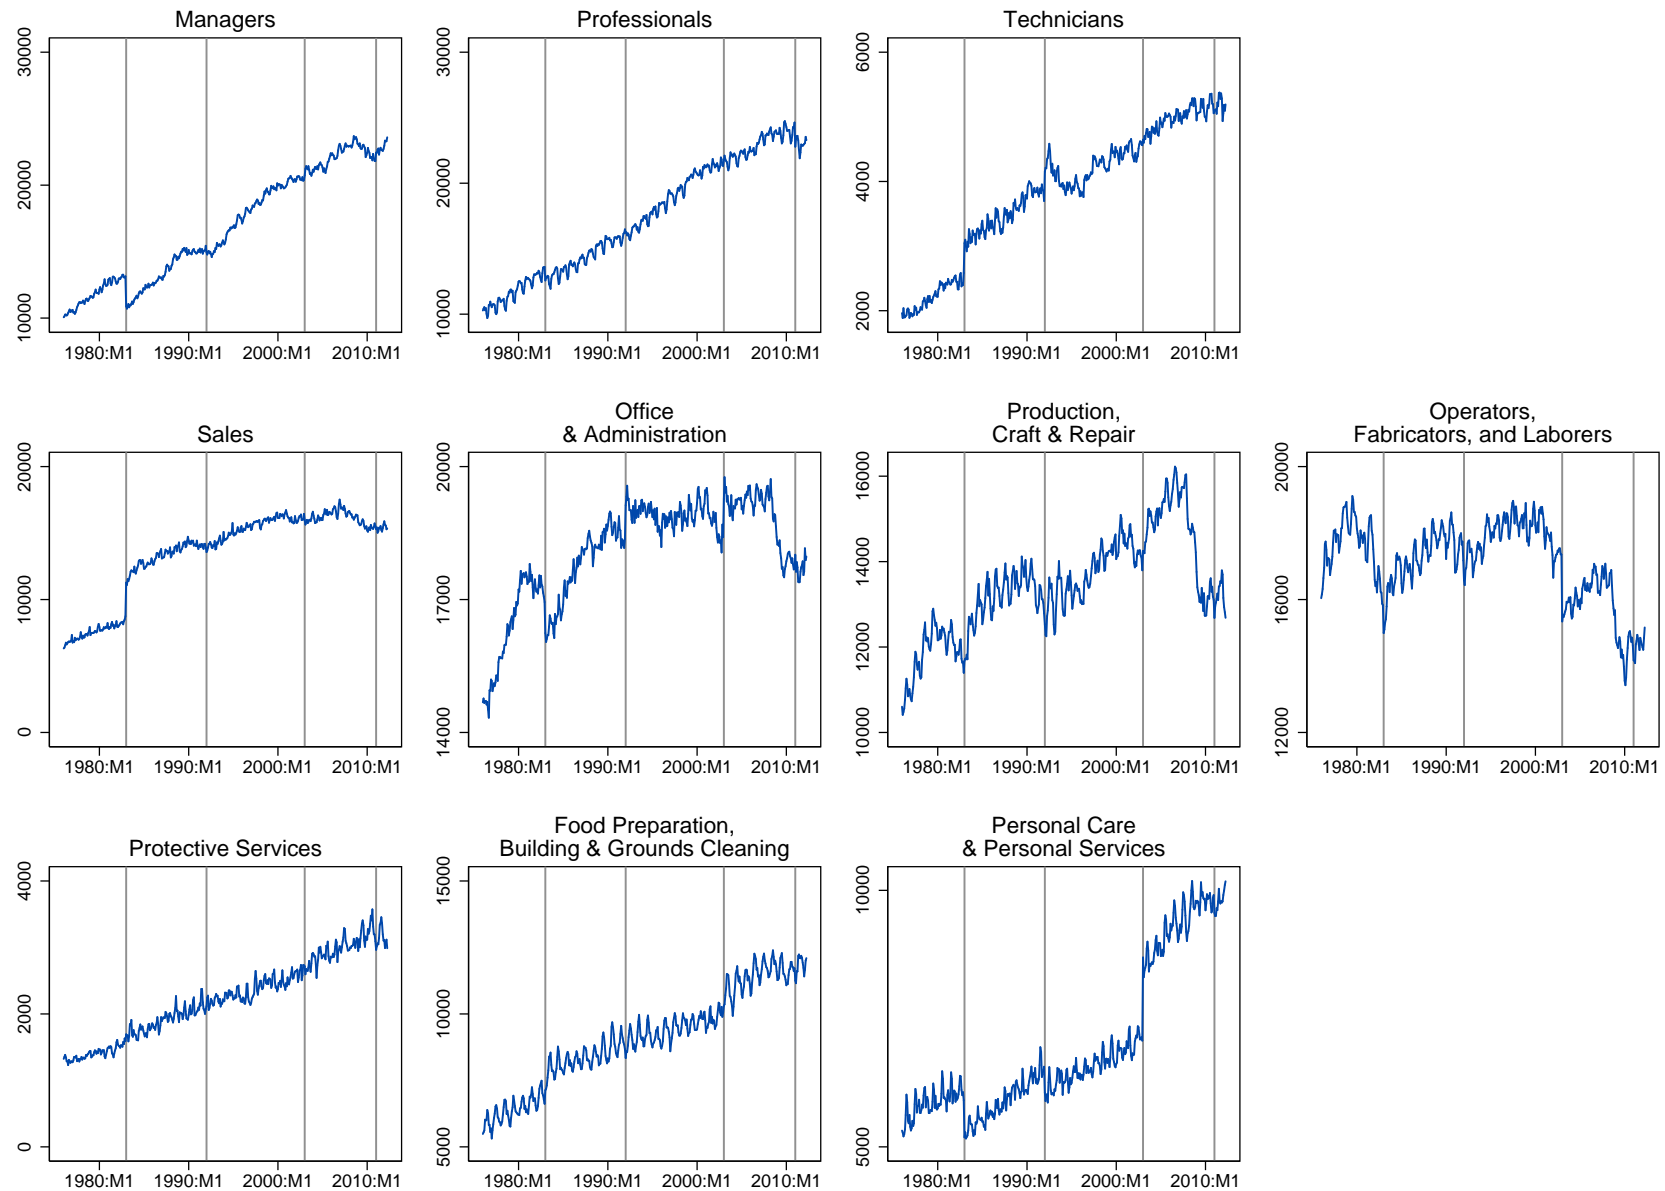


Figure 3. MONTHLY EMPLOYMENT LEVELS BY OCCUPATION, IN THOUSANDS: JANUARY 1976–APRIL 2012. Note: Top row: High-skill employment; second row: middle-skill employment; third row: low-skill employment. Vertical lines correspond to changes in the occupational coding system of the CPS. Data are not seasonally adjusted and vertical scales differ across graphs. Source: Authors' calculations using CPS microdata.

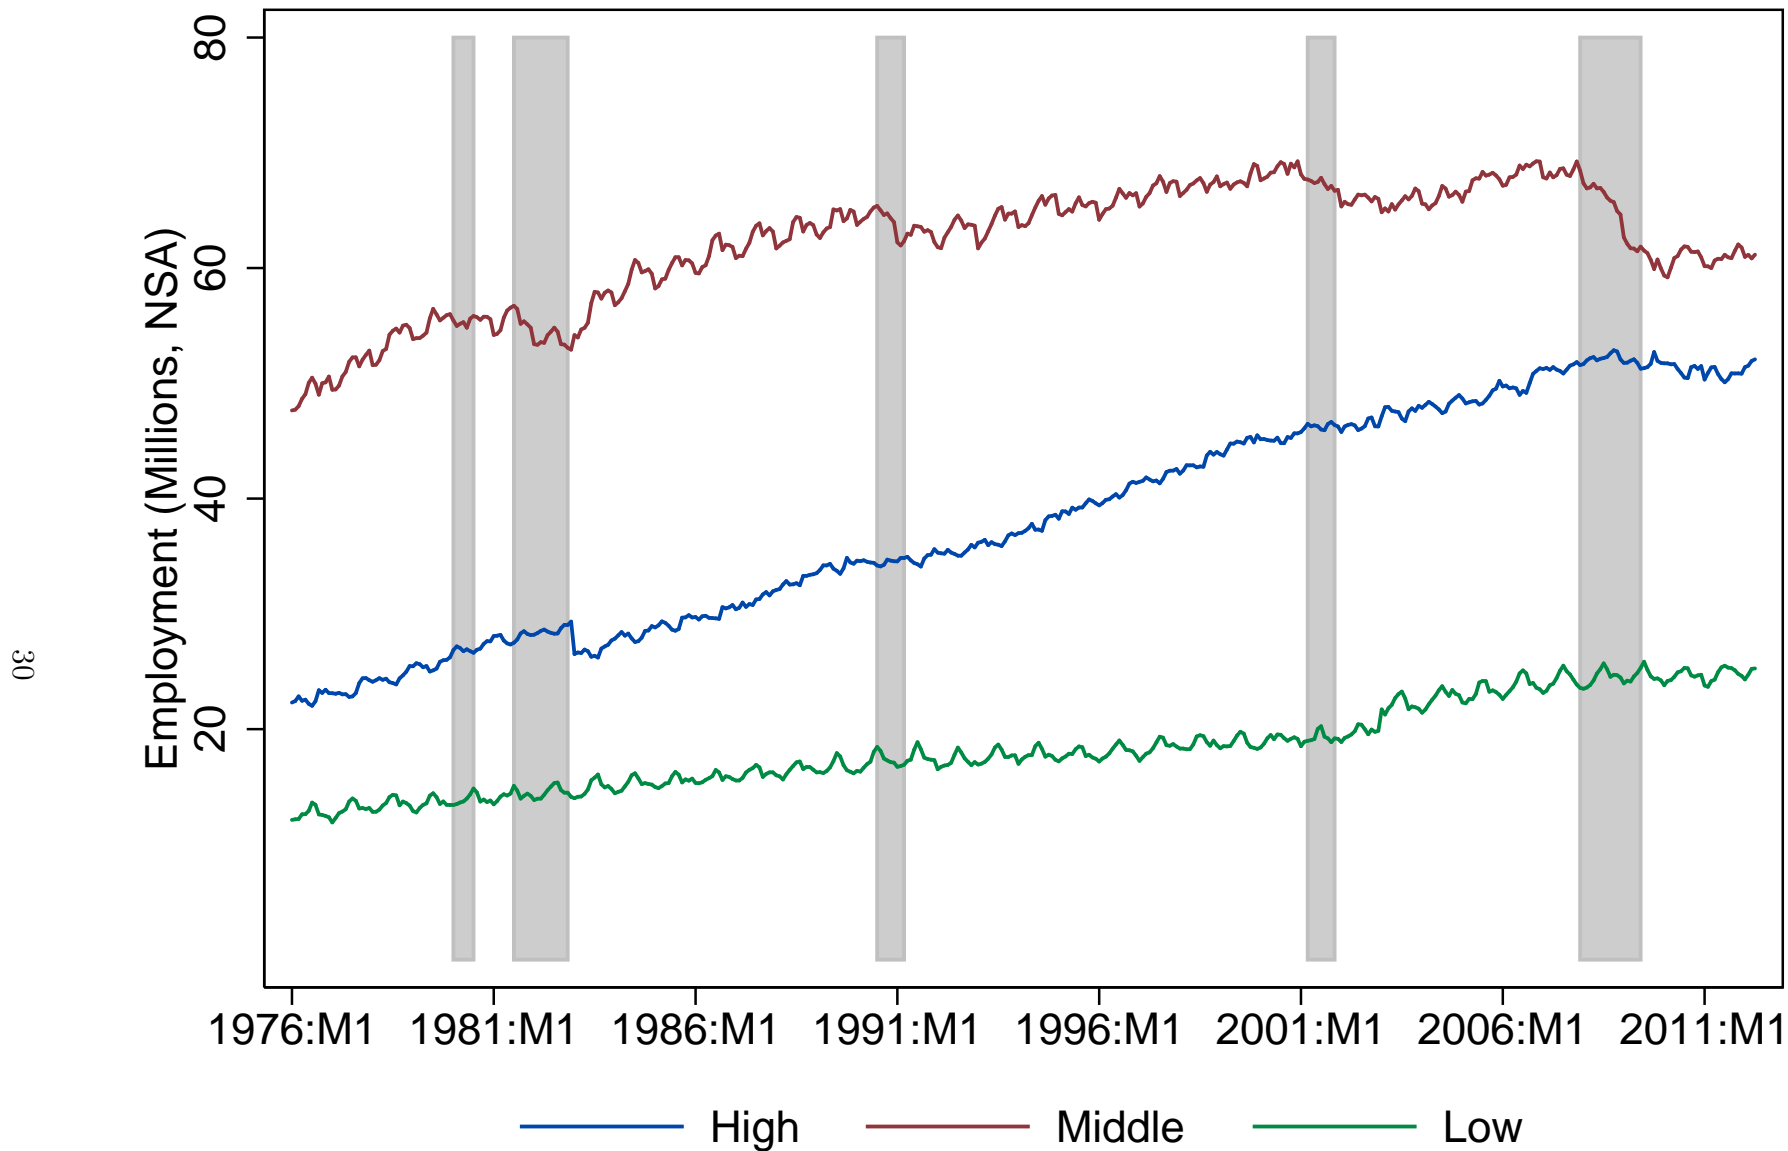


Figure 4. MONTHLY EMPLOYMENT FOR THREE BROAD SKILL CLASSES: JANUARY 1976–APRIL 2012. Note: See Figure 3 for the occupations included in the high-, middle-, and low-skill classes. Recessions are shaded. Data are not seasonally adjusted. Source: Authors’ calculations using CPS microdata.

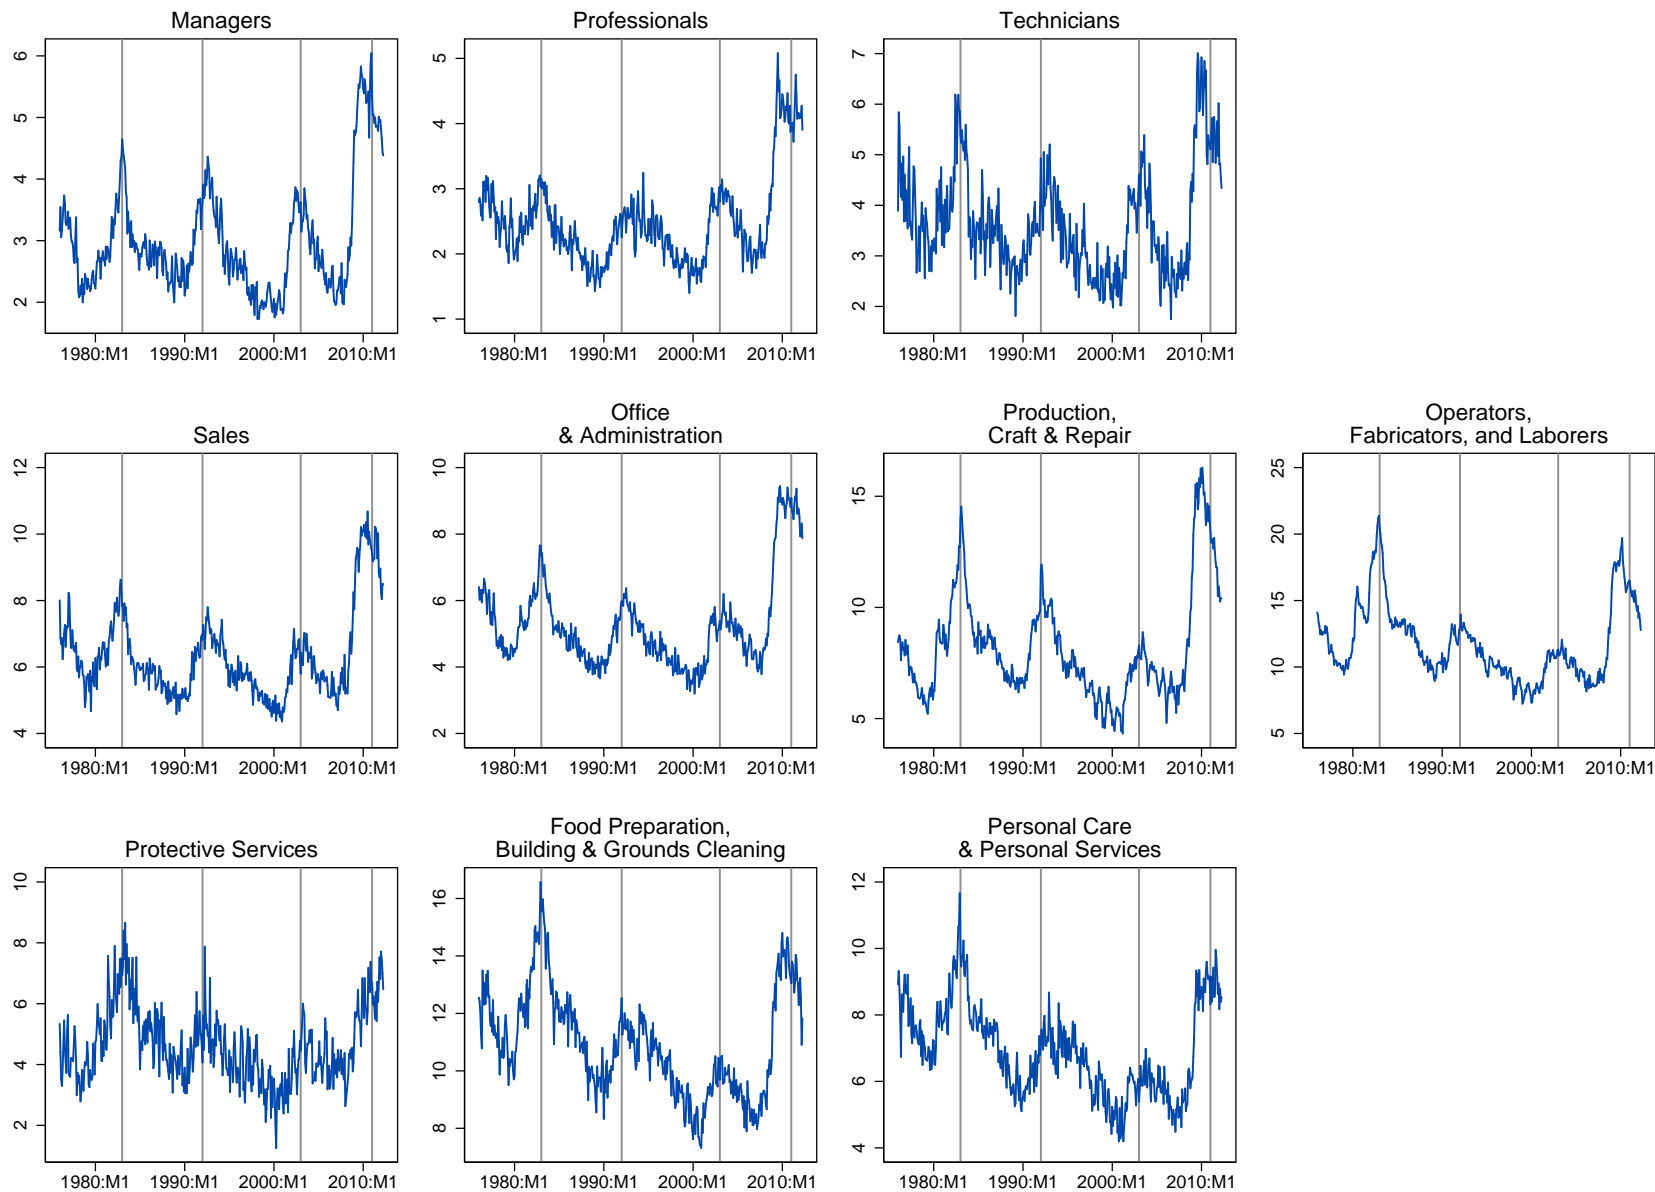


Figure 5. MONTHLY UNEMPLOYMENT RATES BY OCCUPATION, AS PERCENT OF LABOR FORCE: JANUARY 1976–APRIL 2012. Note: Data are seasonally adjusted with a linear regression-based procedure. Vertical lines correspond to changes in the occupational coding system of the CPS. Source: Authors' calculations using CPS microdata.

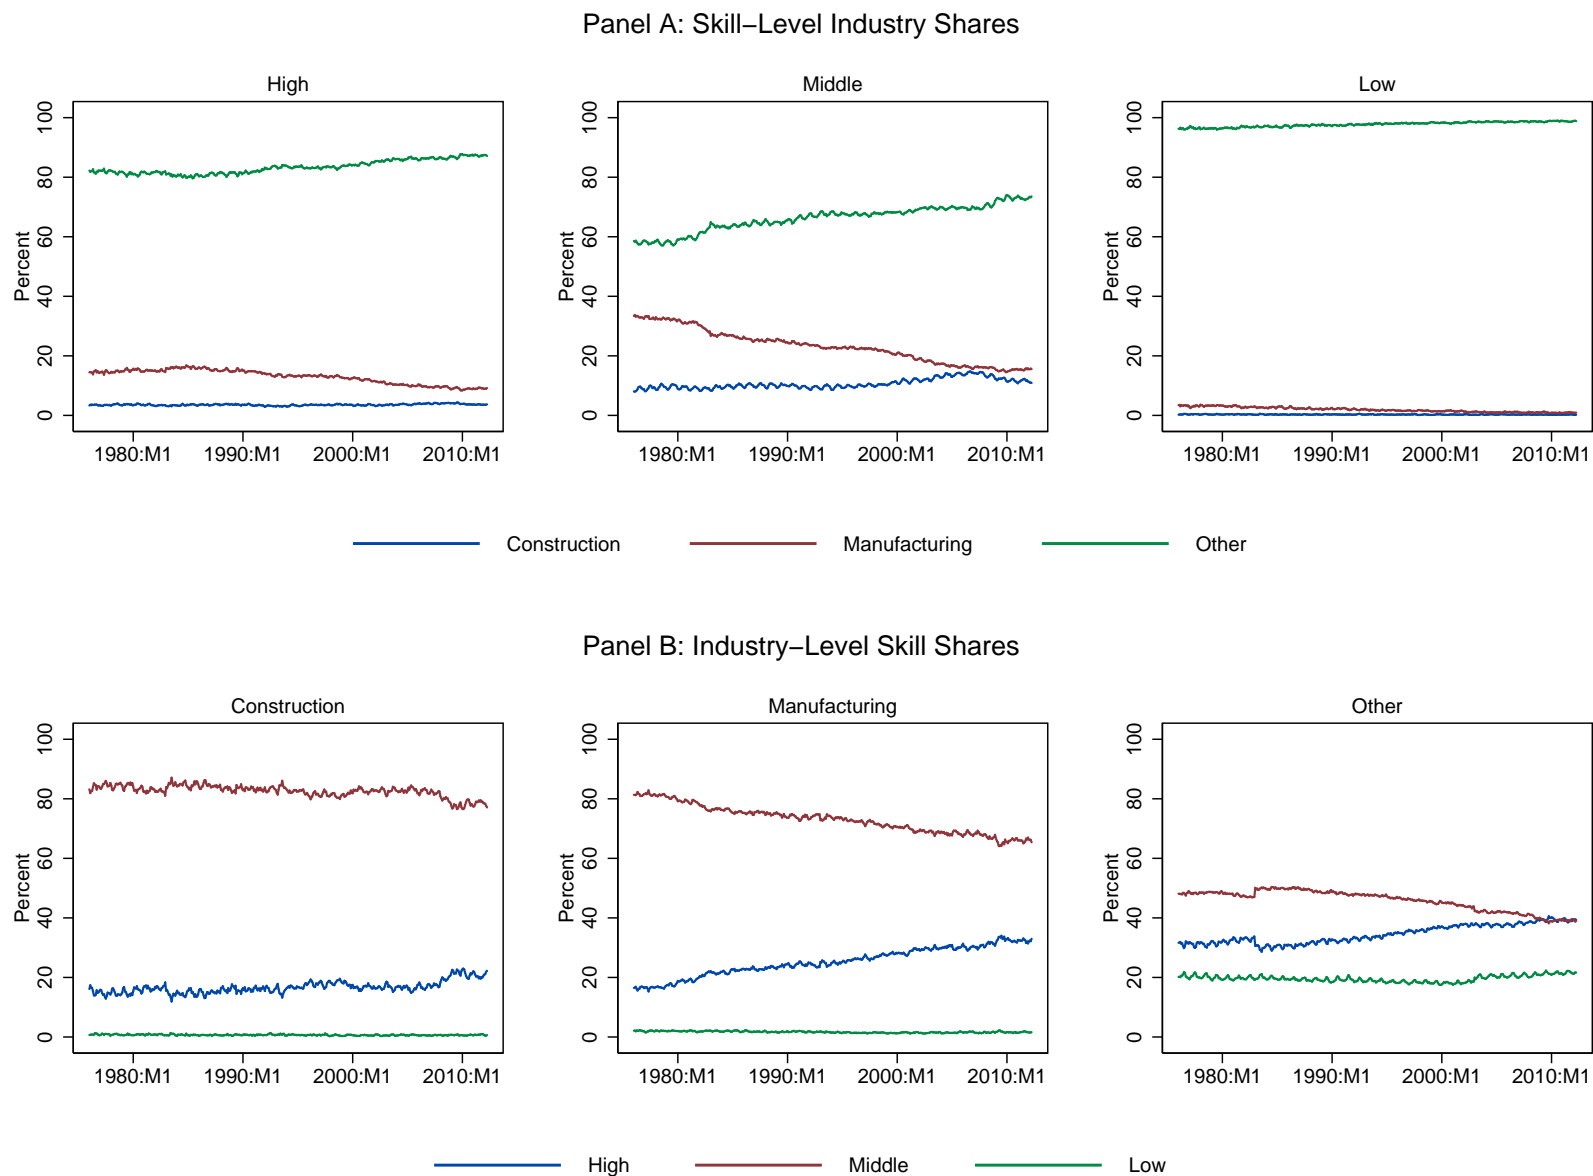


Figure 6. SKILL-LEVEL AND INDUSTRY-LEVEL EMPLOYMENT SHARES: JANUARY 1976–APRIL 2012. Note: “Other” refers to all industries outside of manufacturing and construction. High-skill, middle-skill, and low-skill occupations are as defined in Figure 3. Source: Authors’ calculations using CPS microdata

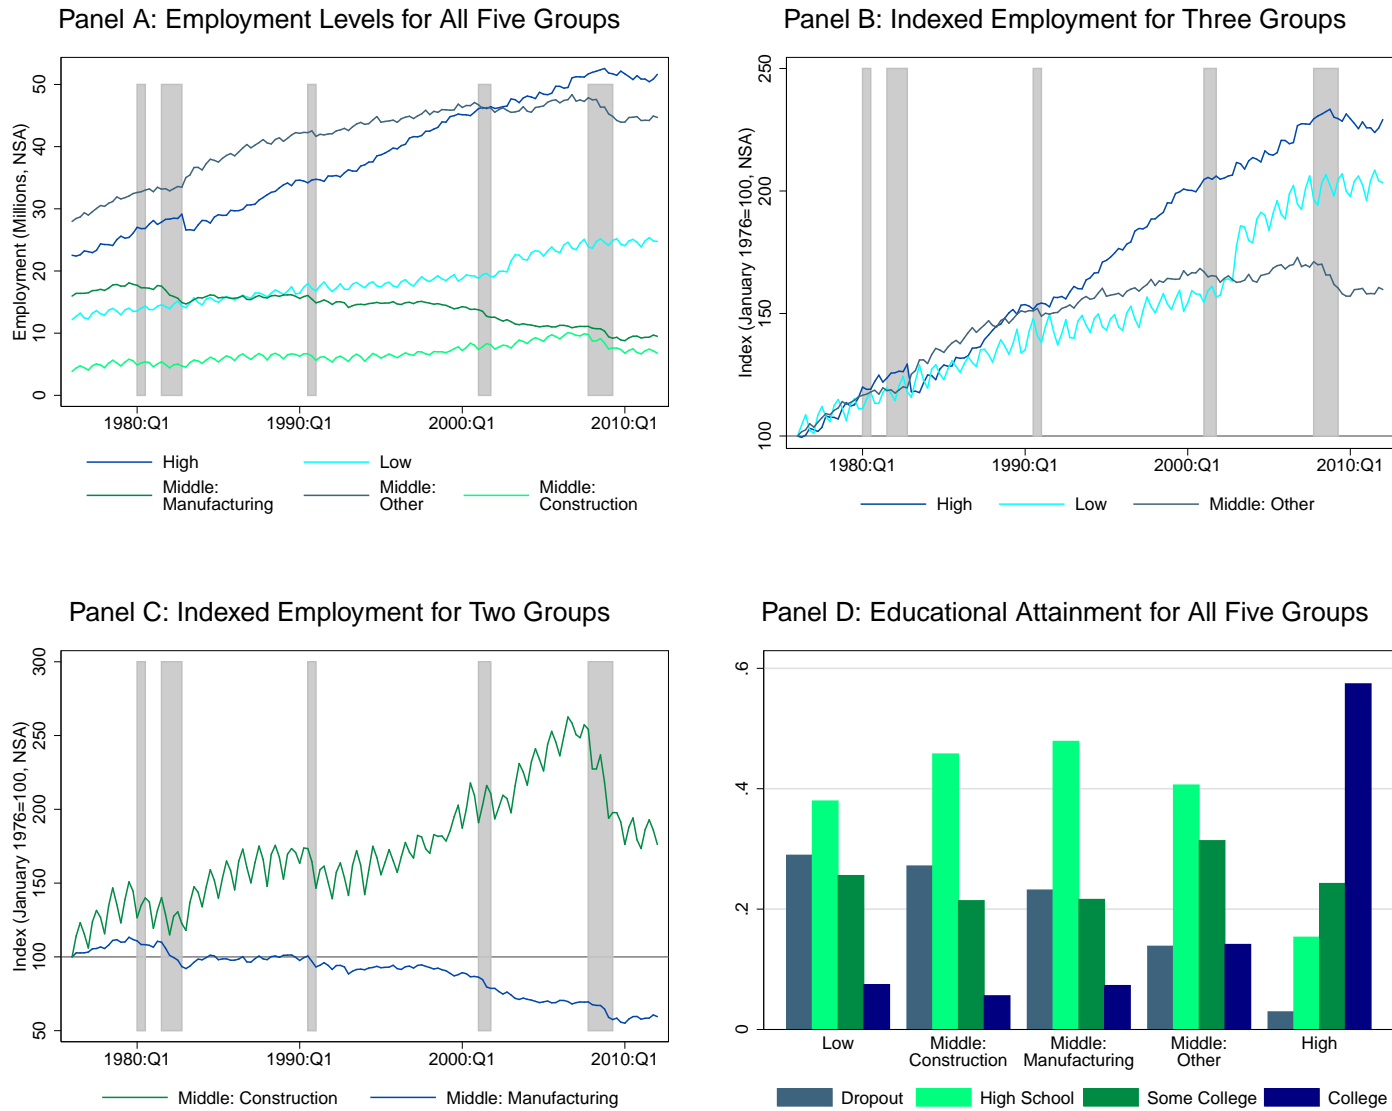


Figure 7. EMPLOYMENT LEVELS AND EDUCATIONAL ATTAINMENT FOR FIVE INDUSTRY-SKILL GROUPS: 1976:Q1–2012:Q1. Note: Employment levels are quarterly averages of monthly data and are not seasonally adjusted. Recessions are shaded. Source: Authors’ calculations using CPS microdata.

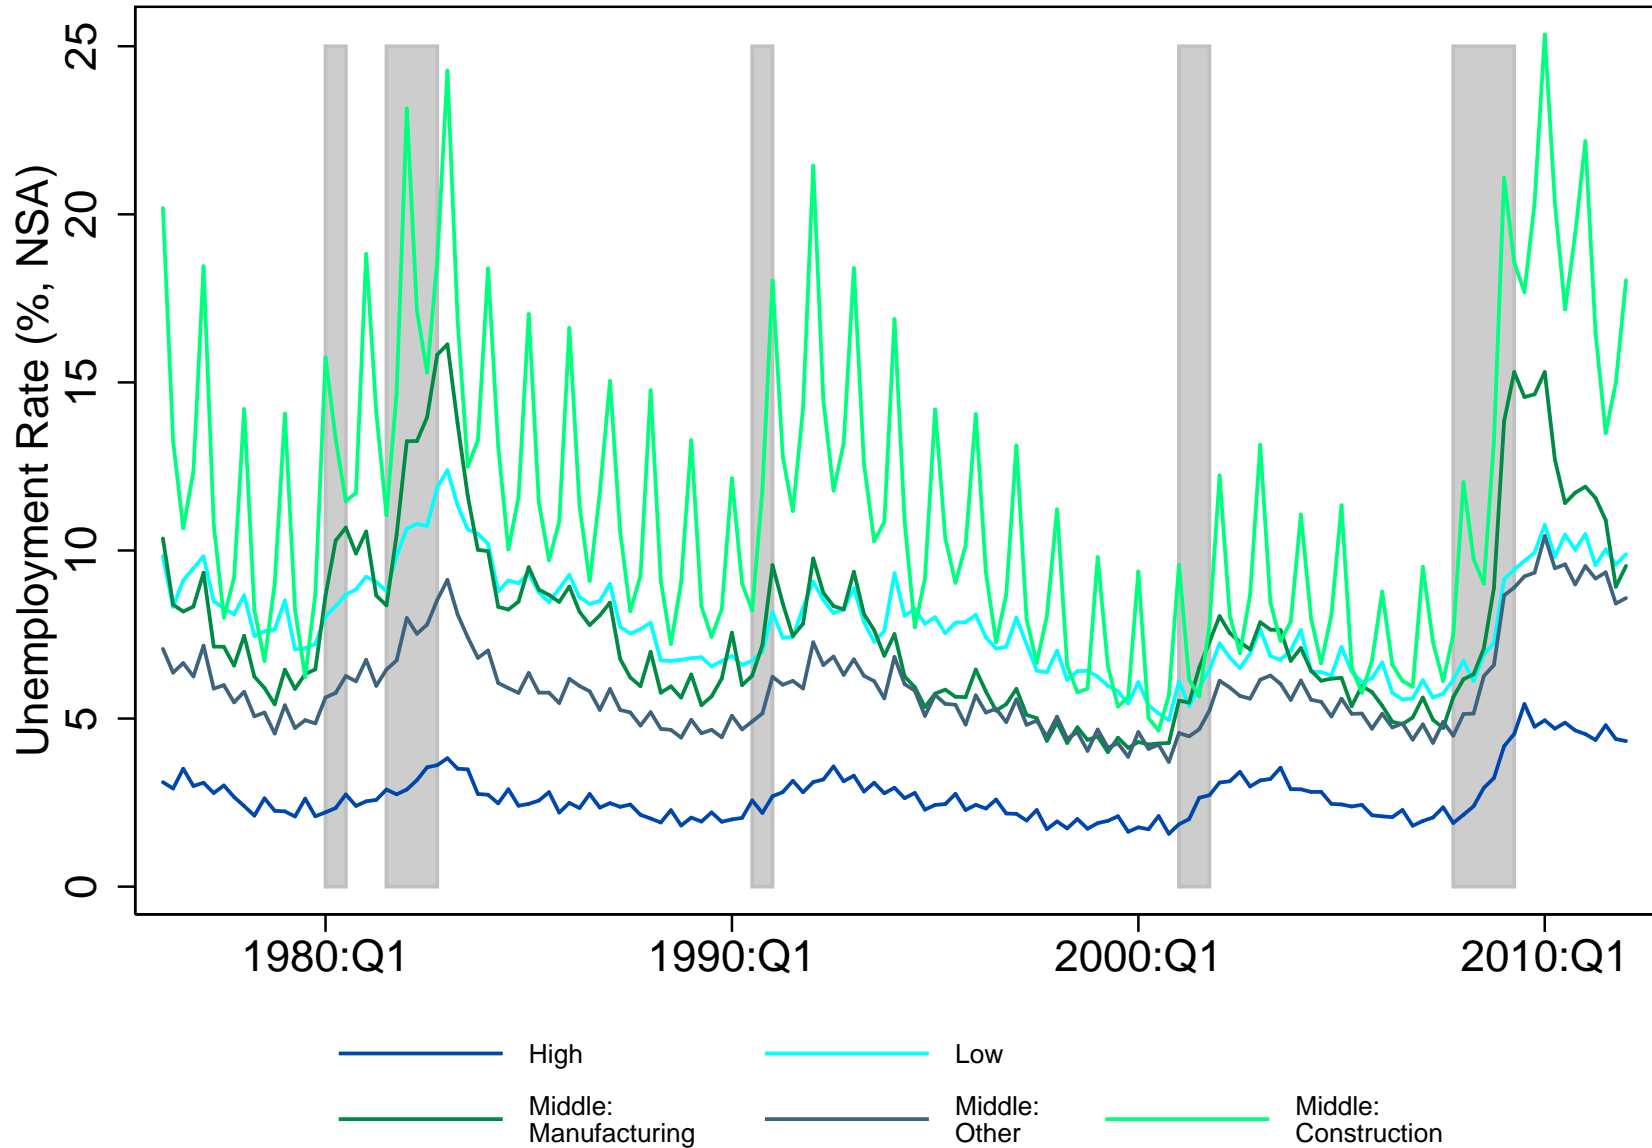


Figure 8. UNEMPLOYMENT RATES FOR FIVE INDUSTRY-SKILL GROUPS: 1976:Q1–2012:Q1. Note: Series are quarterly averages of monthly data. The unemployment rates depicted here are adjusted neither for changing demographic composition nor for seasonal effects. Recessions are shaded. Source: Authors' calculations using CPS microdata.

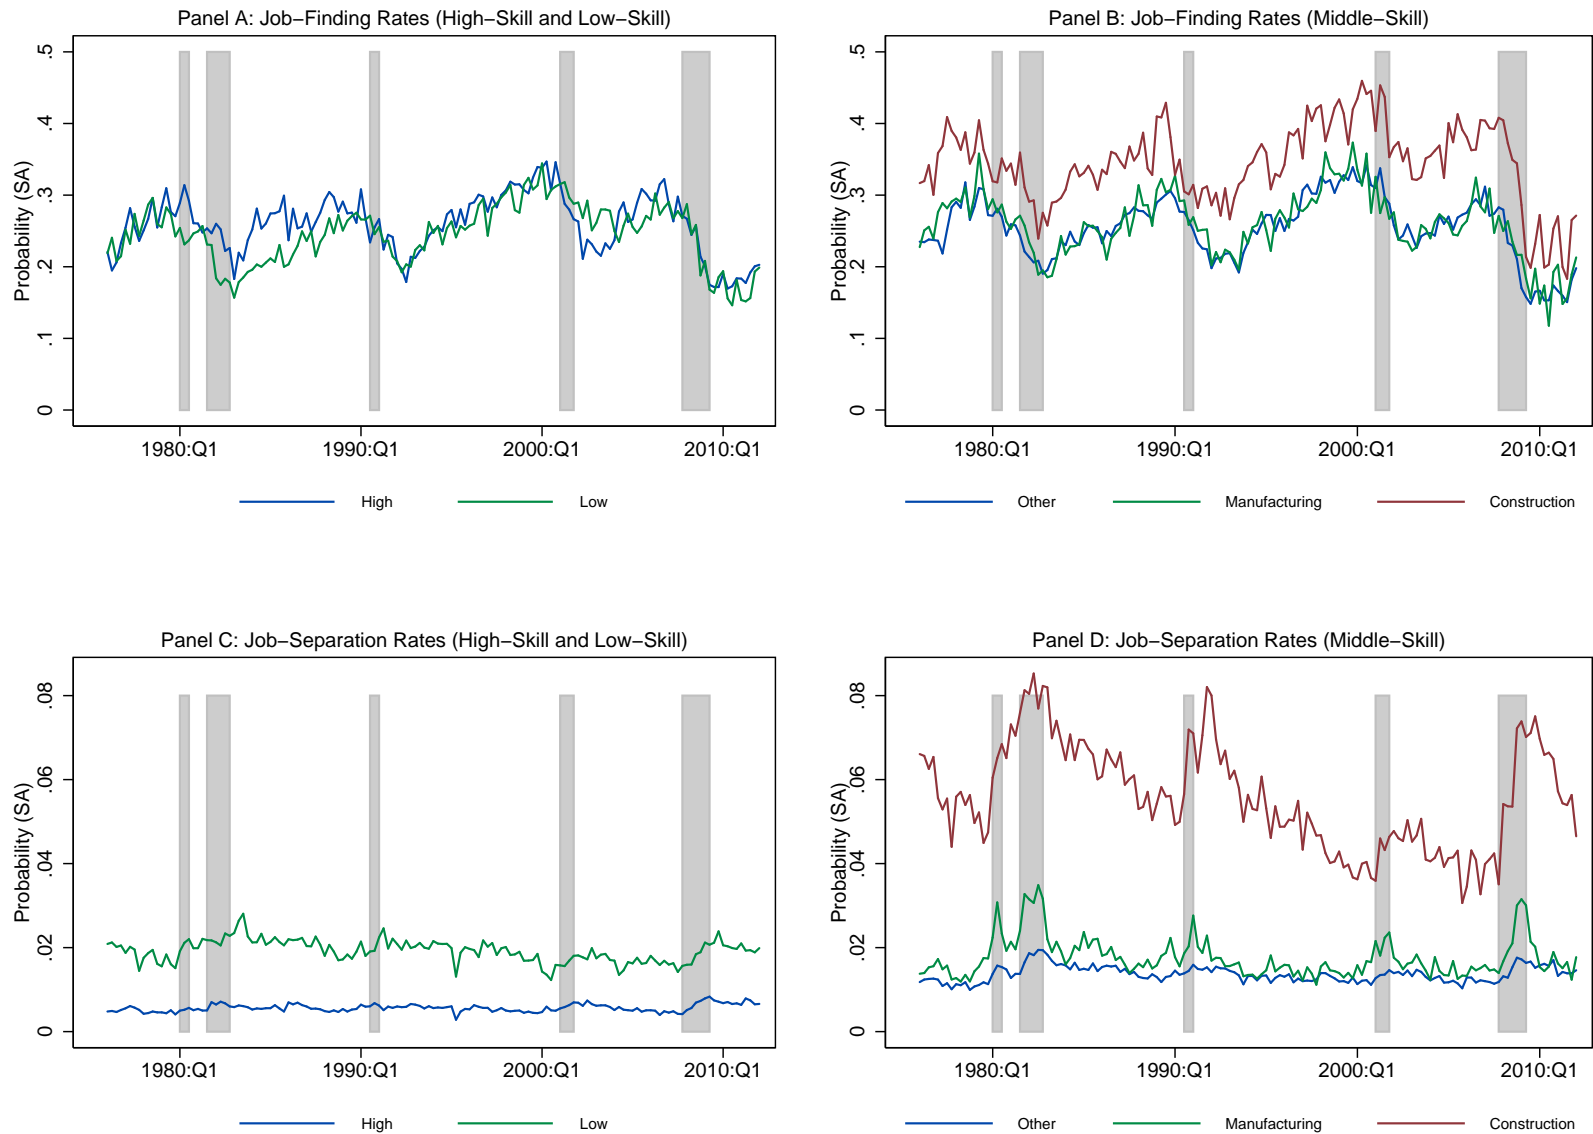


Figure 9. JOB-FINDING AND JOB-SEPARATION RATES FOR FIVE INDUSTRY-SKILL GROUPS: 1976:Q1–2012:Q1. Note: Rates are quarterly averages of monthly data. A job-finding rates is the probability that a worker unemployed in one month will be employed in the next month’s CPS survey. Job-separation rates are defined analogously for employed workers. Both finding and separation rates are adjusted for changes in demographic composition and for time-aggregation. Rates are also seasonally adjusted with a linear regression-based procedure. Recessions are shaded. Source: Authors’ calculations using CPS microdata.

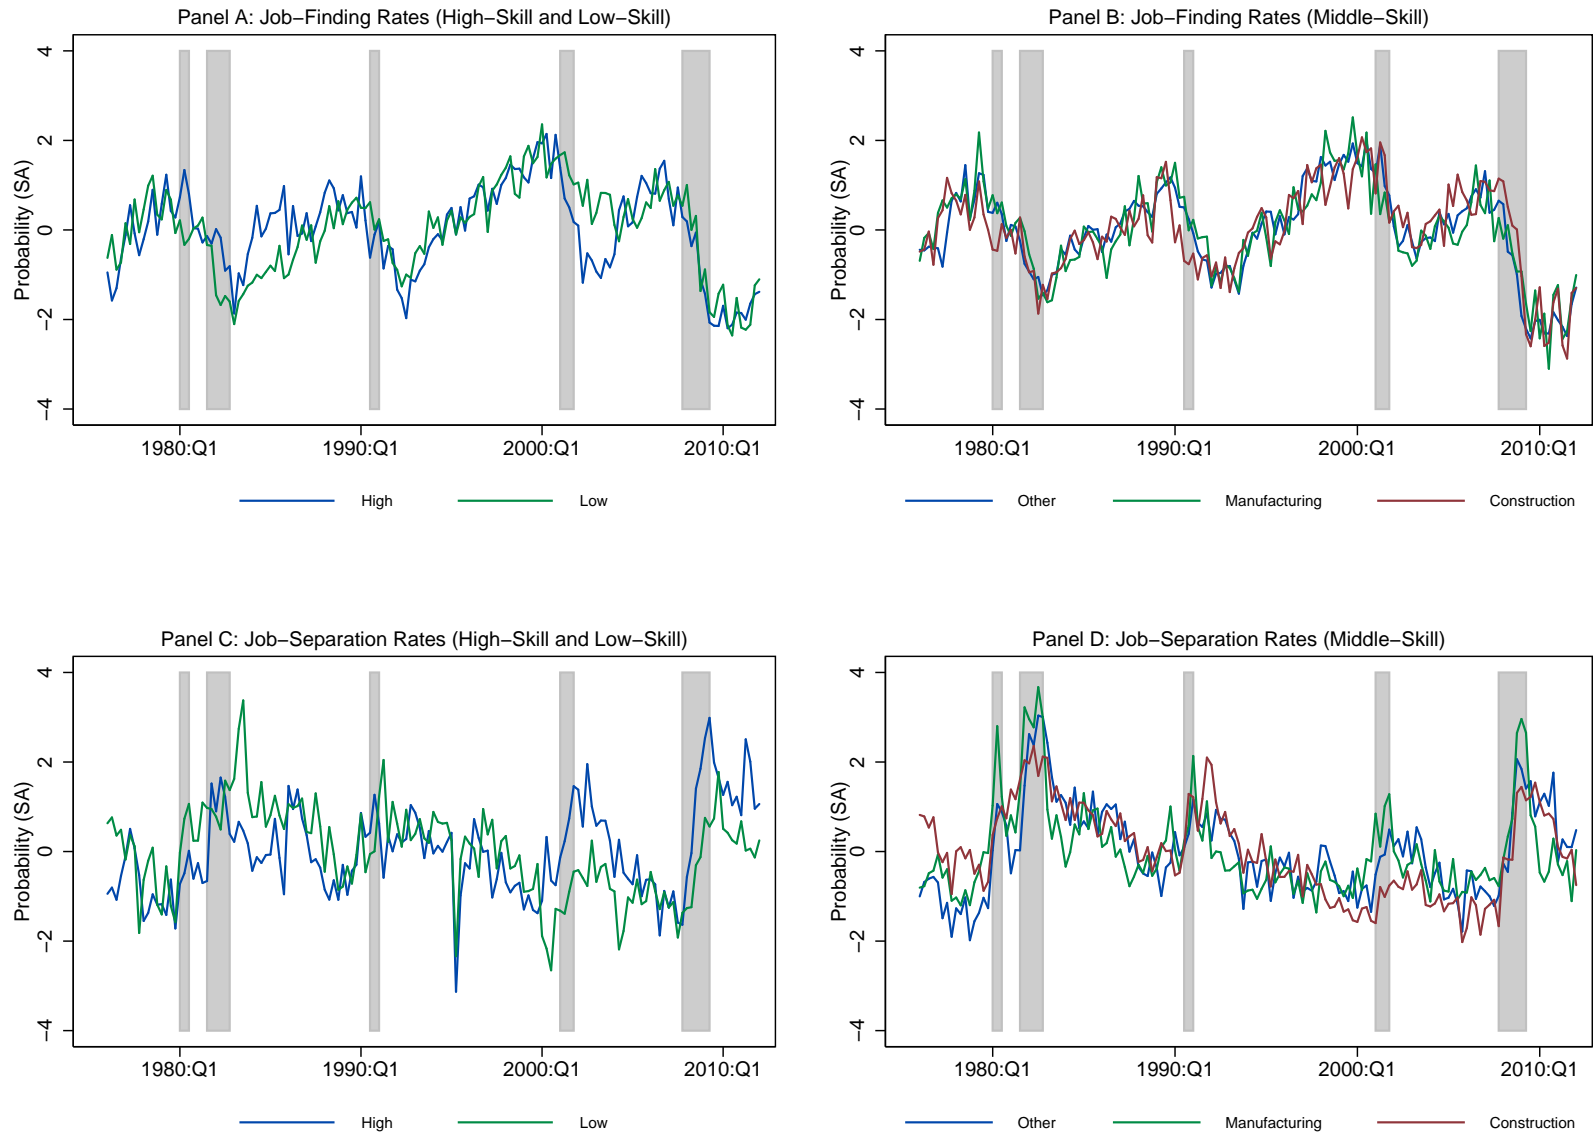
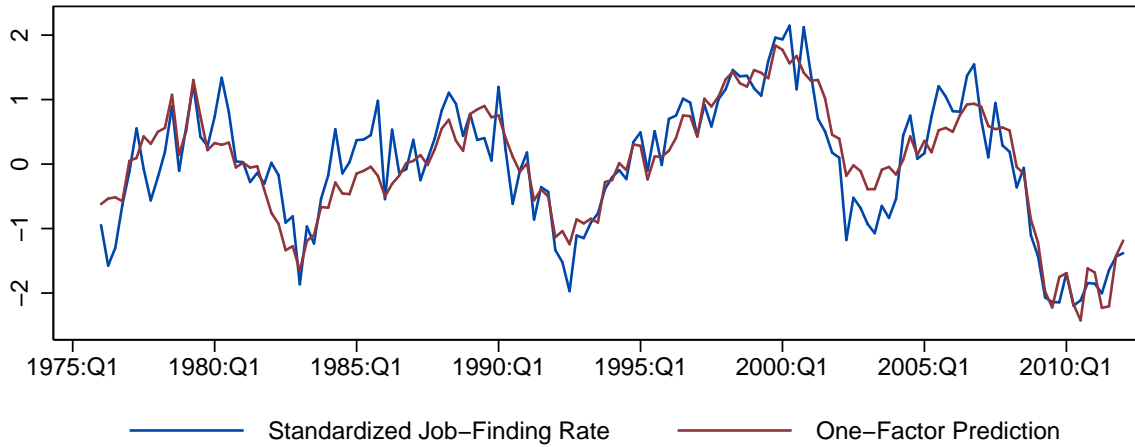


Figure 10. STANDARDIZED JOB-FINDING AND JOB-SEPARATION RATES FOR FIVE INDUSTRY-SKILL GROUPS: 1976:Q1–2012:Q1. Note: Rates are those depicted in Figure 9, standardized so that each series has a mean of zero and a variance of one. Source: Authors' calculations using CPS microdata.

Panel A: Cumulative Share of Total Variance in Standardized Rates Explained by PCA Factors

	Unemployment Rates	Job-Finding Rates	Job-Separation Rates
Factor 1	0.902	0.864	0.669
Factor 2	0.956	0.916	0.808
Factor 3	0.980	0.956	0.927
Factor 4	0.994	0.984	0.969
Factor 5	1.000	1.000	1.000

Panel B: High-Skill Job-Finding Rate and One-Factor PCA Prediction



Panel C: High-Skill Job-Separation Rate and One-Factor PCA Prediction

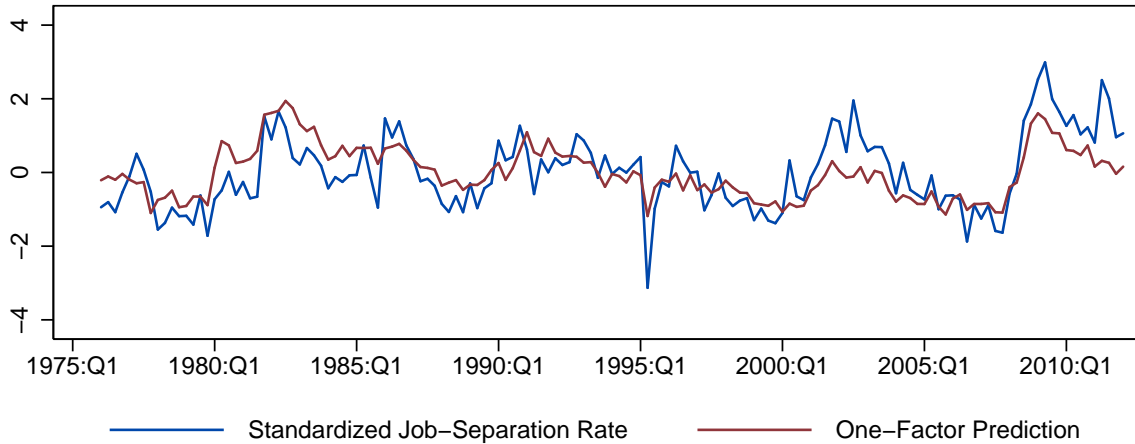
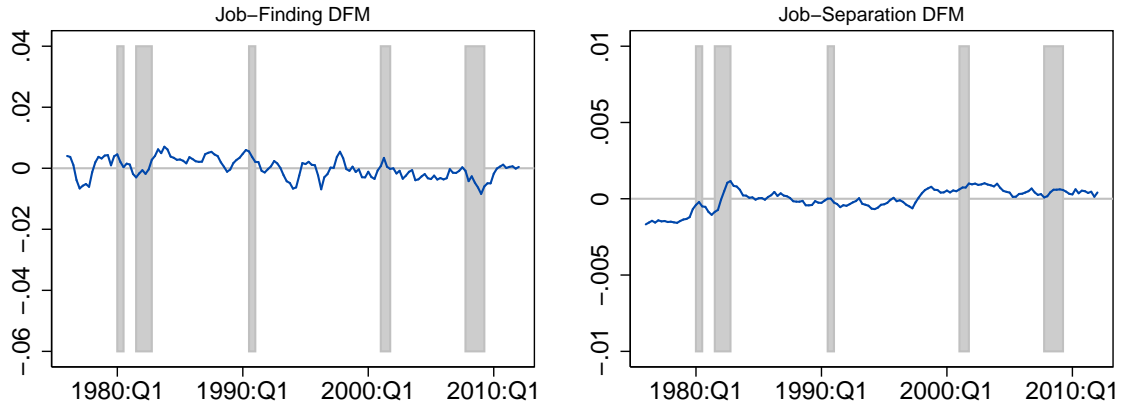


Figure 11. SELECTED RESULTS FROM PRINCIPAL COMPONENTS ANALYSIS (PCA) OF LABOR-MARKET DATA FROM FIVE INDUSTRY-SKILL GROUPS. Note: The top panel shows the fraction of total variance explained by adding additional factors to the PCAs for unemployment, job-finding, and job-separation rates. The lower two panels show the predictions of the PCAs for high-skill job-finding and job-separation rates against their actual standardized values. Source: Authors' calculations using CPS microdata.

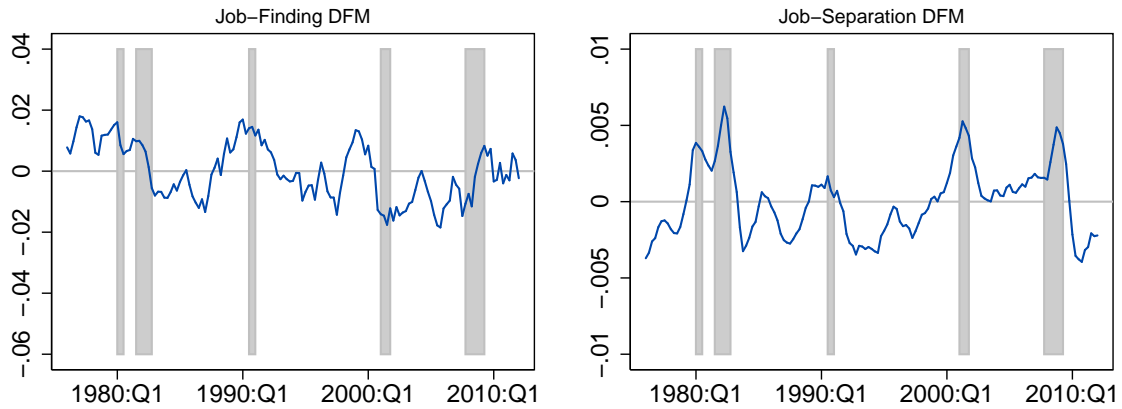
	(1)	(2)
	Job-Finding DFM	Job-Separation DFM
State Equations		
AR1 coefficient for F_t (ρ)	0.959*** (.022)	0.939*** (.027)
Observation Equations		
High-Skill:		
Constant (α_i)	0.251*** (.019)	0.565*** (.024)
F_t coef (ϕ_i)	0.010*** (.001)	0.016*** (.003)
Var(e_{it})	0.00035*** (0.00005)	5.53e-05*** (7.25e-06)
R^2	0.83	0.52
Middle-Skill: Other		
Constant (α_i)	0.245*** (.022)	1.37*** (.068)
F_t coef (ϕ_i)	0.012*** (.001)	0.053*** (.006)
Var(e_{it})	0.00009*** (0.00002)	9.84e-05*** (2.21e-05)
R^2	0.95	0.72
Middle-Skill: Manufacturing		
Constant (α_i)	0.250*** (.022)	1.76*** (.145)
F_t coef (ϕ_i)	0.012*** (.001)	0.108*** (.015)
Var(e_{it})	0.00034*** (0.00004)	0.00111*** (1.55e-04)
R^2	0.86	0.49
Middle-Skill: Construction		
Constant (α_i)	0.333*** (.026)	5.59*** (.476)
F_t coef (ϕ_i)	0.014*** (.001)	0.379*** (.036)
Var(e_{it})	0.00061*** (0.00008)	0.00218** (7.01e-04)
R^2	0.86	0.88
Low-Skill		
Constant (α_i)	0.237*** (.020)	1.92*** (.087)
F_t coef (ϕ_i)	0.011*** (.001)	0.066*** (.008)
Var(e_{it})	0.00036*** (0.00005)	2.74e-04*** (3.67e-05)
R^2	0.83	0.67
Observations	145	145

Table 1. ESTIMATES FROM DYNAMIC FACTOR MODELS (DFMs) OF JOB-FINDING AND JOB-SEPARATION RATES. Note: The symbol ** denotes significance at the 1.0 percent level; *** denotes significance at the 0.1 percent level. Parameter estimates and standard errors in the observation equations of the separation DFM have been multiplied by 100 for notational convenience. The error variances in the state equations of both DFMs have been normalized to 1.

Panel A: Middle-Other



Panel B: Middle-Manufacturing



Panel C: Middle-Construction

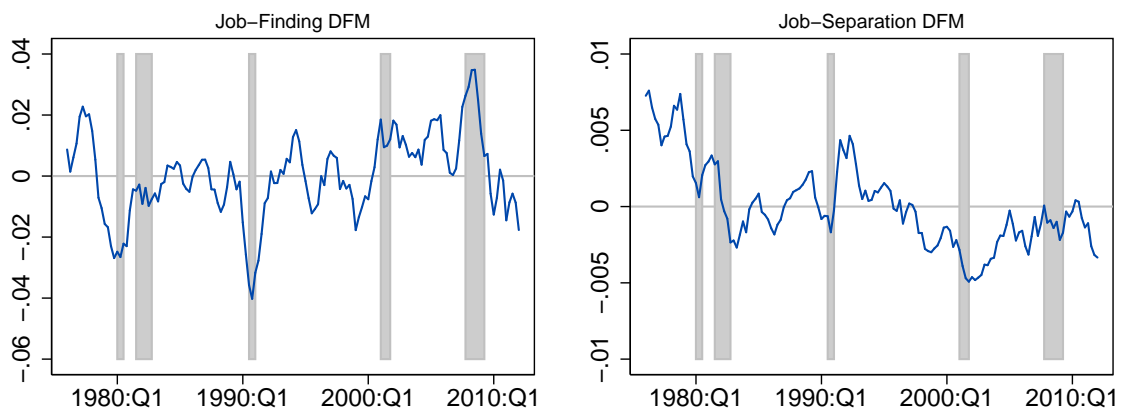
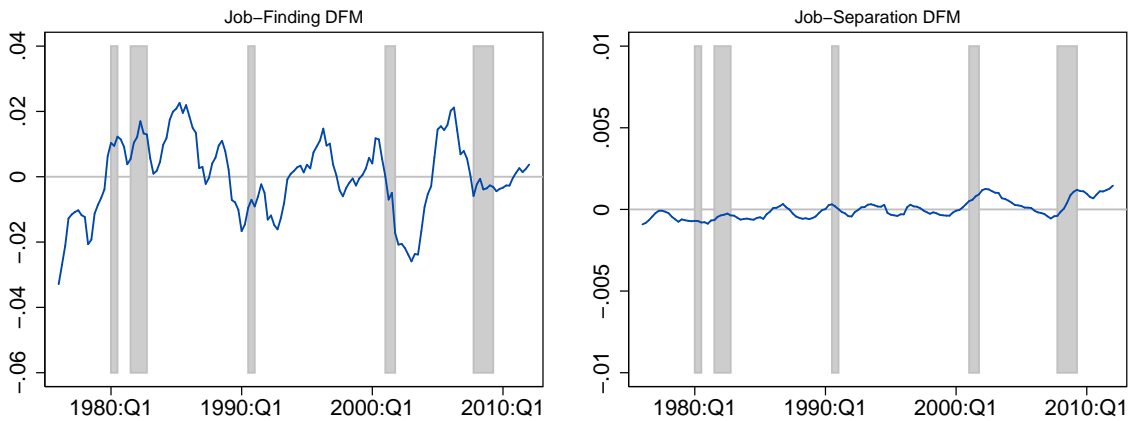


Figure 12. SMOOTHED ERRORS FOR MIDDLE-SKILL WORKERS FROM THE DYNAMIC FACTOR MODELS (DFMs): 1976:Q1–2012:Q1. Note: Errors are smoothed by averaging over months $t - 3$ to $t + 2$ with equal weights. Finding and separation rates are adjusted for demographic composition and time-aggregation before inclusion in the DFMs. See Figure 13 for errors for high-skill and low-skill workers from the same DFMs. Recessions are shaded. Source: Authors' calculations using CPS microdata.

Panel A: High-Skill



Panel B: Low-Skill

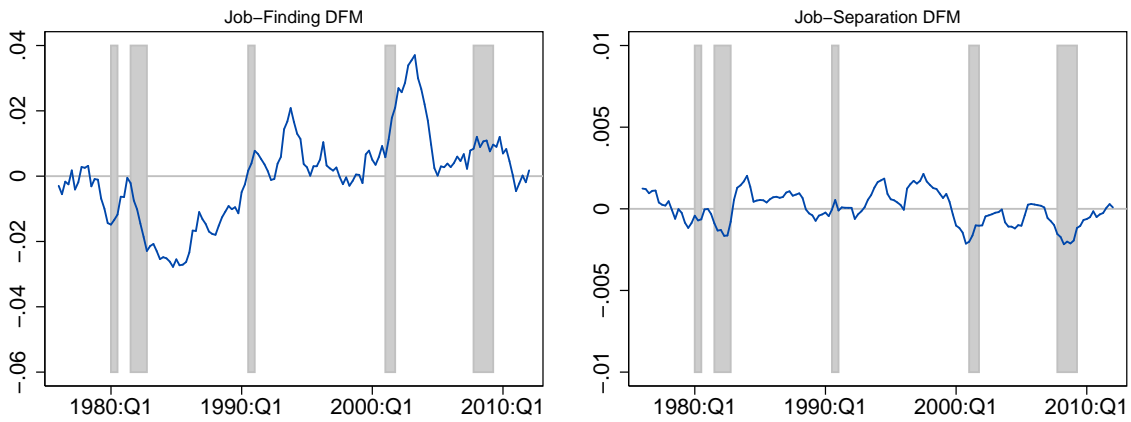


Figure 13. SMOOTHED ERRORS FOR HIGH-SKILL AND LOW-SKILL WORKERS FROM THE DYNAMIC FACTOR MODELS (DFMs): 1976:Q1–2012:Q1. Note: Errors are smoothed by averaging over quarters $t - 3$ to $t + 2$ with equal weights. Finding and separation rates are adjusted for demographic composition and time-aggregation before inclusion in the DFMs. See Figure 12 for errors for middle-skill workers from the same DFMs. Recessions are shaded. Source: Authors' calculations using CPS microdata.

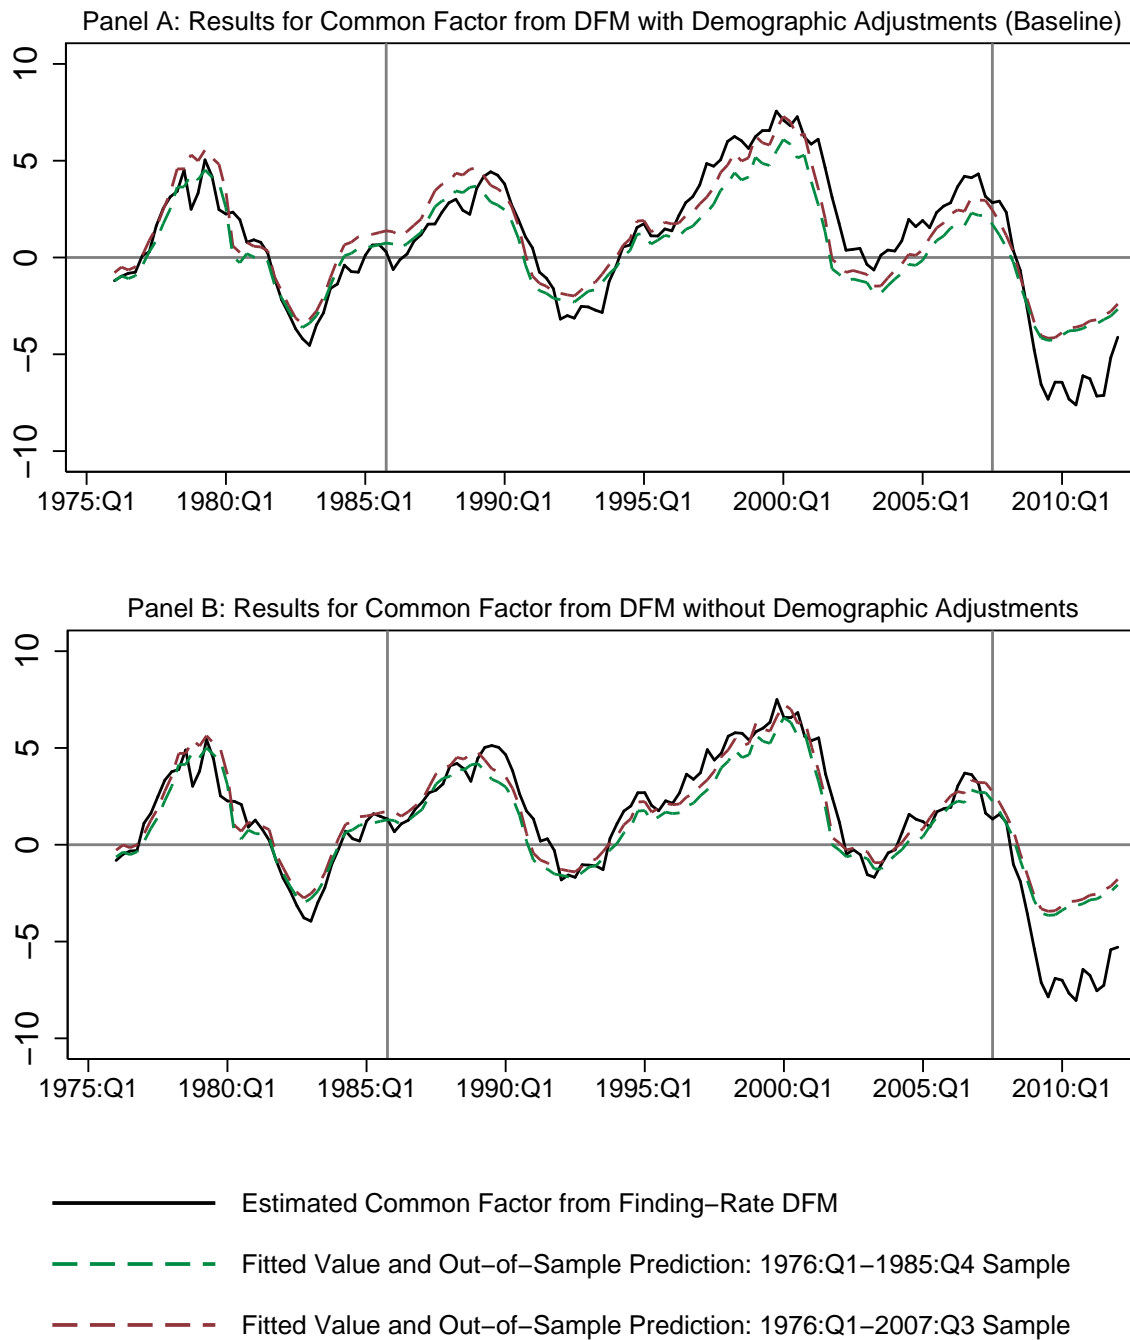


Figure 14. TIME-SERIES BEHAVIOR OF THE HIRING FUNCTION: 1976:Q1–2012:Q1. Note: The dark blue line in each panel is an estimated common factor from a finding-rate DFM. The green and red lines are predictions for these common factors, generated by linear regressions of the common factors on the vacancy–unemployment (VU) ratio and a constant. The green lines use samples that end in 1985:Q4 (first vertical line), while the red lines use samples that end in 2007:Q3 (second vertical line). The common factor in Panel A comes from the baseline finding-rate DFM, for which the five component finding rates are adjusted both for demographic composition and time-aggregation. Panel B uses the factor from a DFM for which the component rates are adjusted only for time-aggregation. The vacancy series in the numerator of the VU ratio is due to Barnichon (2010). The level of unemployment in the denominator of the VU ratio is the seasonally adjusted number of unemployed persons reported by the Bureau of Labor Statistics (BLS). Source: Authors’ calculations using CPS microdata.

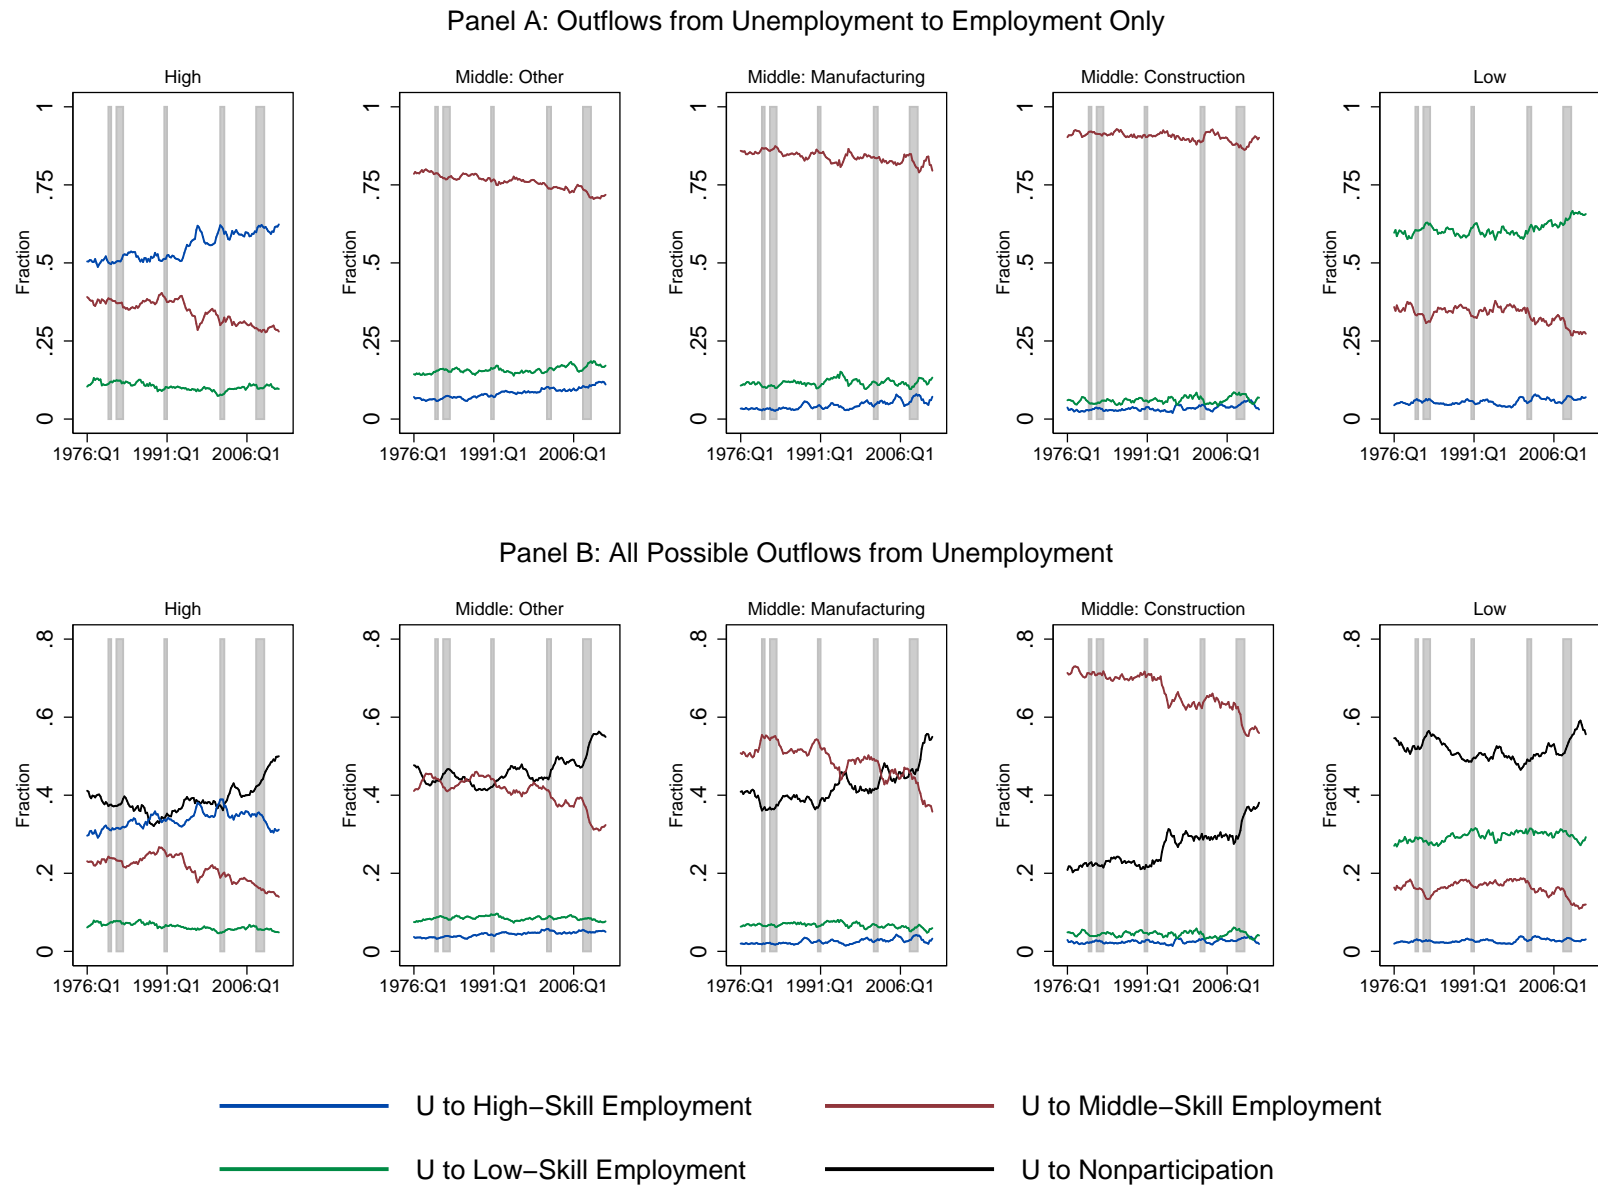
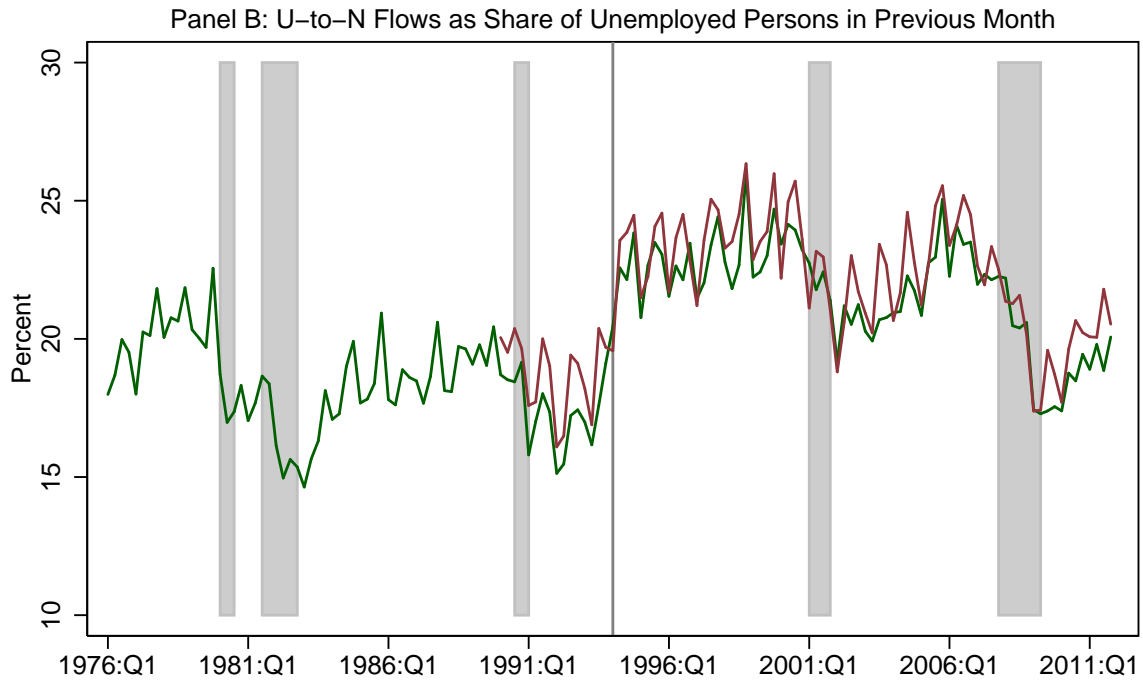
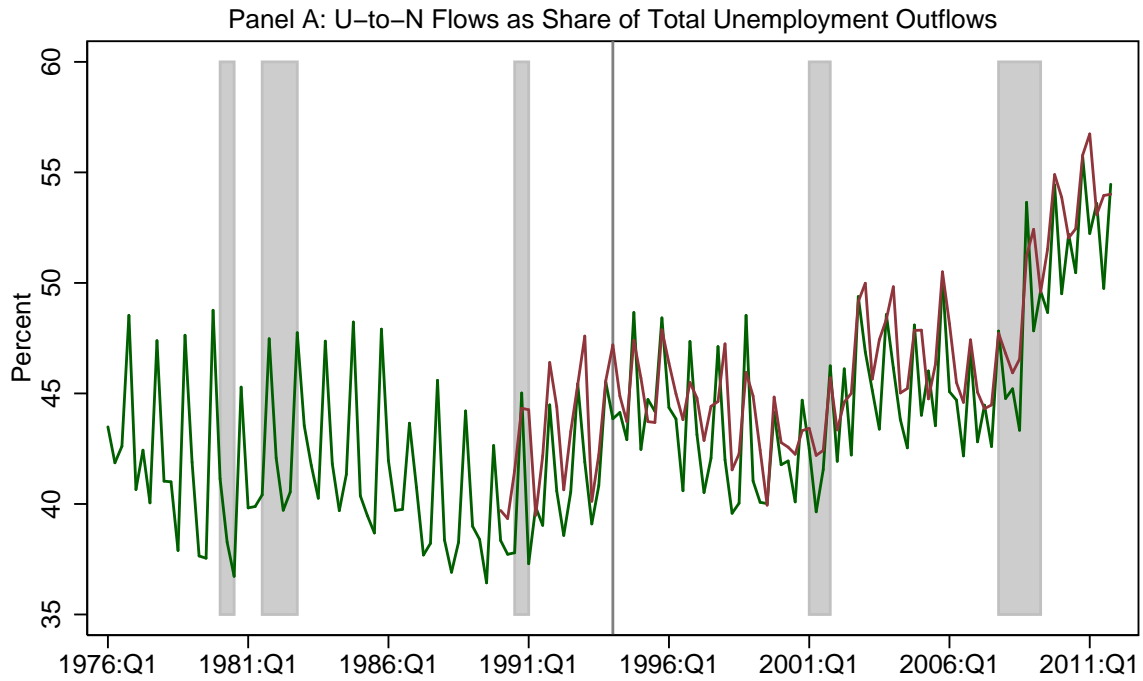


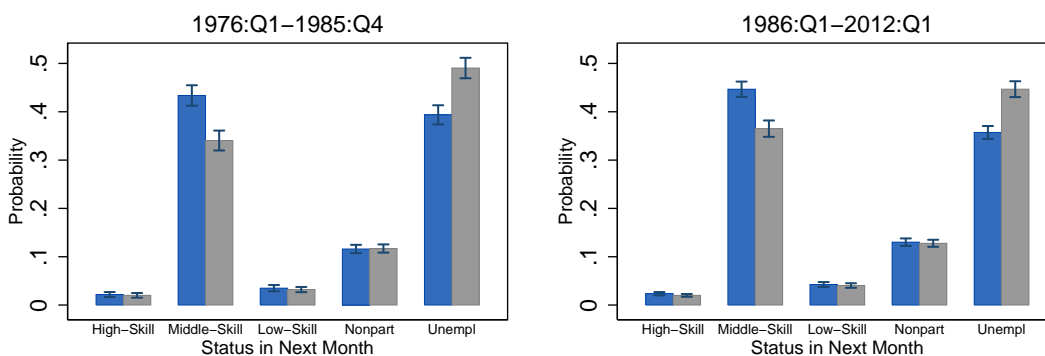
Figure 15. COMPOSITION OF OUTFLOWS FROM UNEMPLOYMENT: 1976:Q1–2012:Q1. Note: Series are quarterly averages of monthly transition probabilities for unemployed workers, conditional on exiting unemployment. Panel A includes exits to one of three types of jobs (high-skill, middle-skill, or low-skill). Panel B includes all possible exits from unemployment, adding exits to nonparticipation. All series are equally weighted moving averages over quarters $t - 3$ to $t + 2$ with equal weights. Recessions are shaded. Source: Authors' calculations using CPS microdata.



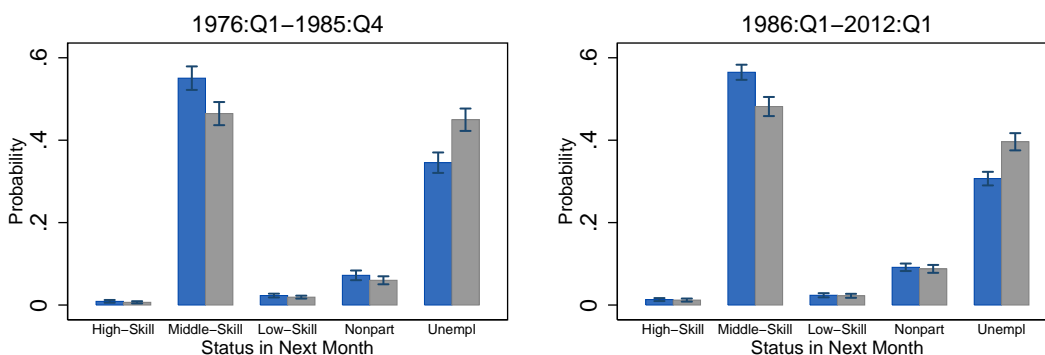
— Using CPS Microdata — Using Published Gross Flows from BLS

Figure 16. FLOWS FROM UNEMPLOYMENT TO NONPARTICIPATION: 1976:Q1–2012:Q1. Note: Series are quarterly averages of monthly not-seasonally-adjusted data. The green lines are generated by the CPS individual-level data used in this paper and the red lines come from gross worker flows and unemployment levels published by the BLS. The vertical line denotes the major redesign of the CPS in 1994. Recessions are shaded. Source: Authors' calculations using CPS microdata.

Panel A: Middle–Other Unemployed



Panel B: Middle–Manufacturing Unemployed



Panel C: Middle–Construction Unemployed

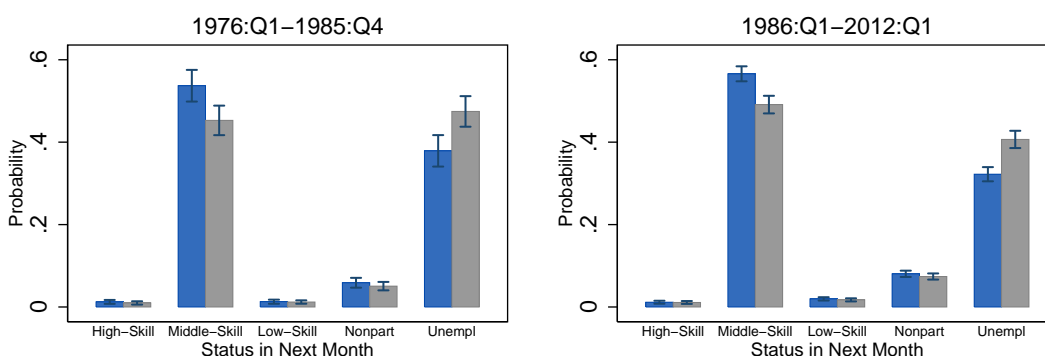


Figure 17. EFFECT OF COMMON JOB-FINDING FACTOR ON TRANSITION PROBABILITIES FOR UNEMPLOYED MIDDLE-SKILL WORKERS: 1976:Q1–1985:Q4 AND 1986:Q1–2012:Q1. Note: Probabilities are estimated from multinomial logits of job-finding rates. Separate logits are run in the two sample periods. Regressors include demographic and duration dummies defined so that the baseline worker is an unmarried 35-year-old white male with a high-school education who reports zero to one week of unemployment duration. Standard errors clustered by quarter. Probabilities for middle-skill workers without the demographic or duration dummies appear in Figure A3. Probabilities for high-skill and low-skill workers from the same logits as those above appear in Figure A4. Source: Authors’ calculations using CPS microdata.

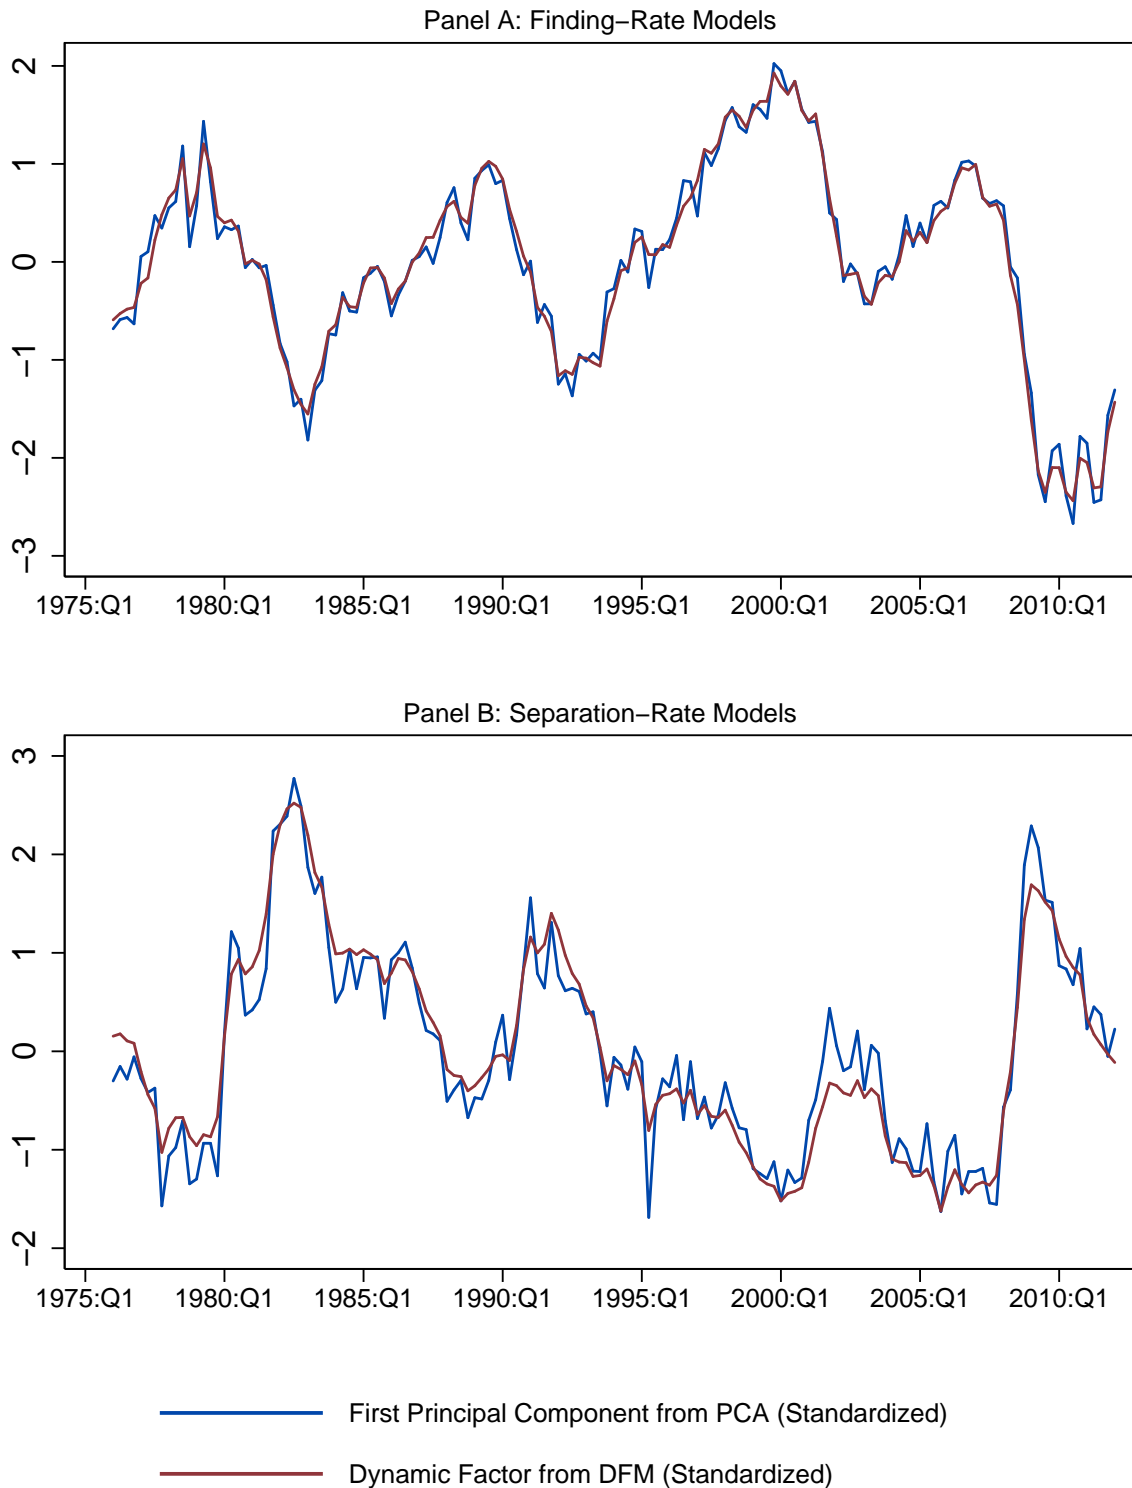


Figure A1. COMPARISON OF FIRST FACTORS FROM THE PCAs WITH COMMON FACTORS FROM THE DYNAMIC FACTOR MODELS (DFMs): 1976:Q1–2012:Q1. Note: The DFMs are estimated on job-finding and job-separation rates that are not zero-one standardized or seasonally adjusted beforehand, though the component rates are adjusted for demographic composition and for time-aggregation. To facilitate comparison of the common factors that emerge from the DFMs with the first factors from the PCAs, both the common factors from the DFMs and the first factors from the PCAs are standardized to have mean zero and variance one in the graphs above. Source: Authors’ calculations using CPS microdata.

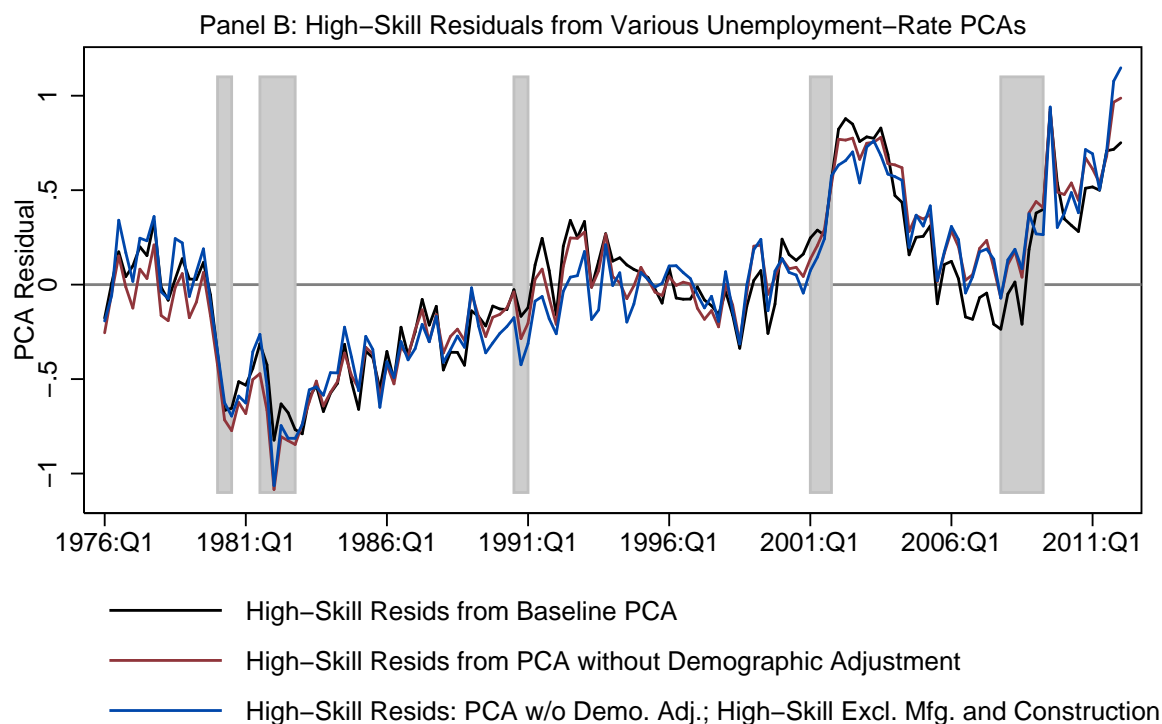
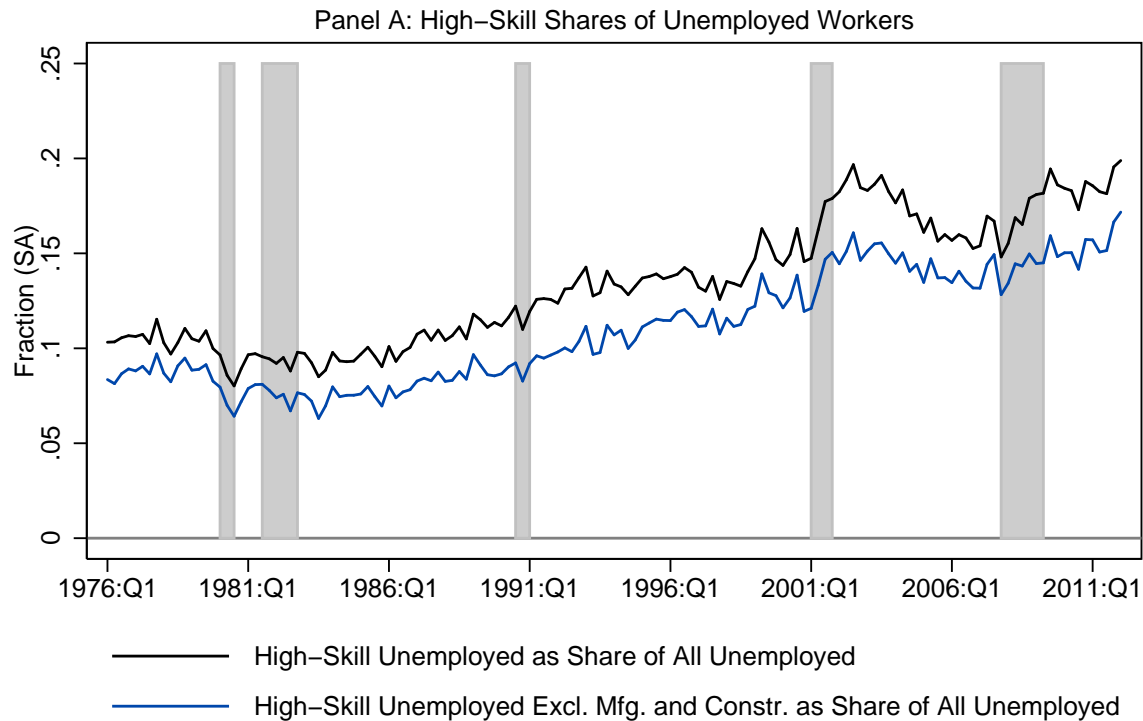
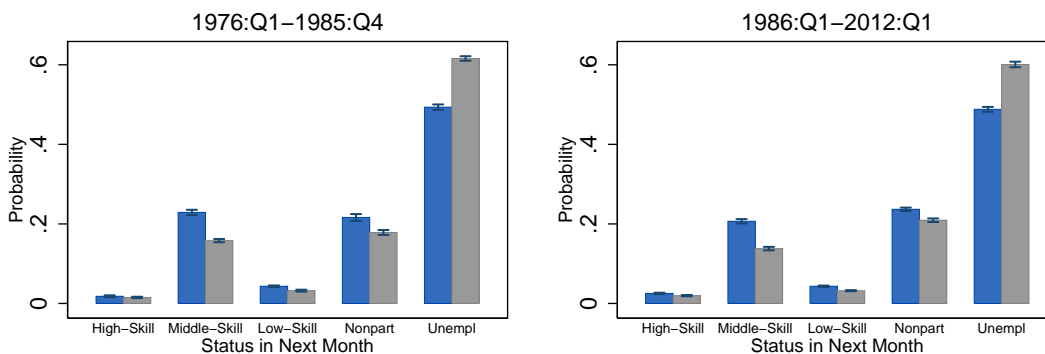
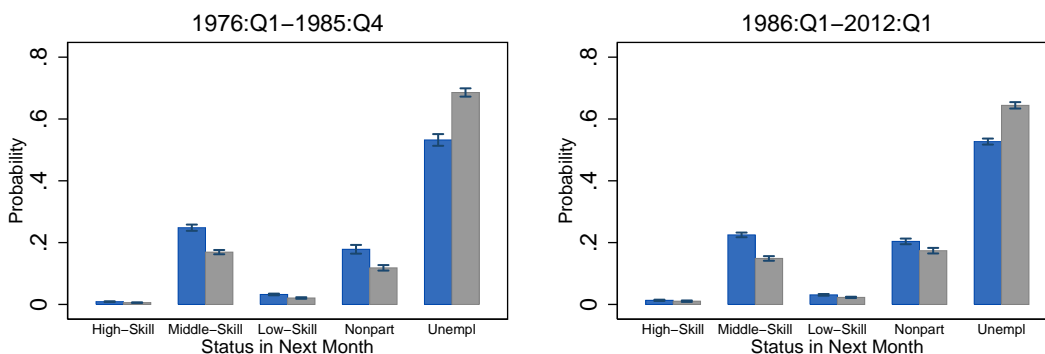


Figure A2. ROBUSTNESS CHECKS FOR HIGH-SKILL UNEMPLOYMENT SHARES AND UNEMPLOYMENT RATES. Note: Panel A shows the shares of unemployed workers that are high-skill, with and without construction and manufacturing workers. Panel B shows high-skill residuals from different PCAs of unemployment rates of the five industry-skill groups. As in all other PCAs, all component unemployment rates have been seasonally adjusted and standardized to have mean zero and variance one. Recessions are shaded. Source: Authors' calculations using CPS microdata.

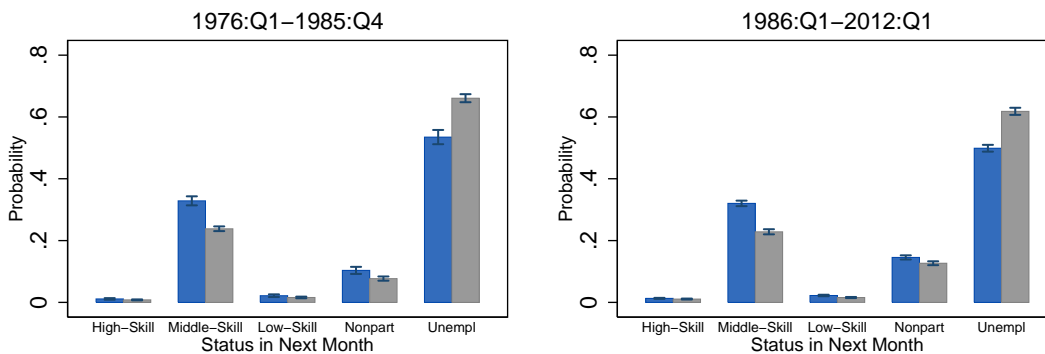
Panel A: Middle–Other Unemployed



Panel B: Middle–Manufacturing Unemployed



Panel C: Middle–Construction Unemployed



High Job-Finding Rate (One Std. Dev. Above Mean)
 Low Job-Finding Rate (One Std. Dev. Below Mean)

Figure A3. EFFECT OF COMMON JOB-FINDING FACTOR ON TRANSITION PROBABILITIES FOR UNEMPLOYED MIDDLE-SKILL WORKERS, WITH NO DEMOGRAPHIC OR DURATION CONTROLS: 1976:Q1–1985:Q4 AND 1986:Q1–2012:Q1. Note: With no demographic or duration controls, a drop in the finding rate appears to reduce the probability of moving from unemployment to nonparticipation. See Figure 17 for probabilities from logits that include both demographic and duration controls. Source: Authors’ calculations using CPS microdata.

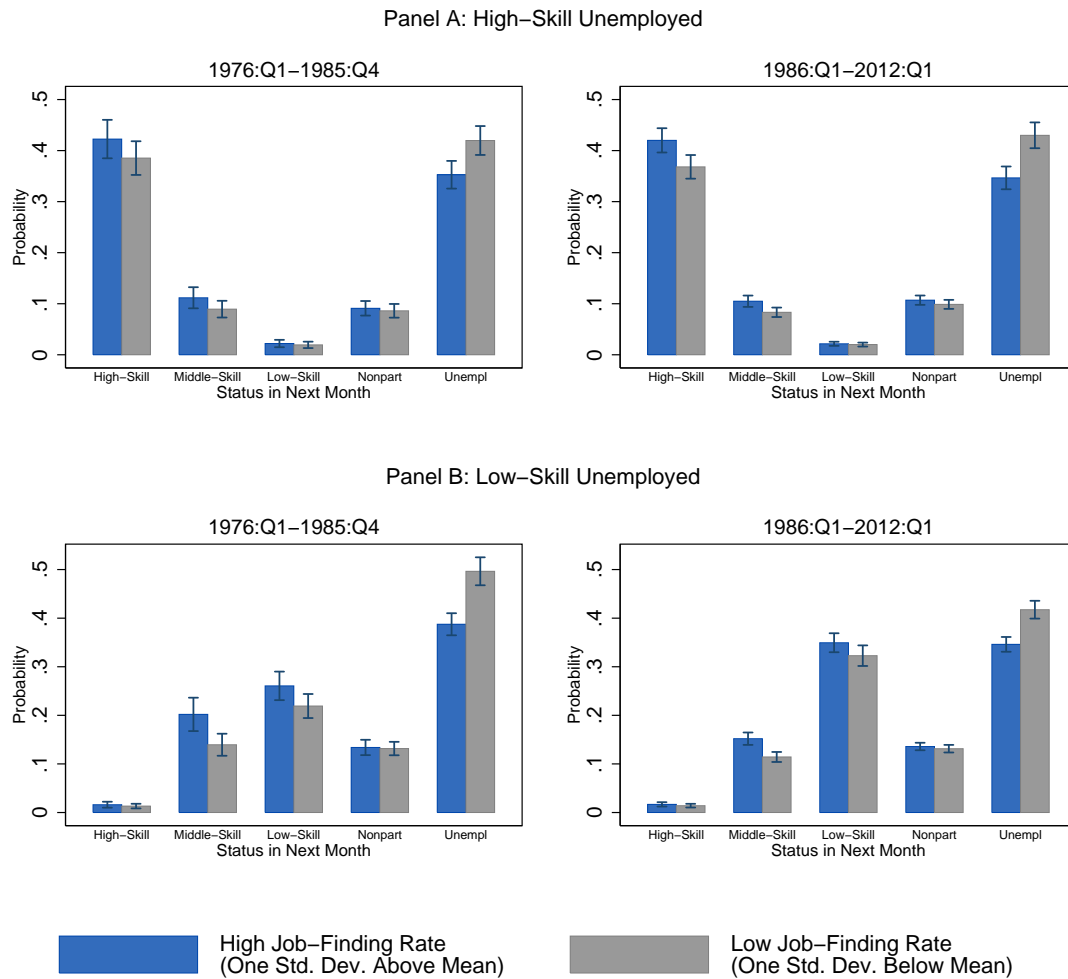
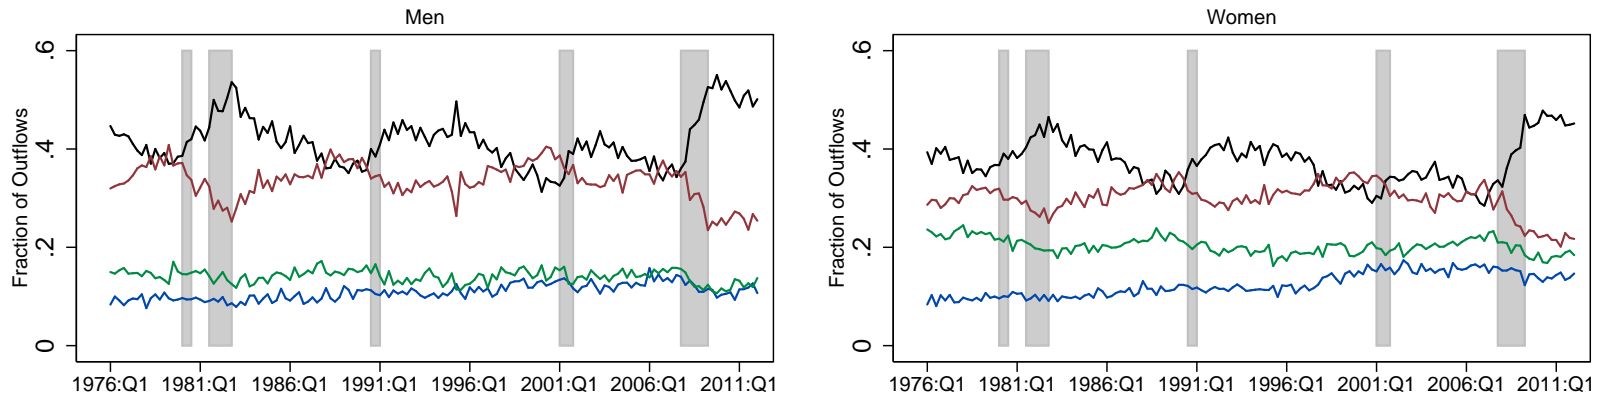
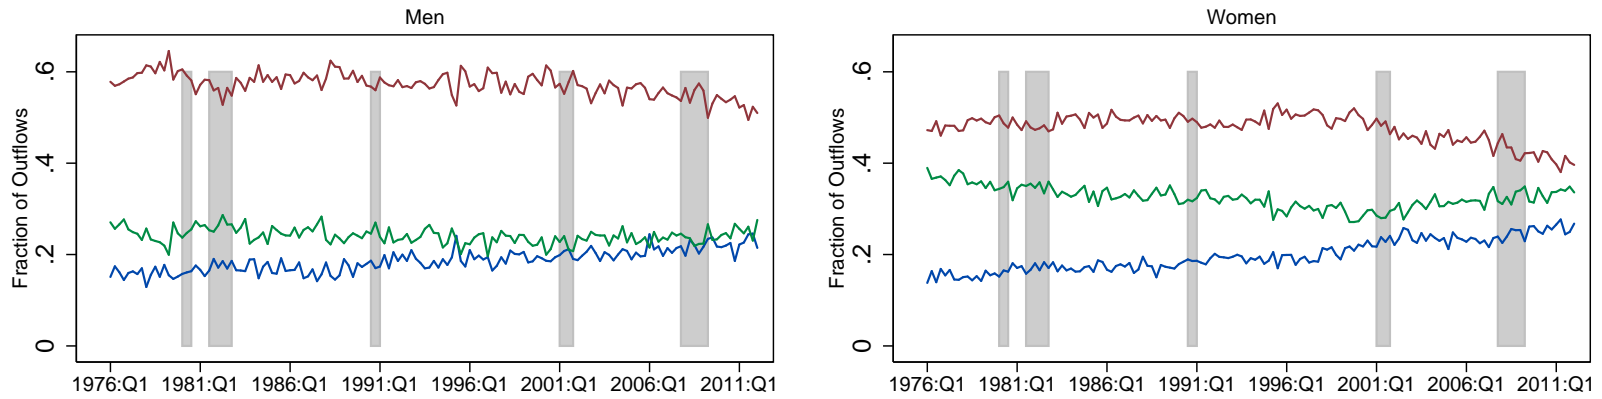


Figure A4. EFFECT OF COMMON JOB-FINDING FACTOR ON TRANSITION PROBABILITIES FOR UNEMPLOYED HIGH- AND LOW-SKILL WORKERS: 1976:Q1–1985:Q4 AND 1986:Q1–2012:Q1. Note: See the notes to Figure 17 for the specification of the multinomial logits that generate these probabilities; that figure includes the relevant probabilities for middle-skill workers generated by the same models. Source: Authors' calculations using CPS microdata.

Panel A: All Possible Outflows from Nonparticipation



Panel B: Outflows from Nonparticipation to Employment Only



N to Unemployment
 N to Middle-Skill Employment
 N to High-Skill Employment
 N to Low-Skill Employment

Figure A5. COMPOSITION OF OUTFLOWS FROM NONPARTICIPATION: 1976:Q1–2012:Q1. Note: Series are seasonally adjusted quarterly averages of monthly transition probabilities for nonparticipating workers, conditional on exiting nonparticipation. Panel A includes all possible exits from nonparticipation (unemployment, a high-skill job, a middle-skill job, or a low-skill job). Panel B includes only exits to one of the three types of jobs. Recessions are shaded. Source: Authors’ calculations using CPS microdata.

References

- Abowd, John, and Arnold Zellner. 1985. "Estimating Gross Labor-Force Flows." *Journal of Business and Economics Statistics* 3: 254–293.
- Abraham, Katharine G., and Lawrence F. Katz. 1986. "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *Journal of Political Economy* 94(3): 507–522.
- Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, eds. Orley Ashenfelter and David E. Card, vol. 4. Elsevier.
- Aghion, Phillipe, and Gilles Saint-Paul. 1998. "Virtues of Bad Times: Interaction Between Productivity Growth and Economic Fluctuations." *Macroeconomic Dynamics* 2(3): 322–344.
- Autor, David H. 2010. "The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings." A paper jointly released by The Center for American Progress and The Hamilton Project. Available at <http://economics.mit.edu/files/5554>.
- Autor, David H., and David Dorn. 2009. "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market." NBER Working Paper No. 15150.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90(2): 300–323.
- Barchenon, Regis, Michael W.L. Elsby, Bart Hobijn, and Aysegul Sahin. 2011. "Which Industries are Shifting the Beveridge Curve?" Federal Reserve Bank of San Francisco Working Paper No. 2010-32. Available at <http://www.frbsf.org/publications/economics/papers/2010/wp10-32bk.pdf>.
- Barchenon, Regis, and Andrew Figura. 2012. "The Determinants of the Cycles and Trends in U.S. Unemployment." University of Pompeu Fabra Working Paper. Available at <https://sites.google.com/site/regisbarnichon/research>.
- Barlevy, Gadi. 2002. "The Sullyng Effect of Recessions." *Review of Economic Studies* 69(1): 65–96.
- Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109(3): 175–178.

- Bowler, Mary, Randy E. Ilg, Stephen Miller, Ed Robinson, and Anne Polivka. 2003. "Revisions to the Current Population Survey Effective in January 2003." Bureau of Labor Statistics. Available at <http://www.bls.gov/cps/rvcps03.pdf>.
- Caballero, Ricardo J., and Mohamad L. Hammour. 2005. "The Cost of Recessions Revisited: A Reverse-Liquidationist View." *Review of Economic Studies* 72(2): 313–341.
- Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2012. "Manufacturing Busts, Housing Booms, and Declining Employment: A Structural Explanation." University of Chicago Booth School of Business Working Paper. Available at http://faculty.chicagobooth.edu/erik.hurst/research/charles_hurst_noto_manufacturing.pdf.
- Cortes, Guido Matias. 2012. "Where Have the Routine Workers Gone? A Study of Polarization Using Panel Data." University of Manchester Economics Department Discussion Paper No. 1224. Available at <http://www.socialsciences.manchester.ac.uk/disciplines/economics/research/discussionpapers/pdf/EDP-1224.pdf>.
- Davis, Steven J. 1987. "Allocative Disturbances and Specific Capital in Real Business Cycle Theories." *American Economic Review* 77(2): 326–332.
- Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2010. "The Establishment-Level Behavior of Vacancies and Hiring." NBER Working Paper 16265. Available at <http://www.nber.org/papers/w16265>.
- Davis, Steven J., and John Haltiwanger. 1990. "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications." *National Bureau of Economic Research Macroeconomics Annual* 5: 123–168.
- Davis, Steven J., and John Haltiwanger. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Quarterly Journal of Economics* 107(3): 819–863.
- Dickens, William T., and Robert K. Triest. 2012. "Potential Effects of the Great Recession on the U.S. Labor Market." *B.E. Journal of Macroeconomics* 12(3).
- Elsby, Michael W. L., Bart Hobijn, and Aysegul Sahin. 2010. "The Labor Market in the Great Recession." *Brookings Papers on Economic Activity* 2010(1): 1–69.
- Elsby, Michael W.L., Bart Hobijn, and Aysegul Sahin. 2012. "On the Importance of the Participation Margin for Labor Market Fluctuations." Unpublished manuscript.
- Faberman, R. Jason, and Bhashkar Mazumder. 2012. "Is There a Skills Mismatch in the Labor Market." Chicago Fed Letter No. 300, Federal Reserve Bank of

- Chicago. Available at http://www.chicagofed.org/digital_assets/publications/chicago_fed_letter/2012/cfljuly2012_300.pdf.
- Federal Open Market Committee of the Federal Reserve System (FOMC). 2012. “Minutes for Meeting of September 12–13.” Available at <http://www.federalreserve.gov/monetarypolicy/files/fomcminutes20120913.pdf>.
- Fleischman, Charles A., and John M. Roberts. 2011. “From Many Series, One Cycle: Improved Estimates of the Business Cycle from a Multivariate Unobserved Components Model.” Federal Reserve Board Finance and Economics Discussion Series Paper No. 2011-46. Available at <http://www.federalreserve.gov/pubs/feds/2011/201146/201146pap.pdf>.
- Goldin, Claudia D., and Lawrence F. Katz. 2008. *The Race Between Education and Technology*. Cambridge, Mass.: Belknap Press of Harvard University Press.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2009. “Job Polarization in Europe.” *American Economic Review* 99(2): 58–63.
- Hall, Robert E. 1991. “Labor Demand, Labor Supply, and Employment Volatility.” *National Bureau of Economic Research Macroeconomics Annual* 6: 17–62.
- Hoynes, Hilary, Douglas L. Miller, and Jessamyn Schaller. 2012. “Who Suffers during Recessions?” *Journal of Economic Perspectives* 26(3): 27–48.
- Jaimovich, Nir, and Henry E. Siu. 2012. “The Trend is the Cycle: Job Polarization and Jobless Recoveries.” National Bureau of Economic Research Working Paper No. 18334. Available at <http://www.nber.org/papers/w18334>.
- Krause, Michael U., and Thomas A. Lubik. 2006. “The Cyclical Upgrading of Labor and On-the-Job Search.” *Labour Economics* 13(4): 459–477.
- Kudlyak, Marianna, and Felipe Schwartzman. 2012. “Accounting for Unemployment in the Great Recession: Nonparticipation Matters.” Federal Reserve Bank of Richmond Working Paper. Available at http://richmondfed.org/publications/research/working_papers/2012/pdf/wp12-04.pdf.
- Lazear, Edward P., and James R. Spletzer. 2012. “The United States Labor Market: Status Quo or A New Normal?” Paper Presented at the Jackson Hole Symposium *The Changing Policy Landscape* sponsored by the Federal Reserve Bank of Kansas City. Available at <http://www.kansascityfed.org/publicat/sympos/2012/el-js.pdf>.

- Lee, Yoonsoo, and Toshihiko Mukoyama. 2012. "Entry, Exit, and Plant-level Dynamics over the Business Cycle." University of Virginia Working Paper. Available at <http://people.virginia.edu/~tm5hs/LM.pdf>.
- Lilien, David M. 1982. "Sectoral Shifts and Cyclical Unemployment." *Journal of Political Economy* 90(4): 777–793.
- Meyer, Peter B., and Anastasiya M. Osborne. 2005. "Proposed Category System for 1960–2000 Census Occupations." U.S. Bureau of Labor Statistics, Office of Productivity and Technology. Working Paper no. 383. Available at <http://econpapers.repec.org/RePEc:bls:wpaper:ec050090>.
- Mincer, Jacob. 1991. "Education and Unemployment." National Bureau of Economic Research Working Paper No. 3838.
- Mukoyama, Toshihiko, and Aysegul Sahin. 2006. "Costs of Business Cycles for Unskilled Workers." *Journal of Monetary Economics* 53(8): 2179–2193.
- Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*. Second edition. Cambridge, Mass.: MIT Press.
- Reicher, Christopher. 2012. "The Aggregate Effects of Long-Run Sectoral Reallocation." Kiel Institute for the World Economy Working Paper. Available at <https://sites.google.com/site/chrisreicher/research>.
- Rissman, Ellen R. 2009. "Employment Growth: Cyclical Movements or Structural Change." *Economic Perspectives*, Federal Reserve Bank of Chicago 2009(4): 40–57.
- Sahin, Aysegul, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2012. "Mismatch Unemployment." Federal Reserve Bank of New York Working Paper. Available at <http://www.newyorkfed.org/research/economists/sahin/USmismatch.pdf>.
- Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95(1): 25–49.
- Shimer, Robert. 2011. "Job Search, Labor Force Participation, and Wage Rigidities." University of Chicago Working Paper. Available at <https://sites.google.com/site/robertshimer/research/workingpapers>.
- Shlens, Jonathon. 2009. "A Tutorial on Principal Component Analysis." Version 3.01. Unpublished manuscript. Available at <http://www.sn1.salk.edu/~shlens/pca.pdf>.

Stock, James H., and Mark W. Watson. 2011. "Dynamic Factor Models." In *Oxford Handbook of Economic Forecasting*, eds. Michael P. Clements and David F. Hendry. Oxford, U.K.: Oxford University Press.

Stock, James H., and Mark W. Watson. 2012. "Disentangling the Channels of the 2007-2009 Recession." Princeton University Working Paper. Available at http://www.princeton.edu/~mwatson/papers/Stock_Watson_Disentangling_May_2012.pdf.

Tasci, Murat. 2012. "The Ins and Outs of Unemployment in the Long Run: Unemployment Flows and the Natural Rate." Federal Reserve Bank of Cleveland Working Paper. Available at <http://www.clevelandfed.org/research/workpaper/2012/wp1224.pdf>.