NBER Macroeconomics Annual (2014) Proposal
"Labor Market Polarization Over the Business Cycle"

Christopher L. Foote and Richard W. Ryan

July 2013

1 Introduction

For the 2014 *NBER Macroeconomics Annual*, we propose a cyclical analysis of occupational labor market data and reallocation patterns using some new data constructed for this purpose. Occupational designations have become increasingly important in labor market analysis amid the so-called polarization of the labor market during the last three decades. Briefly, the polarization hypothesis of Autor (2010), Acemoglu and Autor (2011) and others holds that since the late 1970s, job opportunities have been most favorable at the poles of the skill distribution. The jobs of high-skill workers (such as college professors and technicians) and those of low-skilled workers (such as custodians and cooks) have proven relatively immune to displacement by either automation or by international trade.[1] By contrast, due to their routine nature, many jobs in the middle part of the skill distribution, including assembly-line and clerical work, are more easily replaced by robots or lower-paid workers overseas. So far, most of the academic work on polarization has focused on its long-run impact on wages and employment patterns, with less attention to potential cyclical aspects. Yet a long literature, originating with Lilien (1982), Davis and Haltiwanger (1992), and the *NBER Macro Annual* papers of Davis and Haltiwanger (1990) and Hall (1991), has encouraged economists to consider potential relationships between aggregate business cycles and employment reallocation. A hypothesis in this literature is that much of the higher unemployment experienced in recessions is really the intensified reallocation of workers from one type of job to another, with this reallocation encouraged either by allocational shocks (as opposed to aggregate shocks) or by the lower opportunity cost of suspending production in favor of reallocation when aggregate productivity is low.

Some recent research has analyzed job polarization in the context of this "recessions-as-reallocations" approach. An influential recent paper by Jaimovich and Siu (2012, henceforth JS) claims not only that workers choose to reallocate when aggregate productivity is low, but also that the optimal reallocation of middle-skill workers has intensified during the recessions of the polarization era. Further, JS claim that polarization-era reallocation explains why

[1] Despite advances in artificial intelligence, machines have yet to replace human creativity and skill in occupations such as medicine, higher education, and business management. But modern machines are also unable to do tasks that most humans find easy, such as emptying a dishwasher or carrying a vacuum cleaner up a flight of stairs. And low-skill jobs like building and grounds cleaning are impossible to outsource to other countries.

the last three recoveries have been "jobless." The idea is that middle-skill workers deserted their jobs *en masse* when recent recessions began, because negative productivity shocks had temporarily reduced the opportunity cost of reallocating. And because these middle-skill workers took a long time to establish themselves in new positions, employment growth during the ensuing recoveries was less than robust.

The claim that polarization explains jobless recoveries is controversial. The JS model is firmly nested in the search-and-matching literature, where the main economic friction is the inability of workers and firms to instantly create new job matches. As a result, when aggregate productivity declines and the implicit cost of reallocation falls, employment and unemployment respond. But search frictions may not be the only frictions relevant to the macroeconomy—sticky prices and wages may matter as well. If so, then recent jobless recoveries may instead stem from changes in the time-series behavior of the aggregate shocks to technology, markups, or preferences that drive aggregate output in structural models that include nominal rigidities. Along these lines, Galí, Smets, and Wouters (2012) argue that the so-called jobless recoveries are better characterized as slow recoveries, in the sense that GDP did not bounce back as quickly after recent recessions as it did after earlier ones. To support this view, the authors contend that there has been no fundamental change in the empirical relationship between GDP and employment in the jobless–recovery era. The implication is that recent recoveries should be analyzed by looking for changes in the pattern of the structural shocks that determine GDP, not by looking for changes in the cyclical reallocation behavior of any particular group of workers.

A key goal of our paper is to analyze the potential relationship between job polarization and recent business cycles with some new data and new empirical approaches. Like JS (2012), our paper draws on historical employment data to construct occupational employment series that are appropriate for a polarization study. In particular, we follow JS in disaggregating total CPS employment into:

- One high-skill classification: *nonroutine–cognitive*, which includes managers and professional and technical workers,
- Two middle-skill classifications: *routine–manual*, which includes laborers and machine operatives, among others; and *routine–cognitive*, which are primarily sales workers, clerical workers and other office and administrative workers, and
- One low-skill classification: *nonroutine–manual*, consisting of in-person service workers such as custodians, cooks, and waiters.

Though our classifications are similar to those of JS, we believe that we improve on their empirical work in four respects. First, our occupational aggregates go back to the third

quarter of 1947, so we are able to measure occupational employment movements over virtually the entire postwar era.[2] Second, we take pains to ensure that each of the four employment series we construct are consistent across the numerous definitional changes that occur in occupational data during the past 65 years. Third, we explicitly look for changes in the relationship of occupational employment data to GDP, a question at the heart of the jobless-recovery vs. slow-recovery debate. Finally, in a separate analysis using CPS microdata since 1976, we look for direct evidence of intensified reallocation of unemployed workers in recessions, an implication of recessions-as-reallocations approach.

Our work is still in progress. But so far, the analysis reveals three main findings:

- As we might expect, employment in the middle-skill, routine–manual category, which includes many factory workers and machine operatives, is traditionally more cyclical than employment in the other three categories. But the relationship between routine–manual employment and GDP has been remarkably stable since 1947. While it is true that many routine–manual jobs were shed in recent recessions, the size of the job losses experienced by this group are consistent with the size of coincident changes in GDP. One way to see this pattern is with a simple rolling correlation of detrended GDP and detrended routine–manual employment. This correlation remains close to one at all points in the sample. A more formal way to show the stability of the GDP/routine–manual relationship is with a dynamic factor model (DFM). Our baseline DFM projects the four occupation-level employment series onto a latent "business-cycle" factor that is itself a function of GDP.[3] Estimated on two subsamples split at 1985q4, this methodology shows that the response of routine-manual employment to the business cycle in the jobless-recovery era is almost exactly the same as the response we would expect using pre-1985 data. The implication is that jobless recoveries do not involve changes in the relationship between routine–manual employment and GDP in an important way.

- While the relationship between routine–manual employment and GDP has been stable since the late 1940s, the same is not true of employment in other classes. Employment of nonroutine–cognitive workers (including high-skill managers and professionals) is nearly acyclical in the early part of the sample. But the correlation of this employment series with GDP grows over time, approaching one in the late 1990s and early 2000s, before falling back somewhat recently.[4] A similar pattern is found for employment among routine–cognitive workers (the middle-skill group that includes sales and clerical workers), though the correlation between employment for this group and GDP begins

[2] JS's data begin in July 1967.

[3] The four occupation-level employment equations also include manufacturing output and residential construction, as we explain below.

[4] See the top left panel of Figure 8.

to rise somewhat earlier. We are not the first to notice that the cyclicality of skilled labor input is higher now than before. In particular, Castro and Coen-Pirani (2008) use the Merged Outgoing Rotation Groups (MORG) of the CPS, available since 1979, to show that the hours of college-educated workers became more cyclical starting in 1984, relative to the 1979-1983 period.[5] Our analysis shows that the cyclicality of both types of cognitive employment has been rising for much of the last several decades.[6]

- Finally, the CPS microdata show little evidence of intensified occupational reallocation among unemployed workers in recessions. The microdata allow us more flexibility in classifying workers, relative to the classifications required in the study of employment levels based on historical data. As a result, we follow the Acemoglu and Autor (2011) handbook chapter by grouping workers in the CPS microdata into broad high-skill, middle-skill, and low-skill classes. We further split the middle–skill group along industry lines into middle-manufacturing, middle-construction, and middle-other groups, leaving five industry–skill classes in all. We find that unemployed middle–skill workers rarely move to high- or low-skill jobs at any point in the business cycle. There are secular trends in the reallocation of the unemployed that are consistent with an overall decline in middle-skill jobs, but for the most part these movements appear unrelated to the business cycle.

In general, these findings support the slow-recovery interpretation of recent macroeconomic data rather than the jobless-recovery interpretation. Our results also shed light on some other current topics in empirical macroeconomics. One of them concerns changes in the cyclicality of measured productivity growth during the postwar era. Papers such as Galí and Gambetti (2009) and Gordon (2010) have noted that measured productivity growth in early postwar era was highly procyclical, falling sharply in recessions (see **Figure 1**). In recent years, however, measured productivity growth has become less cyclical, leading these authors to suggest that labor-hoarding has become a less prominent feature of the business cycle. Our findings indicate that it is the labor-hoarding of cognitive workers in particular that has declined the most during the past three decades, an interesting counterpoint to any claim that only the cyclical experiences of routine workers have deteriorated in recent years. Our results are also relevant for the study of the long-term unemployment that has been a

[5]Castro and Coen-Pirani (2008) disaggregate hours data by education, not occupation, so their college-educated group is not identical to our high-skill/nonroutine–cognitive group. Obviously, however, there is substantial overlap.

[6]Most of the analysis in Castro and Coen-Pirani (2008) uses post-1979 MORG files, but they also use data from the Annual Demographic Survey (March CPS) to study the cyclicality of college-educated hours before 1979. In contrast to our results, they find that high-skilled cyclicality was relatively high between 1963 and 1975. The difference between our results and theirs may result from their use of annual rather than quarterly data before 1979, their disaggregation of workers by education rather by occupation, their focus on hours rather than employment, or their exclusion of self-employed workers.

feature of recent labor market data. Soon after the Great Recession ended, the Beveridge Curve underwent an adverse shift, suggesting that the job market was creating fewer job matches for given levels of unemployed workers and job vacancies. Occupation-specific analysis of job-finding and job-separation rates shows that whatever caused this shift, it probably has little to do with accelerated job polarization. As we have found in other work, unemployed workers in all industry-skill classes have struggled to find jobs during and immediately after the Great Recession. And consistent with the earlier evidence on the rising cyclicality of high-skill employment, the job-separation rates of high-skill workers were particularly elevated during and after the recessions of 2001 and 2007-09.[7]

2 Occupational Employment and the Business Cycle Since 1947

2.1 Data issues

We understand that the *Macro Annual* conference is not the place for a long-winded discussion of the government's occupational classification system. But understanding the changes in this system over time is important when analyzing historical labor market data for different types of workers. Like Jaimovich and Siu (2012), we coded the published historical data on employment levels for major occupational groups, as defined by the occupational classification system that was used by the government at the time.[8] High-skill, nonroutine–cognitive workers are typically found in classifications such as "managers," "proprietors," and "professional and technical" workers. Middle-skill, routine–cognitive workers are generally "sales," "clerical," or "administrative" workers. Low-skill nonroutine–manual workers are in found "service" occupations. Lastly, middle-skill routine–manual workers are found in categories such as "craftsmen," "operatives," "laborers," and "production workers." Table 1 illustrates how major occupational groups used in the CPS can be aggregated into task- or skill-based groups.

The complicating factor is that the occupational classification system used by the government changes every 10 years or so, and sometimes these changes are significant enough to affect the major occupation groups on which our employment series are based. The most significant changes in the major occupation groups occur in January 1971, January 1983,

[7]The implication that most recent unemployment is cyclical and not structural is consistent with other work in this area, such as Elsby, Hobijn, and Sahin (2010) and Dickens and Triest (2012).

[8]Starting in July 1959, occupation-level data are available in the BLS's monthly *Employment and Earnings* publication. Before that, the data are found in the Bureau of the Census's *Monthly Report on the Labor Force* (Current Population Reports Series P–57). Until January 1958, the Census Bureau reported occupation data only for the first month of each quarter (January, April, July and October); elsewhere we use quarterly averages of monthly data. We also take account of the 1958 reporting change whenever we seasonally adjust the data. Monthly, machine-readable occupational data are available from the BLS website starting in January 1983 for employment and January 2000 for unemployment rates.

Prototypical Major Occupation Groups (subject to change about every 10 years)	**Task-based Classifications**	**Skill-based Classifications**
Managers & Proprietors Professionals Technicians	Nonroutine Cognitive	High-Skill
Sales Clerical Office and Administrative	Routine Cognitive	Middle-Skill
Operatives Laborers Other "blue-collar" workers	Routine Manual	
Service workers (custodians, cooks, etc.)	Nonroutine Manual	Low-Skill

Table 1. Potential Occupational Classifications in CPS Data. The empirical work in this section uses four employment series defined along "task-based" lines, as in the middle column.

and January 2003. The 2003 change does not cause too many problems, because the BLS has constructed consistent occupational employment aggregates that extend back to January 1983 (the second major change). But when we compile our series from the historical sources we find appreciable level shifts in 1971q1 and 1983q1.

To deal a classification break, we first estimate a VAR for the log employment growth rates of the four employment series: nonroutine–cognitive, routine–cognitive, routine–manual, and nonroutine–manual. The sample period for the VAR extends up to the given break, so the VAR implies out-of-sample predictions of log employment growth for four types of workers across the break, and thus predictions for the four levels of employment as well. The published employment data after the break are then adjusted by a constant amount that joins them to the VAR predictions. The effects of the VAR adjustments are shown in **Figure 2.** In 1971q1, the most important adjustment occurs in the low-skill, nonroutine–manual category depicted in the lower right panel. In 1983q1, however, nontrivial adjustments are made to all but the high-skill, nonroutine–cognitive workers shown in the upper left panel.

Two questions are relevant for this methodology. First, after all VAR adjustments are made, does the implied *total* of the four employment series match the historical series for nonagricultural employment available on the BLS website? This question is answered by **Figure 3**, which shows that our implied totals with and without VAR adjustments closely resemble total postwar nonagricultural employment as reported by the BLS. In particular, the totals match the BLS's published aggregate for *nonagricultural industry* employment from 1948q1 through 1982q4 and they match the published aggregate for *nonagricultural occupational* employment starting in 1983q1.[9]

[9]Like the historical data we collected, both of the BLS nonagricultural aggregates in Figure 3 are generated

A second question to ask about the VAR methodology is whether the four individual series it generates can be corroborated by another source of information. To check this, we bring the CPS microdata into the analysis. Specifically, we recoded the occupation variable in the microdata using the OCC1950 classification that has been constructed as part of the University of Minnesota's IPUMS project.[10] The CPS microdata begin in January 1976, so we can determine whether the four occupational employment series constructed using the OCC1950 code bear much resemblance to the corresponding series constructed with the published data and VAR adjustments. **Figure 4** shows that for the most part, the correspondence between the different series is good. The largest gaps between the OCC1950-based series and the VAR-adjusted series occur among nonroutine–cognitive workers (top left panel) and routine–cognitive workers (top right panel) in the period between the 1983 and 2003 classification changes. We hope to learn more about why these discrepancies exist, but in any event they become much less important after the classification change of 2003. So while it might be conceptually cleaner to use the OCC1950-coded microdata in our empirical work below, the occasionally abrupt changes in these series when classifications change indicates that they are not perfect, either. Given the good match between our aggregates and the BLS aggregates, as well as the generally good fit between our individual series and the OCC1950-based series for most of their period of overlap, we feel comfortable in using the VAR-adjusted series below.[11]

by the CPS (often called the household survey). These data should not be confused with the BLS's separate *establishment* survey, which does not collect data from nonagricultural firms. Polarization studies generally ignore agricultural employment, so we did not collect agricultural data from the CPS-based historical sources, even though these data are available. Recall also that the BLS does not currently publish any occupational data prior to January 1983. Consequently, before that date, we can match our implied total only to total nonagricultural–industry employment, not to nonagricultural–occupation employment. Fortunately, Figure 3 shows that the fits between either type of BLS aggregate and our implied totals are good.

[10] See https://usa.ipums.org/usa/volii/occ_ind.shtml for details of the IPUMS OCC1950 designation.

[11] Aside from the VAR adjustments, we made two other small corrections to the four employment series. First, data for 1953q3 were for some reason unavailable in the relevant issue of the Census Bureau's *Monthly Report on the Labor Force*. We therefore estimated a simple linear regression for each employment series using a trend and quarterly dummies and data through 1971q1 (the date of the first major classification change). We then imputed the 1953q3 observations for each of the four series using the resulting predicted values as well interpolations of residuals from 1953q2 and 1953q4. Second, the BLS changed the lower age limit for most employment and labor force series from 14 to 16 when it undertook a significant revision of the CPS in January 1967. Fortunately, however, the BLS then began to publish occupation-specific employment levels for 14-15-year-olds, doing so through 1982. We estimated an employment-to-population model for 14-15-year-olds over the 1967-1982 period, using the employment-to-population ratio of 16-17-year-olds and the overall unemployment rate as explanatory variables. We then estimated separate models for the four occupational *shares* for 14-15-year-olds using the aggregate unemployment rate as a regressor. These shares turn out to be remarkably constant, with about three-quarters of 14-15-year-olds employed in the low-skill, nonroutine–manual group throughout the 1967-1982 period. Multiplying the actual 14-15-year-old population level with the estimated employment-to-population ratio for this group, then multiplying the resulting product with the relevant estimated occupational share gives the 14-15-year-old employment in each of the four occupational categories starting in 1947q3. These employment levels can then be subtracted from the published 14+ historical data to generate 16+ data consistent with data found in BLS publications starting in 1967.

2.2 Changes in cyclical responses of occupational employment to GDP

Once we have constructed the four employment series we can investigate their cyclical properties. **Figure 5** graphs each of the series along with an HP trend that is generated with a smoothing parameter of $\lambda = 100,000$.[12] The HP-detrended series are then seasonally adjusted by regressing them on quarterly dummies and appropriate interactions.[13] A simple way to investigate the cyclical relationship between the resulting four employment series is with the following dynamic factor model (DFM):

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

$$E_{it} = \alpha_i + \phi_i F_t + \beta_{1i} Mfg_t + \beta_{2i} ResConstr_t + e_{it}. \tag{2}$$

Here, F_t is a single unobserved common factor, which depends on GDP and which is constrained to be an AR(1) process with error term ν_t.[14] The four employment series E_{it} depend on the common business-cycle factor F_t, group-specific constants α_i, manufacturing output Mfg_t (available from the Federal Reserve Board's industrial production release), real residential construction $ResConstr_t$ (taken from the national accounts), and an error term e_{it}. To estimate this model, we assume that the disturbance terms ν_t and e_{it} are normally distributed, which allows joint estimation of the system as a state-space model via maximum likelihood and the Kalman filter.[15]

Figure 6 takes the first step in the analysis by plotting the estimated F_t from a DFM estimated over the entire sample period along with each of the four E_{it}'s in the system.[16] The most striking takeaway from this figure is the close agreement between the business-cycle factor F_t and the E_{it} series corresponding to routine–manual employment in the lower left panel. There is no evidence that this correlation becomes any stronger when either job-market polarization or jobless recoveries become important—it is strong throughout the entire

[12]Note that this choice of a smoothing parameter λ is larger than those typically chosen in macroeconomic work. But to us this choice of λ seems to give the best decomposition between trend and cycle for these particular series. We use the same choice of λ when detrending all other variables in our empirical models below (such as GDP).

[13]Because the Census Bureau reported occupational employment data only for the first month in each quarter through 1957, we interact the quarterly dummies in the seasonal adjustment regression with a $YEAR < 1958$ indicator to allow a different seasonal cycle over this period. Additionally, visual inspection of the nonroutine-manual category in the lower right panel of Figure 5 indicates that the seasonal cycle for this group appears to change after the classification change in 1983. We therefore interact the quarterly dummies in each of the four seasonal-adjustment regressions with a $YEAR >= 1983$ indicator as well.

[14]The variance of ν_t is normalized to equal one.

[15]Maximum likelihood (ML) estimation is feasible here because the number of employment series in the system is small. ML estimation 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 principal components analysis (PCA), then treat them as data for further analysis (Stock and Watson 2011). For an application of this approach, see Stock and Watson (2012).

[16]Thus, the same F_t is pictured in each of the four panels of Figure 6.

sample. A second takeaway from Figure 6 is the rising correlation of the estimated F_t with both nonroutine–cognitive employment (top left panel) and routine–cognitive employment (top right panel). As time passes the high-frequency movements in both types of cognitive employment become more closely synchronized with the overall business cycle summarized by F_t. Finally, the lower right panel shows that high-frequency movements in low-skill, nonroutine–manual employment appear uncorrelated with the business cycle throughout the sample period.

The DFM can also generate dynamic predictions for individual employment series during various historical episodes conditional on the observed paths of GDP, manufacturing output, and residential construction. Obviously, these forecasts are not structural: output depends on employment just as employment depends on output. But by estimating DFMs over different time periods, and then using the estimated parameters to generate forecasts over a single historical period, we can determine whether the response of any employment series to real-side variables has changed over time.

Figure 7 presents dynamic forecasts for occupation-specific employment levels during the most recent three recessions and recoveries.[17] In each panel, actual employment data and forecasts from three DFMs are presented: one from the DFM estimated on the entire sample, another estimated on an early sample that ends in 1985q4, and a third estimated on a later sample that begins in 1986q1. In each panel, the vertical line corresponds to the quarter that both the recession and the dynamic forecast begin. Each row of the figure corresponds to the same business cycle, and all panels in a given row share the same vertical scale.

Consider first the dynamic forecasts for routine–manual employment, which appear in the third column of the figure. Because this employment series is relatively cyclical, both actual employment and the dynamic forecasts tend to decline in recessions more than they do for other categories. Of more interest is the close agreement among all three dynamic forecasts for routine–manual employment in most panels. In particular, the graph in the last row of the third column shows that routine–manual employment fell sharply and remained low during and immediately after the Great Recession of 2007-09. The similarity of the dynamic forecasts in this panel shows that given the actual changes in real-side variables, the time path for routine–manual employment in the Great Recession and the recovery that followed should not have been surprising. The model estimated on the post-1985 data generates a dynamic forecast for routine–manual employment that is almost identical to the forecasts implied by the other DFMs, including the one generated by pre-1985 data. It is therefore difficult to blame this particular jobless recovery on a change in the response of routine–manual employment to real-side data after the mid-1980s.

[17] While the dynamic forecasts use actual values of GDP and manufacturing and construction output, they do not use any data on employment once the dynamic forecasts begin.

A potential criticism of the DFM approach is that subtle changes in the structure of the baseline DFM might generate different results. We experimented with various versions of the model to check robustness, for example by treating GDP as a right-hand-side variable in all regressions and running the model as four individual regressions rather than as a DFM. The results were similar to those graphed in Figure 7. Perhaps the simplest way to prove robustness, however, is to dispense with any type of model and simply graph the rolling correlations of the four employment series with GDP. **Figure 8** does just this, by depicting the rolling 10-year (40-quarter) correlation of the four employment series with aggregate output. The lower left panel shows that routine–manual employment has been very closely correlated with GDP throughout the sample period. At the other extreme, the correlation of nonroutine–manual employment with GDP (lower right panel) has hovered around zero. The top two panels show that the correlation of the two cognitive categories has grown over time, with the correlation for routine–cognitive employment (top right panel) growing somewhat sooner that for nonroutine–cognitive employment (top left panel).

2.3 Implications of occupational-employment patterns

The analysis of detrended employment levels shows that changes in the cyclicality of various employment series do not line up neatly with long-run polarization trends. Polarization has favored the nonroutine workers at the ends of the skill distribution, but changes in the cyclicality of high- and low-skill employment have not been uniform. High-skill, nonroutine–cognitive employment has become more cyclical over time while low-skill nonroutine–manual employment has remained acyclical. Additionally, there appear to be no significant changes in the cyclical behavior of middle-skill, routine–manual employment, a group that has been significantly and adversely affected by polarization trends. Taken together, these patterns suggest that a recessions-as-reallocations explanation of jobless recoveries, of the type suggested by Jaimovich and Siu (2012), may have a hard time explaining the data.

3 Analysis of Occupational Reallocation Since 1976

3.1 Data issues

In this section, we analyze another component of the recessions-as-reallocations hypothesis using CPS microdata that has been matched across consecutive months for a given worker. The CPS includes information on each unemployed worker's last job (for those workers who have been previously employed), so we can measure how often unemployed workers in one skill class take jobs in different skill classes. In particular, how likely is it that unemployed middle-skill workers enter other skill groups? And does this type of reallocation have a systematic cyclical pattern?

As was the case with the employment series, a reallocation study at the worker level requires some discussion of data issues. Because this study is based on CPS microdata, we can be more flexible in the assignment of workers to different classifications than we were when working from published historical aggregates. We follow Autor (2010) and Acemoglu and Autor (2011) in aggregating an occupational system originally devised by Meyer and Osborne (2005) into the three broad groups: high-skill, middle-skill and low-skill. In an earlier analysis of this approach (Foote and Ryan 2012) we found that using this classification revealed obvious seams in *employment levels* whenever the BLS classifications change, just as we did above when using the BLS's published occupational data (and the OCC1950 code in an earlier robustness check). Yet we also found that these level shifts were not apparent when looking at *unemployment rates and worker flows*, suggesting that the workers who are "moved around" when the classifications change behavior similarly to those left behind. We therefore feel confident that the classification system matters less for the study of worker-level reallocation than it does for the study of high-frequency changes in detrended employment levels.

3.2 Reallocation of unemployed workers

Figure 9 provides direct evidence on reallocation propensities by showing where unemployed workers from different skill classes tend to go when they leave the unemployment pool. As noted above, we start with the grouping of workers into high-skill, low-skill, and middle-skill classes, but we further disaggregate the latter group along industry lines into middle-manufacturing, middle-construction, and middle-other groups. The top panel of Figure 9 shows the destinations of workers in the five industry-skill classes *conditional* on some type of exit from unemployment to employment. As we would expect, the far left graph shows that unemployed high-skill workers tend to take high-skill jobs, while the far right graph shows that unemployed low-skill workers tend to take low-skill jobs. But the three middle graphs of the top panel show that unemployed middle-skill workers are even more likely to exit unemployment for jobs in their previous skill class. There is some evidence that unemployed workers of any type are increasingly less likely to take middle-skill jobs as the years pass, consistent with the key implication of polarization: the share of total employment accounted for by middle-skill jobs is declining. However, there appear to be no pronounced changes in the probability that workers of any type will take middle-skill jobs when recessions occur. The bottom panel of Figure 9 shows reallocation propensities allowing for the unemployed worker to exit either to employment or to nonparticipation. Interestingly, the share of middle-skill workers who exit unemployment for nonparticipation is rising over time. Additionally, there appears to be cyclical pattern to these exits, as they rise during recessions and early recoveries. However, a similar pattern is also present for high-skill workers, making it hard

to blame this U-to-N pattern on polarization *per se*.

In some earlier work, we have estimated discrete-choice models of employment reallocation among unemployed workers, which allow us to construct precise probabilities for how reallocation probabilities are affected either by macroeconomic forces (such as the overall job-finding rate) or by individual-specific influences (including demographic variables and unemployment duration).[18] We would certainly be willing to present these models at the *Macro Annual* conference. But the bottom line of the more formal economic models remains the same as in Figure 9: reallocation across skill groups does not seem to change much over the business cycle, especially for middle-skill workers.

4 Occupation-Specific Unemployment Rates

We believe that a study of the cyclical aspects of labor-market polarization (or lack thereof) would interest a variety of macroeconomists, including those interested in the recessions-as-reallocations literature, those who examine the conditional correlations of aggregate employment and productivity with output, those who study the macroeconomics of the labor market more broadly, and those concerned from a policy standpoint about potential structural aspects of current the unemployment rate. But we think that there is a final piece of the puzzle that needs to be filled in before our study is complete: an analysis of occupation-specific unemployment rates (URs).

We have already begun to augment the employment data with URs, which also come from historical sources. **Figure 10** gives a progress report in this effort. The unemployment rates between 1958 and 1972 are based on seasonally adjusted data published in the February 1973 BLS *Employment and Earnings*, which happens to feature historical, seasonally adjusted data for many CPS-based series. (Data for later years are available in other *Employment and Earnings* issues and is now being coded.) Starting in 1976, the figure depicts the unemployment rates generated by the CPS microdata coded with the OCC1950 designation. The top panel of the figure shows the URs in level form, along with the overall UR reported by BLS, while the lower panel depicts the occupation-specific rates as fractions of the overall rate. One interesting feature of the lower panel is that the ratio of the low-skill, nonroutine–manual unemployment to the overall UR declines in every recession since 1970, consistent with the relatively low cyclical sensitivity for employment levels we found for that sector. Additionally, in the same panel, the ratio of high-skill-to-overall unemployment rises in the past two recessions, consistent with the rising cyclicality of high-skill employment we have previously seen.

Yet the unemployment rates also provide some additional perspective on the employment

[18] For details of this research, see Foote and Ryan (2012).

data discussed above. For example, during the recessions of the late 1950s and early 1970s, the ratio of high-skill unemployment to overall unemployment changed little, indicating that high-skill unemployment rose proportionately with the unemployment rate of the overall economy. According to our earlier analysis, however, this period was not one in which changes in high-skill employment changes were particularly well-aligned with changes in GDP. What accounts for this seeming discrepancy in the cyclicality of high-skill employment and unemployment rates? And is this discrepancy related to the an ongoing debate on whether the Okun's Law relationship between output, unemployment, and productivity has changed?[19] We plan on investigating these questions in the weeks ahead.

[19]For disparate views on this question, see Gordon (2010) and Ball, Leigh, and Loungani (2013).

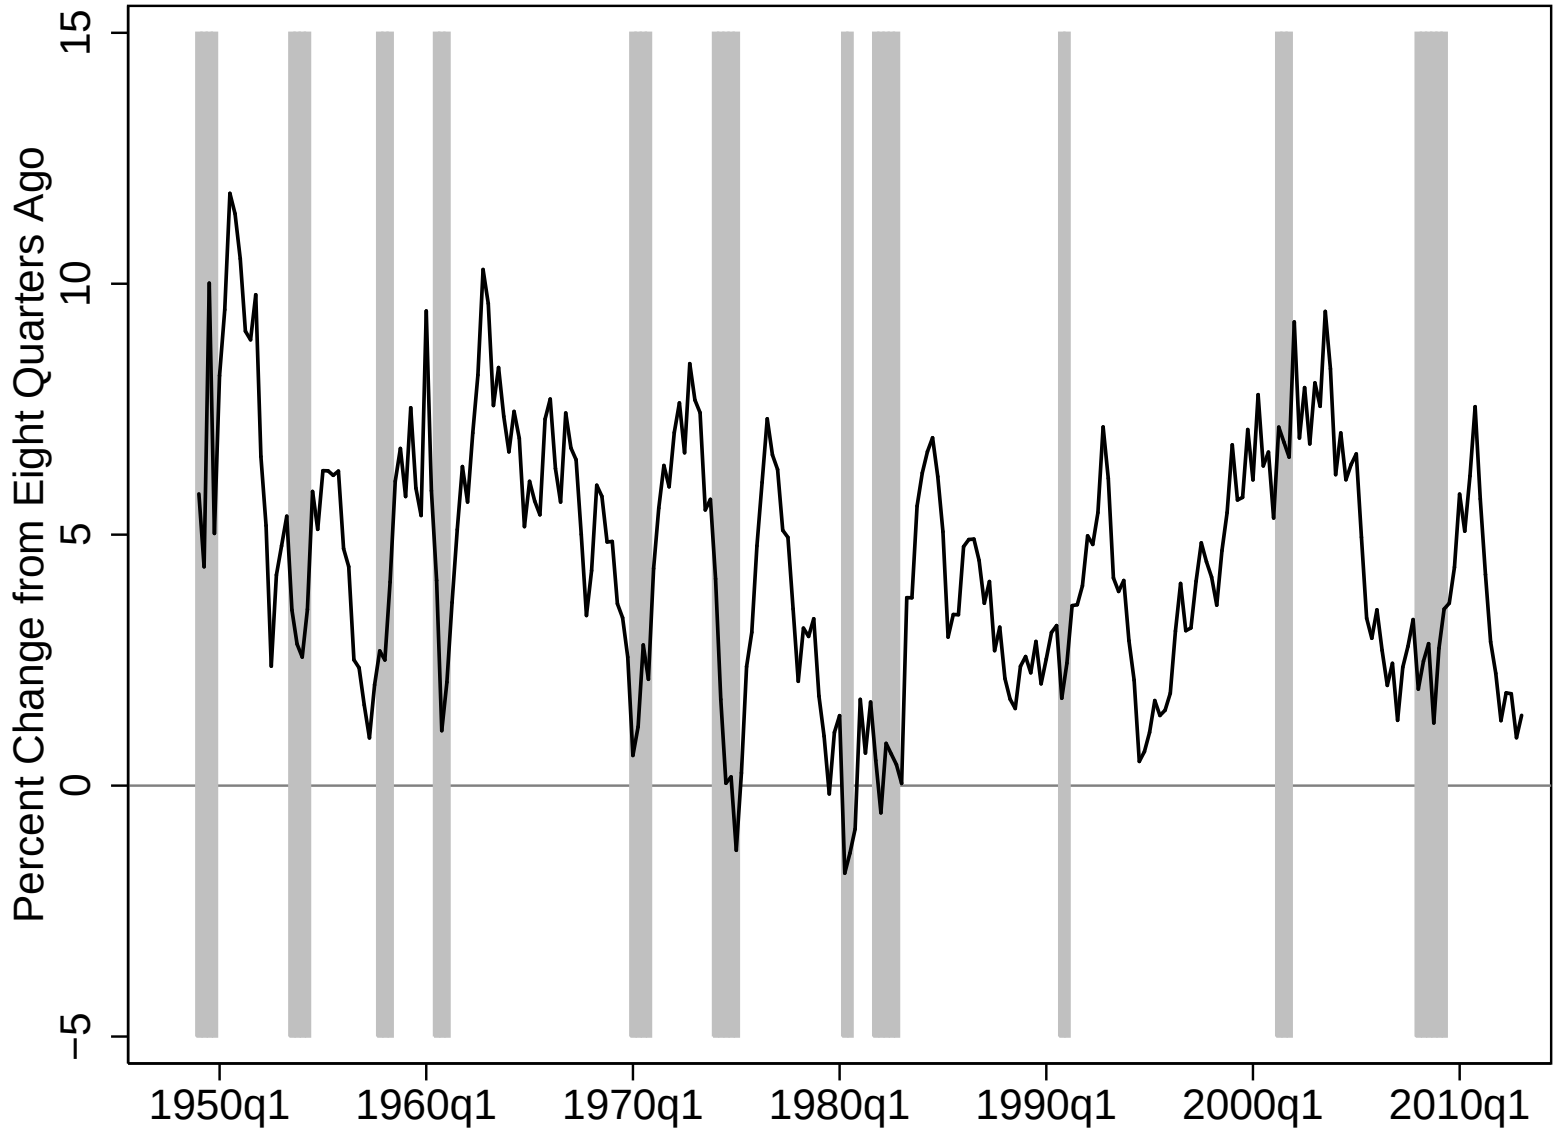


Figure 1. Cyclical Movements in Average Labor Productivity in the Nonfarm Business Sector.

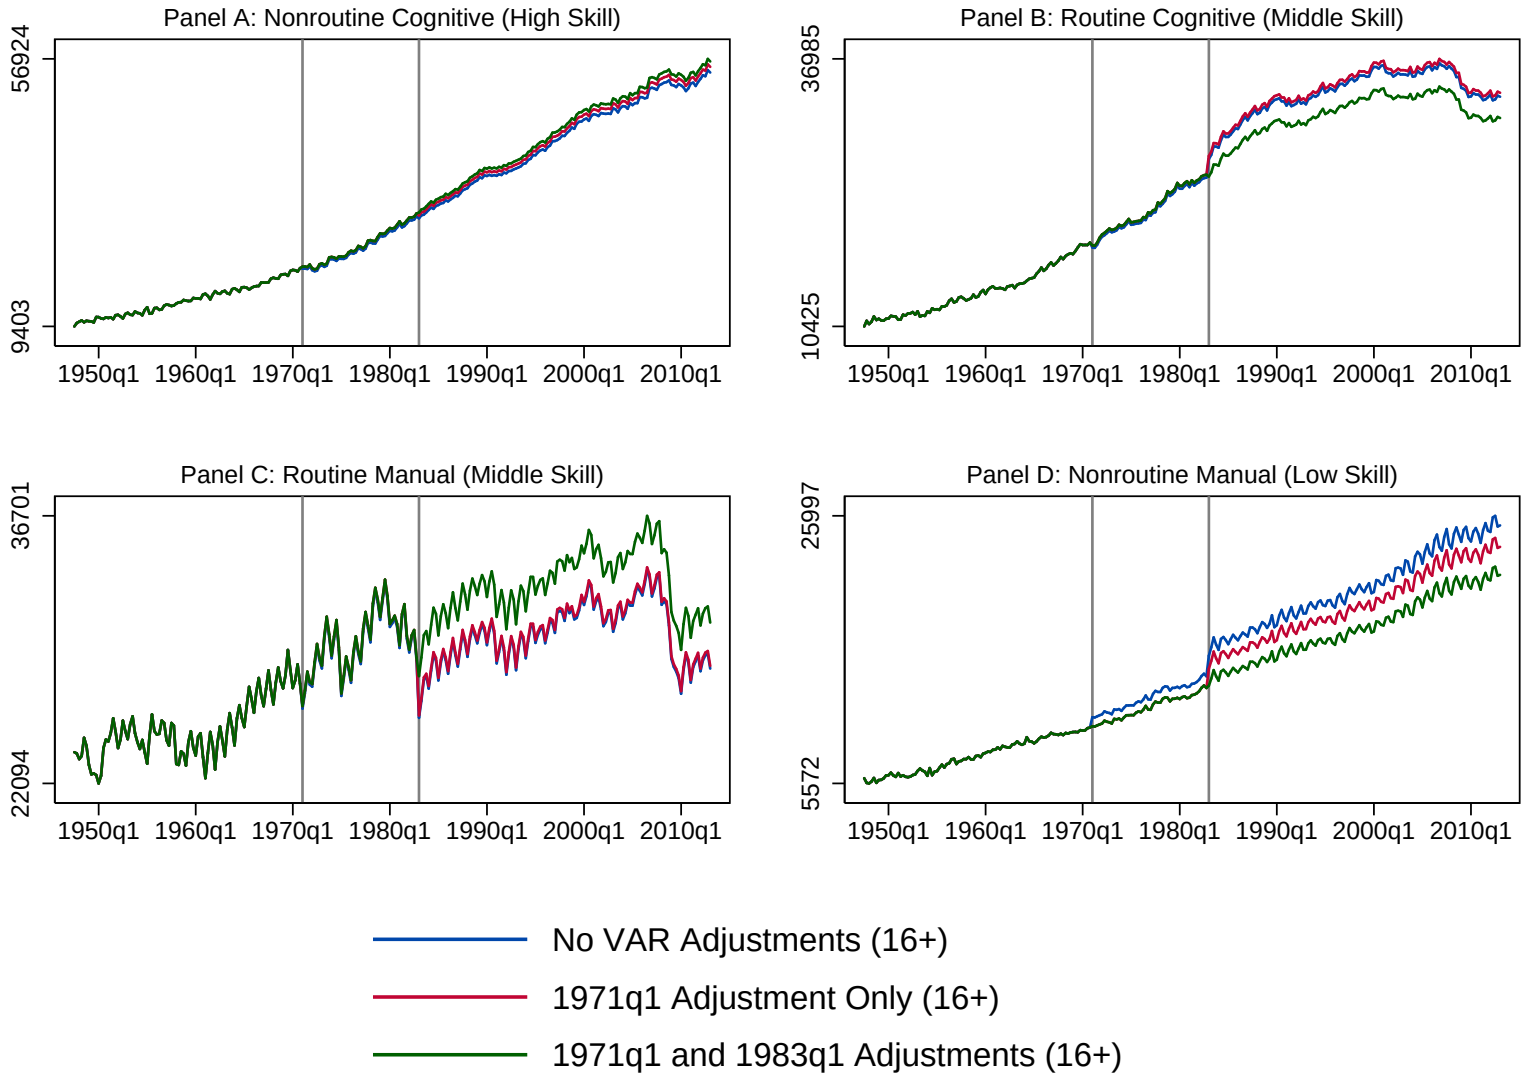


Figure 2. Effects of VAR-based Adjustments on Occupation-Specific Employment Series. Note: Data correspond to persons aged 16 and older. Vertical lines correspond to BLS occupational classification changes in 1971q1 and 1983q1.

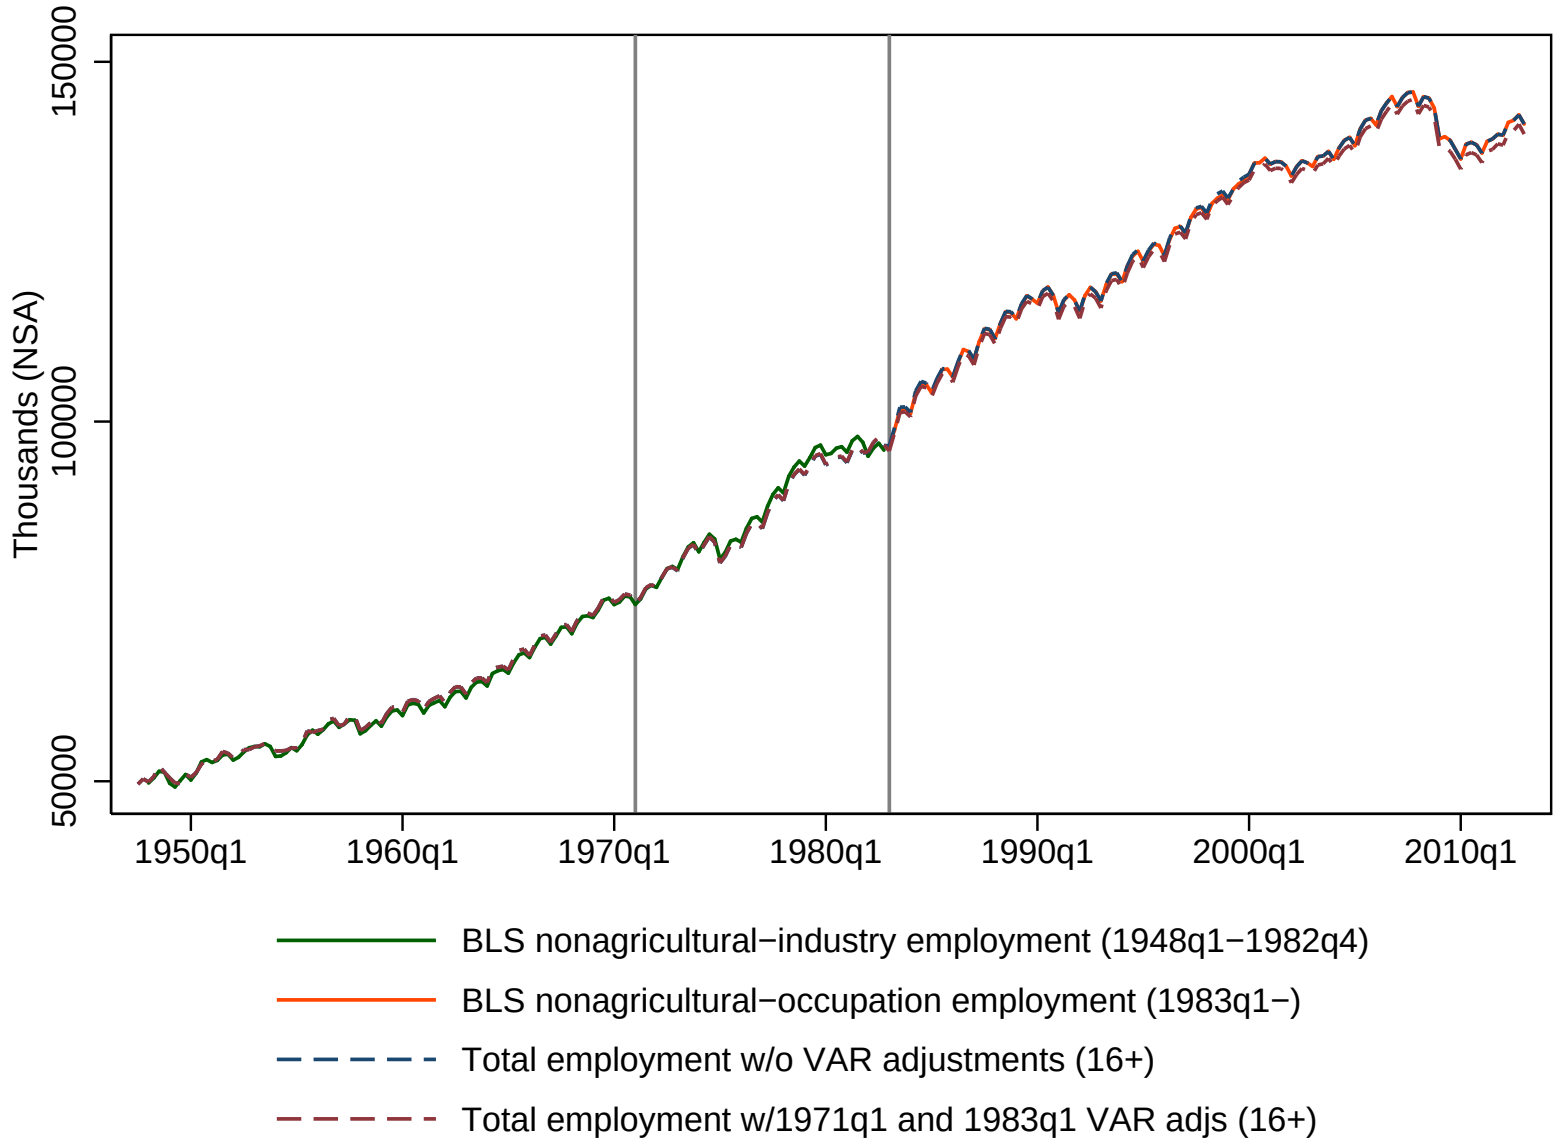


Figure 3. Comparison of BLS-Reported Nonagricultural Employment Totals and Aggregation of Occupation-Specific Employment Series. Note: Data correspond to persons aged 16 and older. Vertical lines correspond to BLS occupational classification changes in 1971q1 and 1983q1.

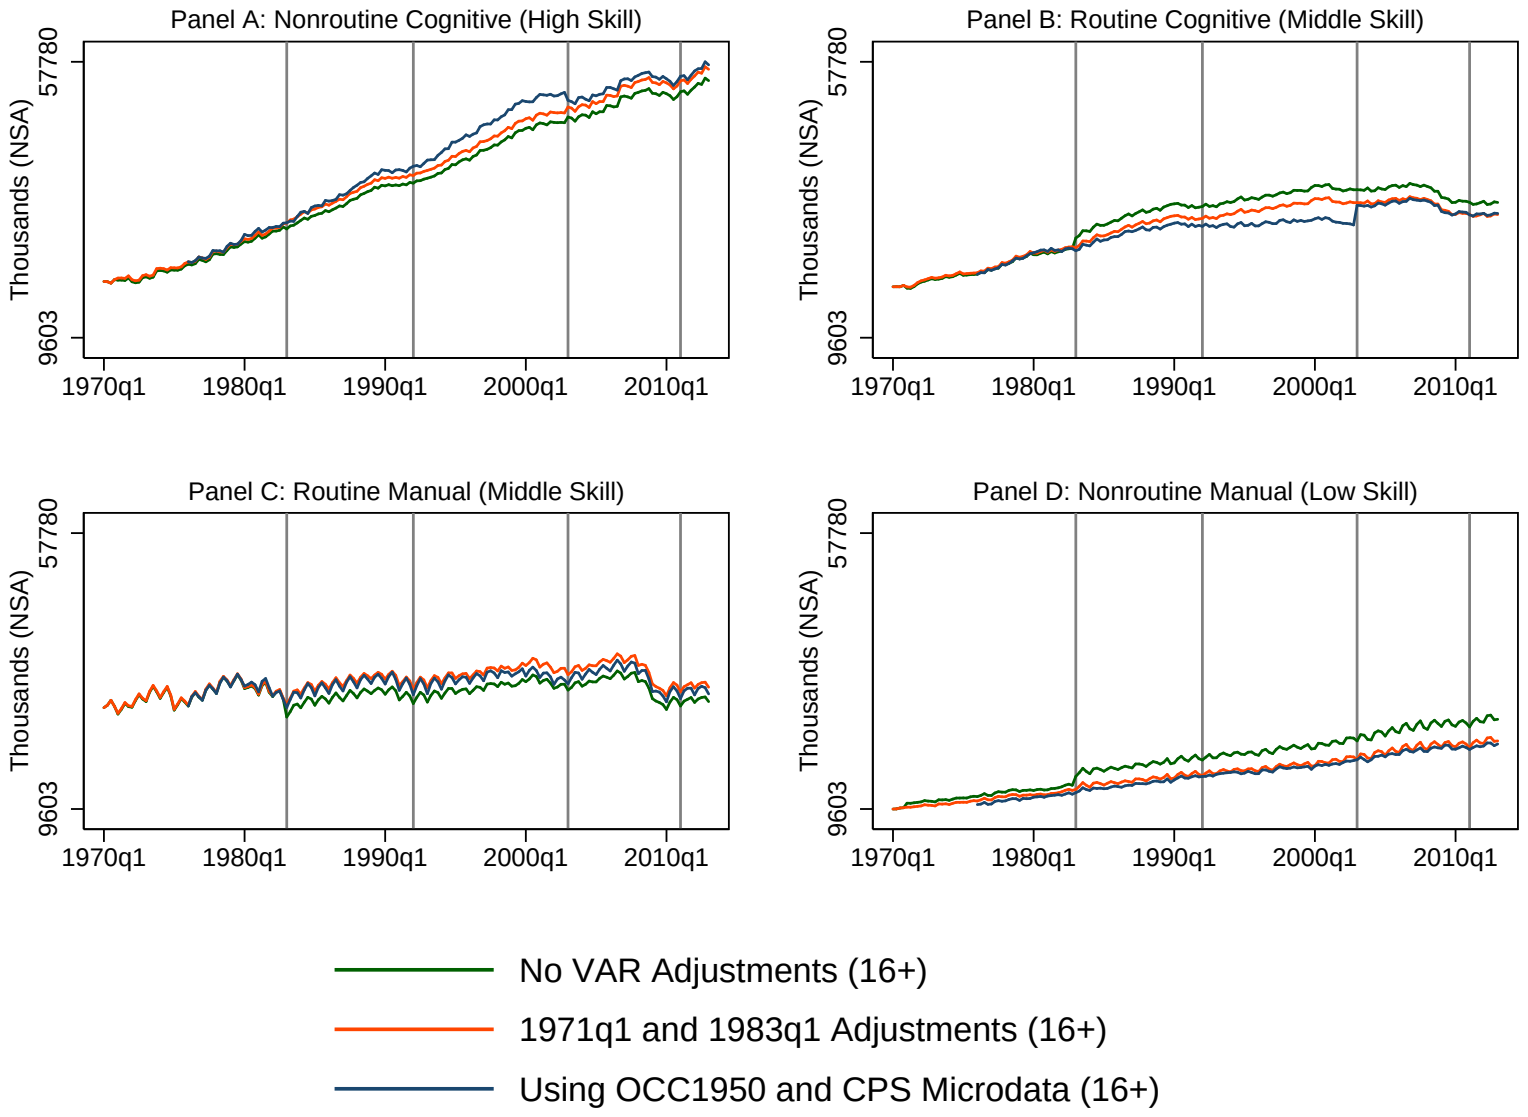


Figure 4. Comparison of VAR-Adjusted Employment Series and Alternative Series Constructed Using IPUMS–OCC1950 Designation and CPS Microdata. Note: Data correspond to persons aged 16 and older. Vertical lines correspond to all BLS occupational classification changes starting in 1983q1.

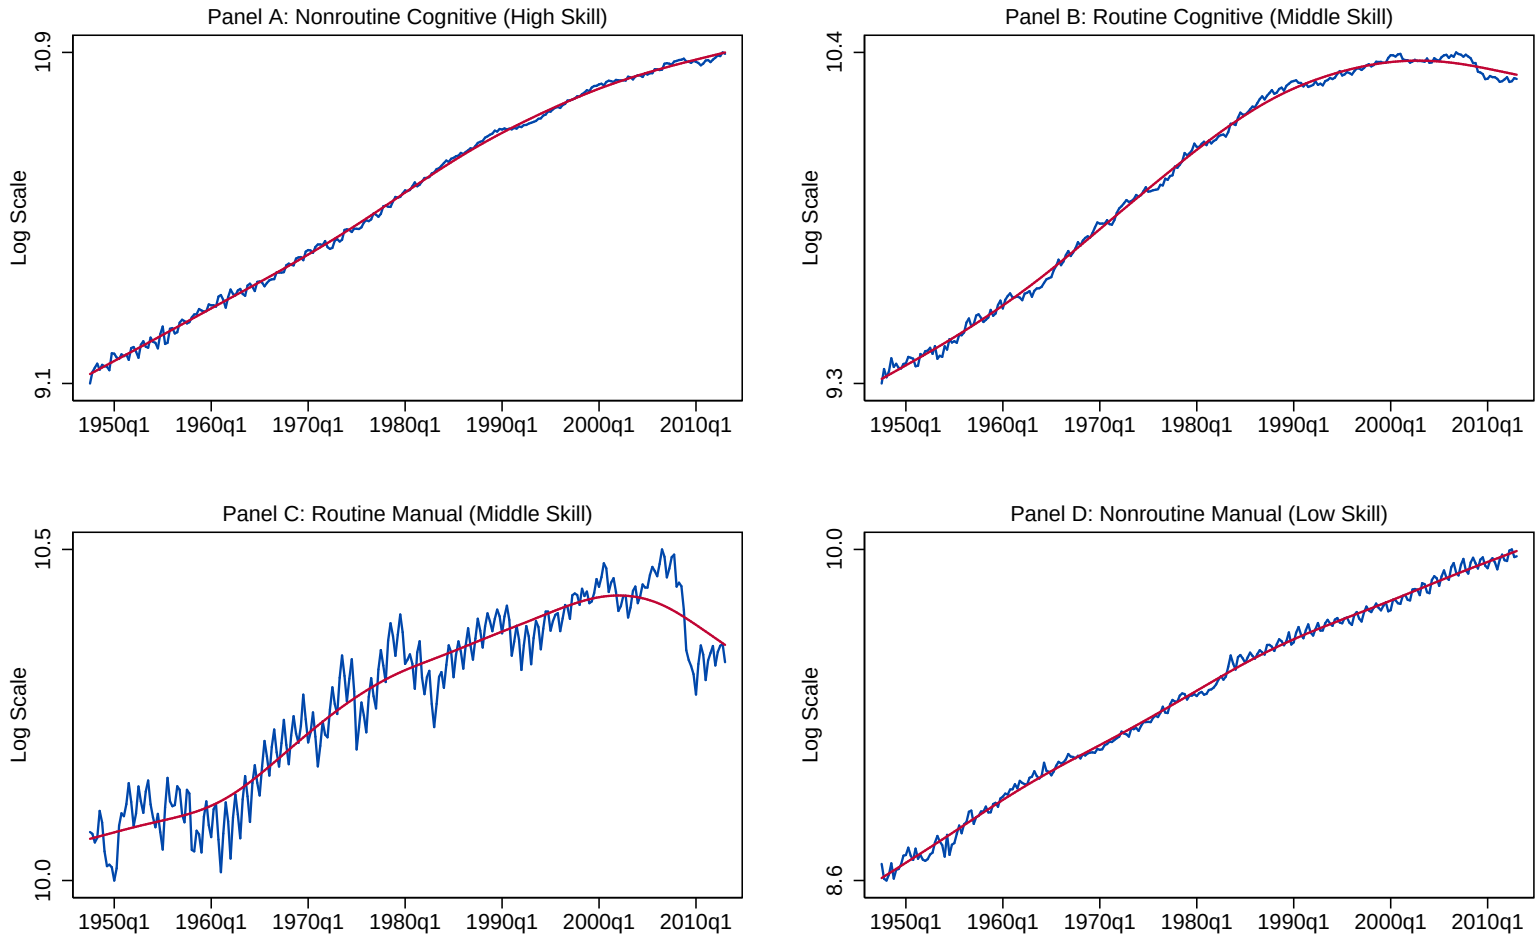


Figure 5. HP Trend ($\lambda = 100,000$) for Occupation-Specific Employment Data. Note: Data are not seasonally adjusted.

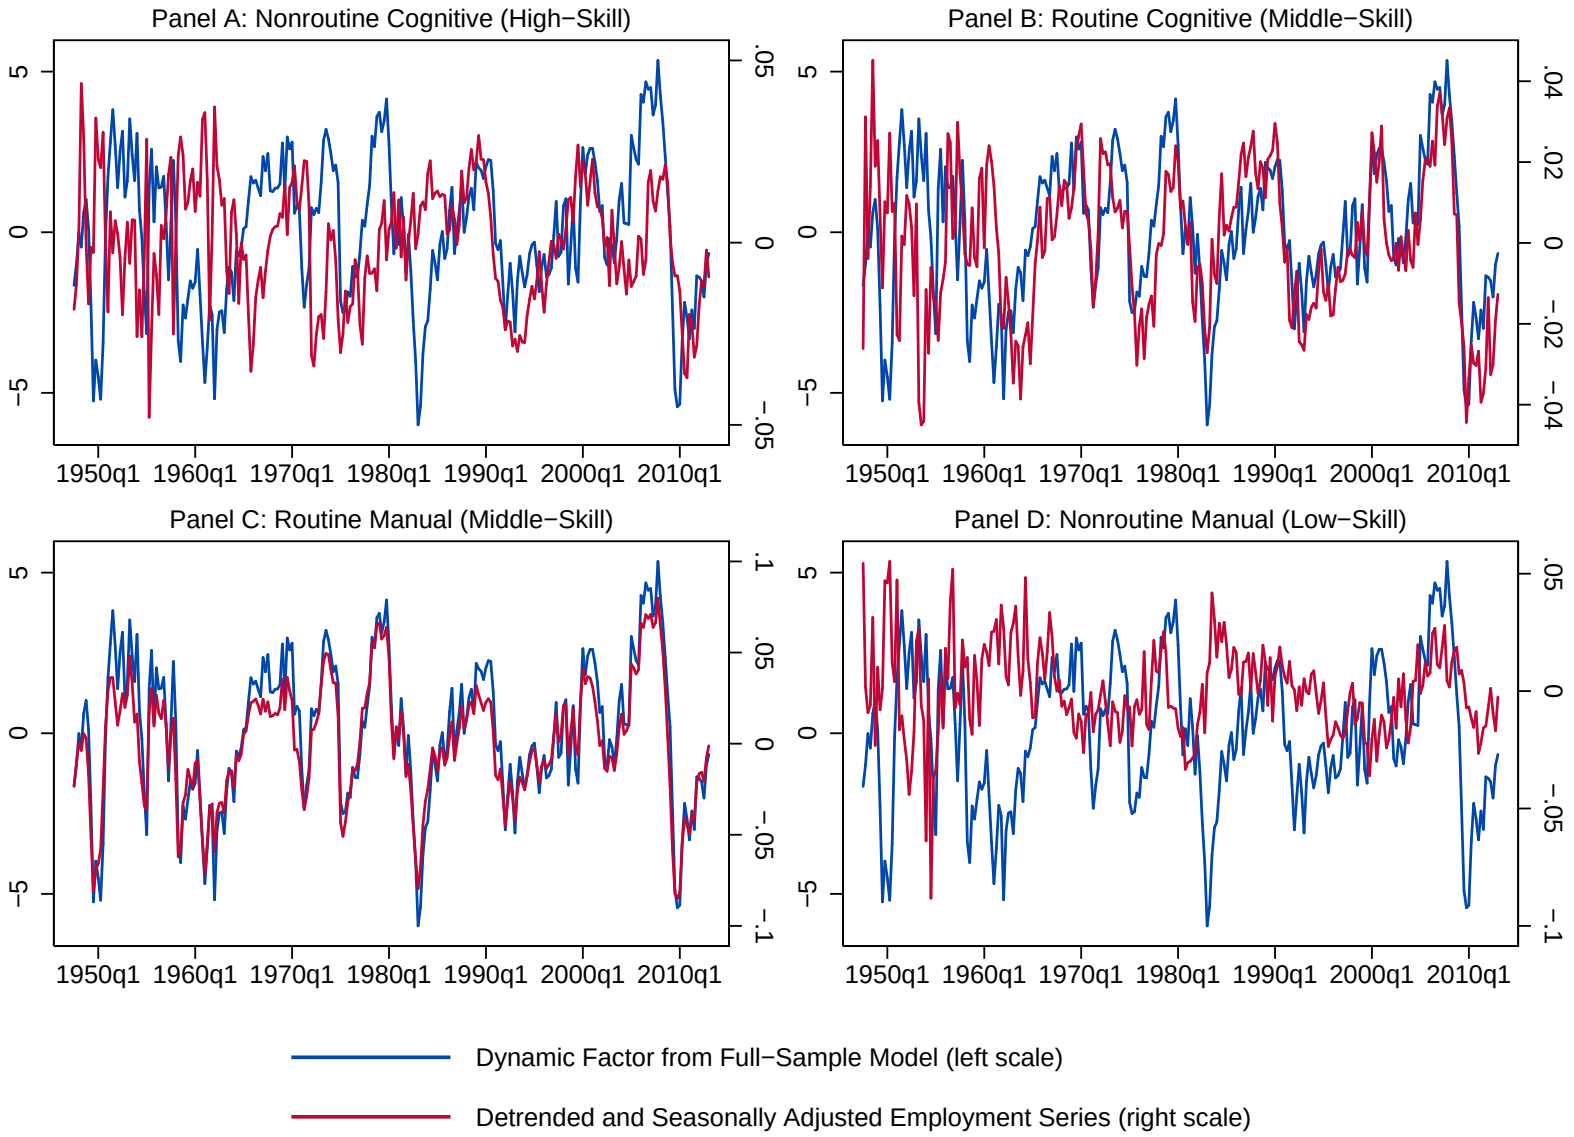


Figure 6. Comparison of Estimated F_t from Dynamic Factor Model and Detrended and Seasonally Adjusted Occupation-Specific Employment Series.

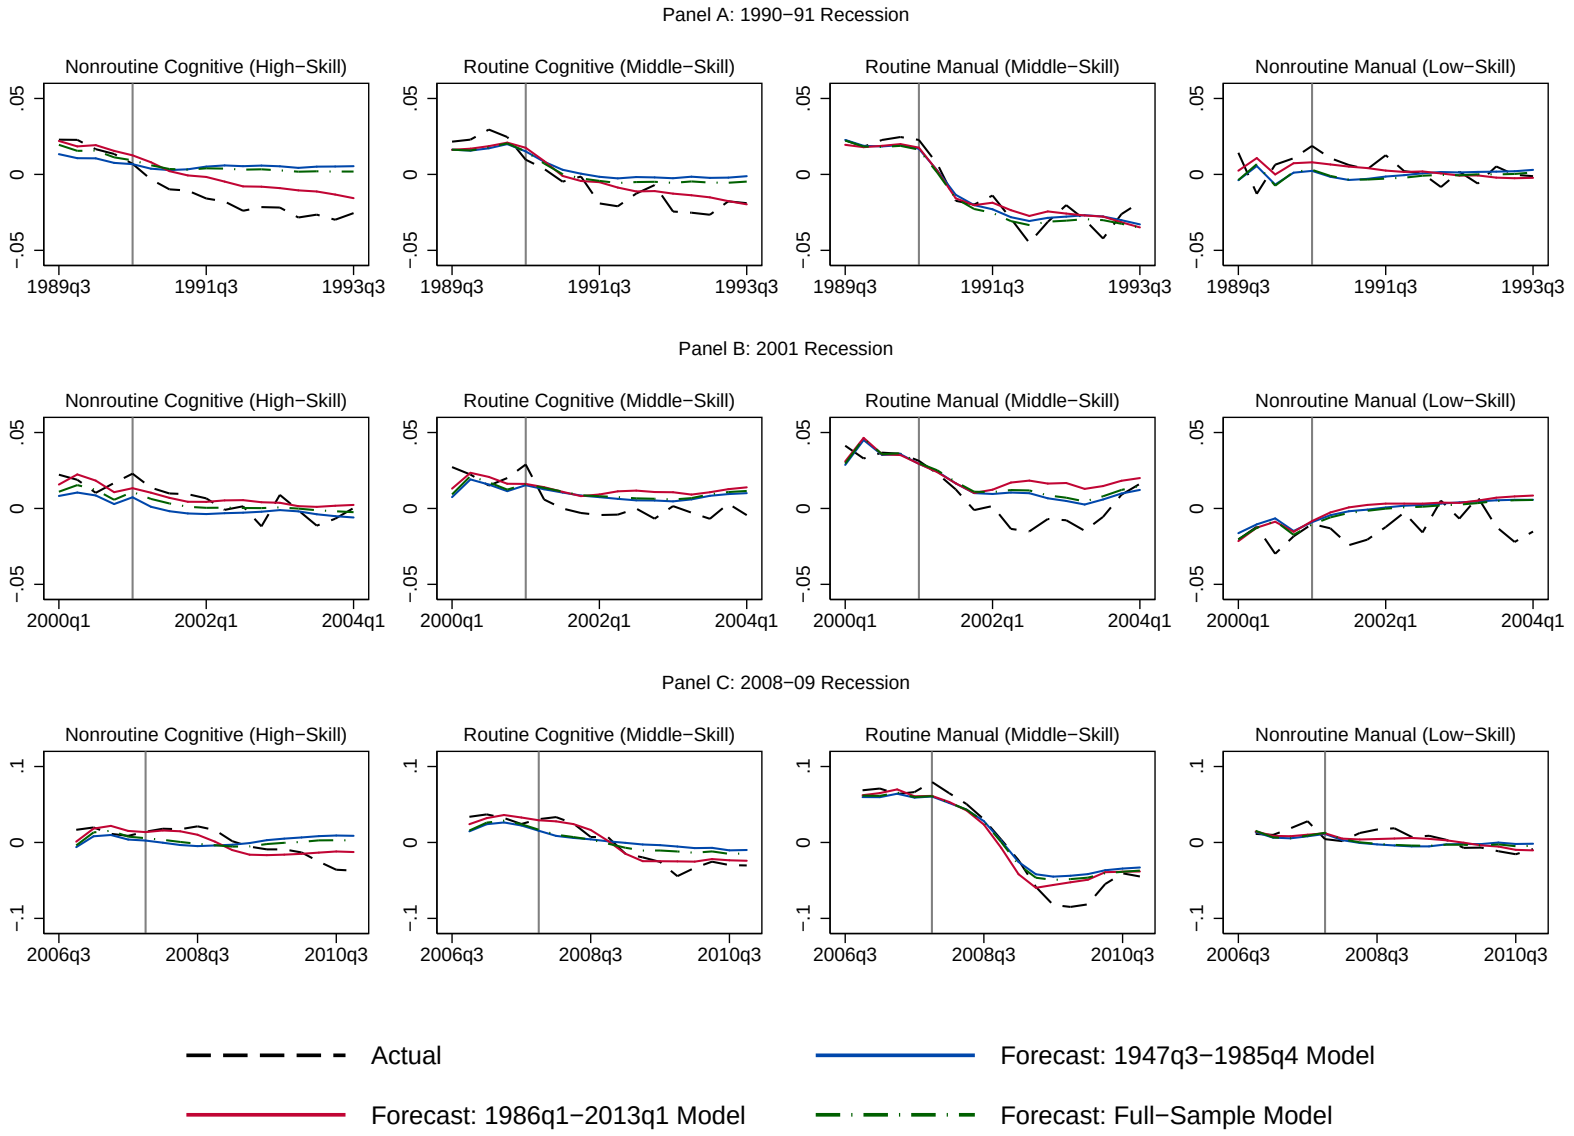


Figure 7. DFM-Generated Dynamic Forecasts for Occupation-Specific Employment in Three Recent Recessions and Early Recoveries. Note: Dynamic forecasts begin with the onset of recessions, which are marked with vertical lines.

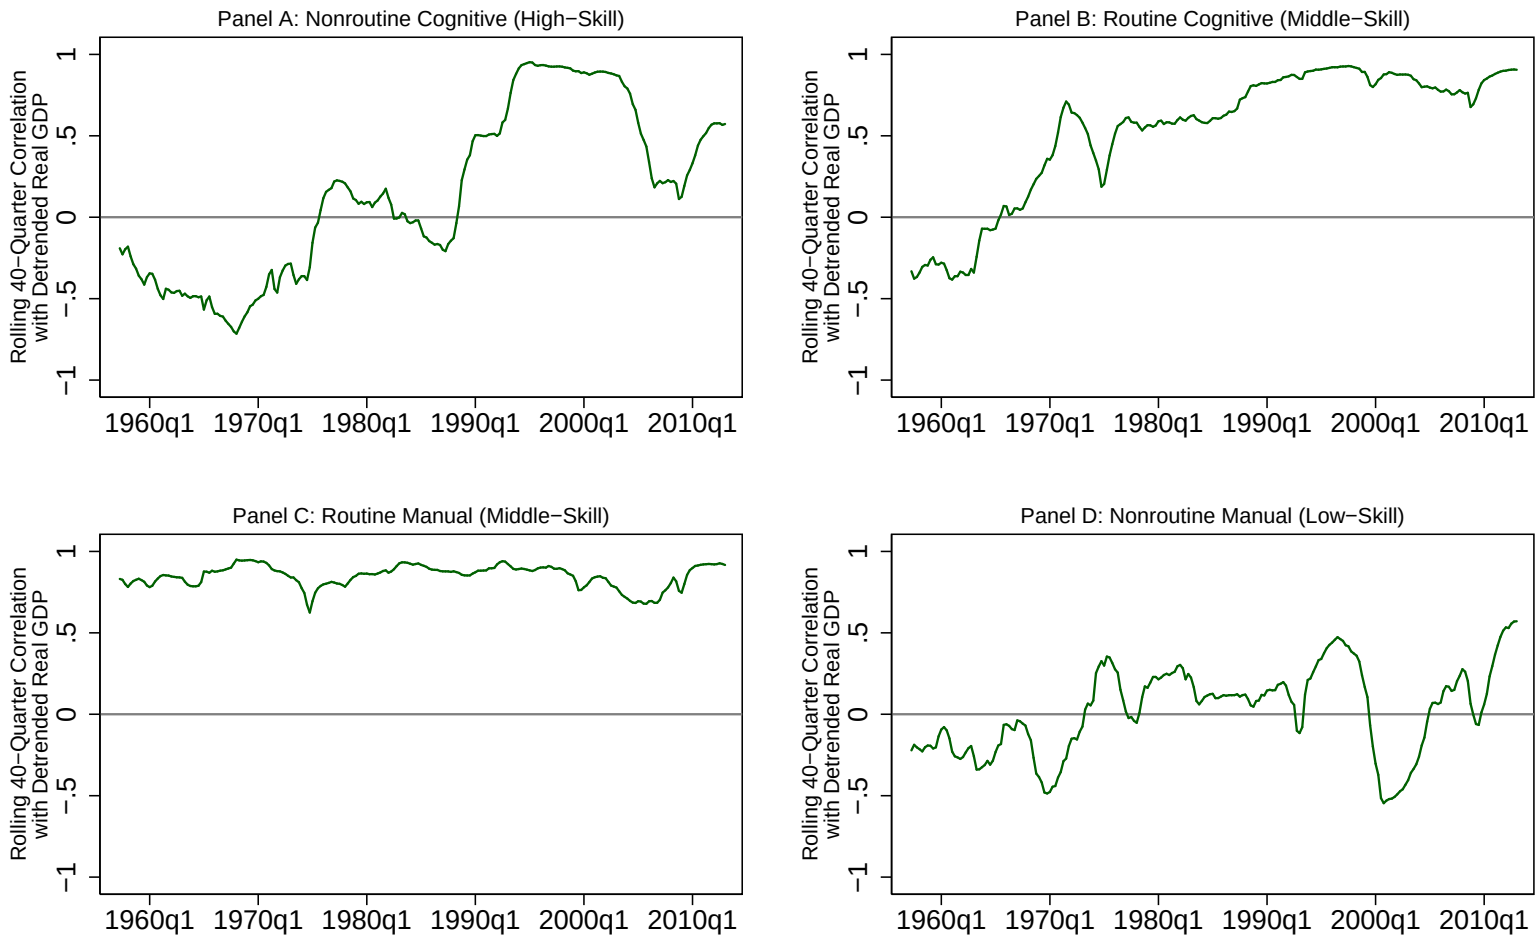


Figure 8. Rolling 10-Year Correlations of Detrended and Seasonally Adjusted Occupation-Specific Employment Series with Detrended GDP.

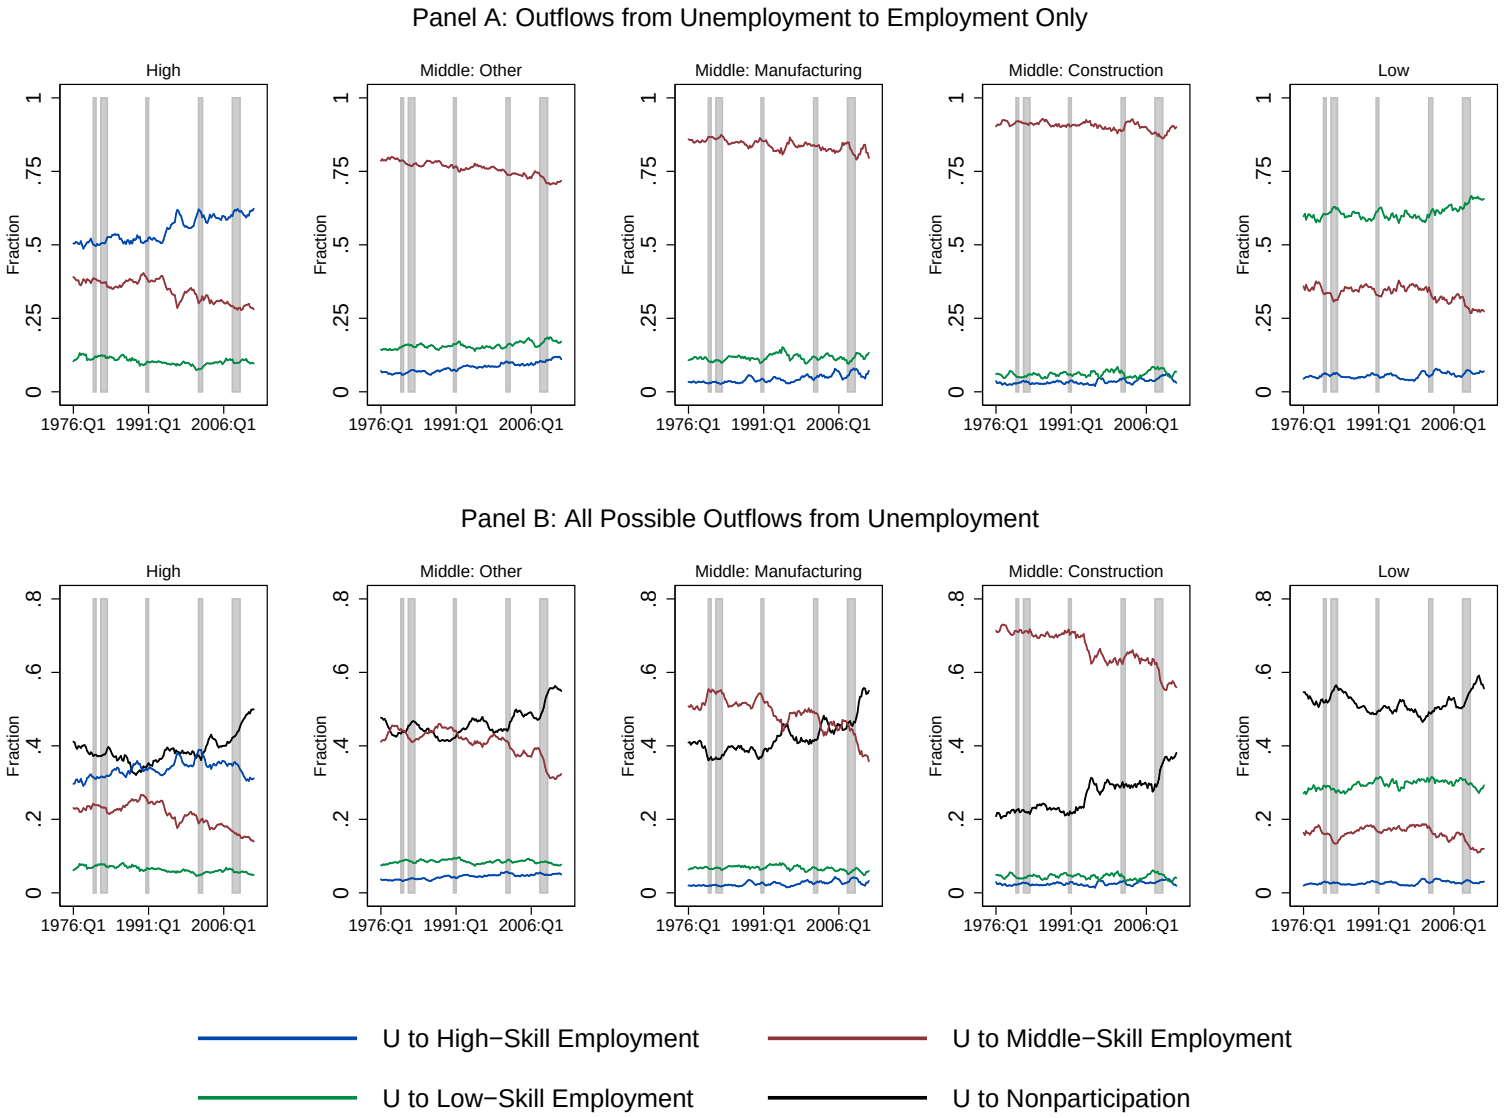


Figure 9. Reallocation Patterns of Unemployed Workers in Matched CPS Microdata.

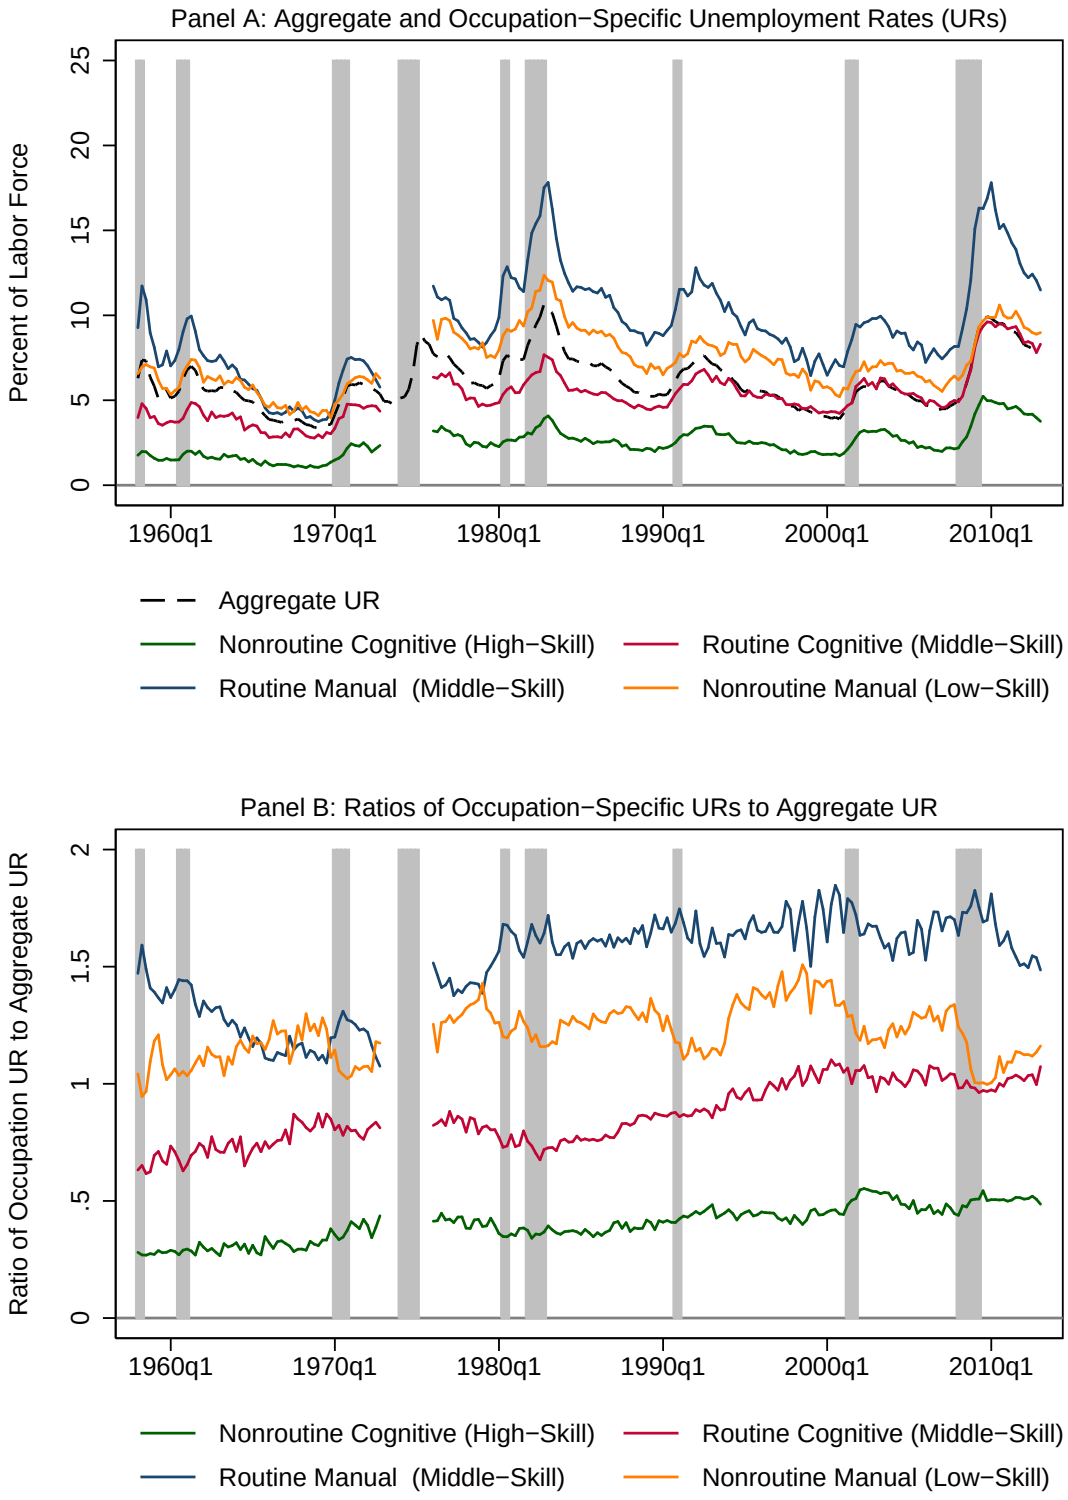


Figure 10. Occupation-Specific Unemployment Rates. Note: Data from 1958-1972 come from BLS's *Employment and Earnings* and are seasonally adjusted by BLS. Data starting in 1976 are generated from CPS microdata and the IPUMS-OCC1950 occupational classification and are seasonally adjusted by the authors. (Data from 1973-1975 are available in various BLS publications and are now being coded.)

References

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.

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.

Ball, Lawrence M., Daniel Leigh, and Prakesh Loungani. 2013. "Okun's Law: Fit at Fifty?" National Bureau of Economic Reseach Working Paper No. 18668. Available at http://www.nber.org/papers/w18668.

Castro, Rui, and Daniele Coen-Pirani. 2008. "Why Have Aggregate Skilled Hours Become So Cyclical Since the Mid-1980s?" *International Economic Review* 49(1): 135–184.

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.

Foote, Christopher L., and Richard W. Ryan. 2012. "Labor Market Polarization over the Business Cycle." Federal Reserve Bank of Boston Working Paper No. 12-09. Available at http://www.bos.frb.org/economic/ppdp/2012/ppdp1208.htm.

Galí, Jordi, and Luca Gambetti. 2009. "On the Sources of the Great Moderation." *American Economic Journal: Macroeconomics* 1(1): 26–57.

Galí, Jordi, Frank Smets, and Rafael Wouters. 2012. "Slow Recoveries: A Structural Interpretation." *Journal of Money, Credit and Banking* 44: 9–30.

Gordon, Robert J. 2010. "Okun's Law and Productivity Innovations." *American Economic Review* 100(2): 11–15.

Hall, Robert E. 1991. "Labor Demand, Labor Supply, and Employment Volatility." *National Bureau of Economic Research Macroeconomics Annual* 6: 17–62.

Jaimovich, Nir, and Henry E. Siu. 2012. "The Trend is the Cycle: Job Polarization and Jobless Recoveries." National Bureau of Economic Reseach Working Paper No. 18334. Available at http://www.nber.org/papers/w18334.

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.

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.