

Labor-Market Polarization Over the Business Cycle*

Christopher L. Foote[†] Richard W. Ryan[‡]

April 10, 2014

Conference Draft (Revised)

Abstract

Large numbers of workers in the middle of the skill distribution lost their jobs in the Great Recession. Middle-skill occupations have also faced the weakest job growth during the past several decades due to the “polarization” of job growth at the ends of the skill distribution. To what extent are these two facts related? Search-based models of the labor market imply that firms and workers will take account of poor long-run prospects when considering efficient job separations. Rising wages in high-skill jobs may have also encouraged the efficient dissolution of middle-skill matches. Using a newly constructed dataset of historical occupation-level employment and unemployment, however, we find that recent middle-skill job losses were not far out of line with those in other postwar recessions, once the large drop in GDP is taken into account. Using more recent microdata to assess relevant alternatives for unemployed middle-skill workers, we find that few of them are able to transition directly to high-skill jobs, in large part because they lack the formal education to do so. This inability to obtain high-skill jobs, combined with the lower wages paid in low-skill service occupations such as custodial work and food preparation, suggests that large numbers of middle-skill workers will respond to polarization by dropping out of the labor force. An investigation of changes in participation rates for prime-age and older males supports this assertion. All told, these findings suggest that large middle-skill losses in recessions—including the most recent one—owe more to the cyclical nature of the industries that employ middle-skill workers rather than to optimal job-search considerations.

*Prepared for the *NBER Macroeconomics Annual* conference, April 11-12, 2014. The views in this paper are not necessarily those of the Federal Reserve Bank of Boston or the Federal Reserve System. Matthew Curtis and Ryan Ruxlow provided outstanding research assistance. For helpful comments we thank, without implicating, the editors (Jonathan Parker and Michael Woodford), Daron Acemoglu, David Autor, Ryan Michaels, Claire Reicher, Paul Willen, and seminar participants at a number of institutions. Thanks are also due to Autor, David Dorn, and Melanie Wasserman for sharing their occupational classification codes. Remaining errors are our responsibility.

[†]Federal Reserve Bank of Boston. Email: Chris.Foote@bos.frb.org.

[‡]University of Michigan. Email: RichRyan@umich.edu.

1 Introduction

The Great Recession was a widespread labor market calamity, but workers in some skill classes were hurt worse than others. Figure 1 depicts average annual employment growth for 10 consistently defined occupations over various time periods.¹ The two left-most bars for each occupation depict job growth over relatively long periods, 1979–1999 and 1999–2007. These bars show that workers in the middle part of the skill distribution experienced the smallest increase in job opportunities from 1979 to the onset of the recession in 2007. The concentration of long-run job growth at the ends of the skill distribution has been called the “polarization” of the labor market and has attracted a great deal of attention in both academic and policy circles.² Polarization is thought to stem from two complimentary forces: automation, which allows workers in routine occupations like factory work to be replaced by machines, and international trade, which encourages firms to offshore routine tasks to countries where wages are lower.³ The right-most bar in Figure 1 corresponds to the Great Recession (2007–2010).⁴ These bars reveal a striking pattern: middle-skill jobs, which have the worst long-run trends, were also the occupations experiencing the worst recessionary job losses. In this paper, we investigate whether the trend and cycle in middle-skill employment are related in a fundamental way.

A connection between poor long-run trends and large recessionary job losses is suggested by a version of the search-and-matching model typically used to study employment fluctuations. In the so-called endogenous-destruction version of this model, unemployment rises in recessions not only because the job-finding rate for unemployed workers declines, but also because the job-separation rate for employed workers rises. In this model, originated in the seminal paper by Mortensen and Pissarides (1994) and outlined in Pissarides (2000, Ch. 2), worker–firm employment matches are dissolved when their productivities fall below an endogenously defined productivity threshold.⁵ When this occurs the resulting separation is privately efficient, because both the firm and the worker agree that they are better off parting ways. Two important determinants of the productivity threshold are the *option value* of preserving the match in hopes of a future productivity improvement and the *outside alternatives* available to a worker should the match dissolve. For workers in middle-skill jobs,

¹The figure is a variant of one that originally appeared in Autor (2010).

²Examples of empirical and theoretical work on polarization include Autor, Katz, and Kearney (2008), Autor (2010), Acemoglu and Autor (2011), and Autor and Dorn (2013). In the public-policy realm, commentators often reference polarization research when discussing the “hollowing out” of the middle class.

³Jobs in the middle part of the skill distribution are thought to suffer more from automation and trade because these jobs tend to encompass more routine tasks than high- or low-skill jobs; it is hard to build a robot that can either argue a legal case or clean a house.

⁴Officially, the Great Recession ended in June 2009, but employment continued to decline until February 2010.

⁵Endogenous-destruction models are sometimes called cleansing models because they imply that negative productivity shocks “cleanse” the economy of its least-productive job matches.

both of these determinants could encourage relatively high separations in recessions. The negative trend in middle-skill job opportunities reduces the option value of maintaining a low-productivity middle-skill match. And if middle-skill workers improve their chances of landing a high-skill job by leaving their current job—for example, by getting more education, or by searching in a different job market—then the outside alternatives for middle-skill workers are enhanced when the productivity and wages of other jobs trend higher. A calibrated search model with an endogenous separation margin could therefore be consistent with the occupational pattern displayed in Figure 1.

Yet there is an alternative explanation for this pattern that does not rely on option values or outside alternatives: middle-skill workers may simply work in more cyclical industries. As we will see, the two middle-skill occupations experiencing the largest recent employment losses, production workers and operatives, are disproportionately employed in manufacturing and construction. Employment in these two industries is traditionally cyclical, and the close association of these industries with the the relatively volatile investment and consumer-durables sectors is no doubt part of the reason why.⁶ To be sure, the effect of large investment and/or consumer-durable shocks on middle-skill separations would be tempered if middle-skill wages were flexible. But an established body of literature argues that some degree of real wage rigidity is needed to explain first-order facts about the labor market, most notably the high cyclicality of the job finding rate (Shimer 2005; Hall 2005). The combination of wage rigidity and very large demand shocks is therefore an alternative explanation for the large middle-skill job losses experienced in the Great Recession.

Empirical evidence on the relative importance of outside alternatives and industry effects for middle-skill workers has both theoretical and practical relevance. Blanchard and Galí (2010) describe a growing theoretical literature that attempts to blend four elements into a general model of unemployment: labor-market frictions, real wage rigidity, sticky prices, and concave utility over consumption and leisure. Explaining occupation-level movements in employment and unemployment requires that we make this model more complex, but how? Search considerations point to the inclusion of an endogenous separation margin, which is lacking in the prototype model that Blanchard and Galí (2010) describe. As explained by Pissarides (2009), incorporating such a margin into a fully calibrated model is technically challenging, due to the difficulty of calibrating the idiosyncratic shocks needed to “spread out”

⁶Any macroeconomic theory must be consistent with the higher volatility of investment relative to consumption that is readily apparent in the national accounts. Real business cycle models imply that periods of high productivity growth are good times to accumulate capital, and the resulting surge in investment is especially large when productivity innovations are specific to the manufacture of investment goods (Fisher 2006). In sticky-price models, Barsky, House, and Kimball (2007) show that agents’ ability to delay purchases of consumer durables and investment goods makes the demand for these goods highly sensitive to monetary shocks. And the literature on uncertainty shocks begun by Bloom (2009) contends that uncertainty innovations generate “wait-and-see” recessions in which investments are temporarily delayed.

workers relative to the endogenous-separation threshold before an aggregate shock occurs.⁷ On the other hand, allowing for differential labor-demand shocks would require adjusting preferences to encompass durable goods or adding capital and investment to the model. All in all, focusing on the most important determinants of middle-skill employment movements would appear to be necessary if model tractability is to be preserved. Policymakers also need to understand how middle-skill employment unfolds in recessions. If middle-skill jobs are separating efficiently due to high outside alternatives and/or low option values, then there is little that monetary policymakers can (or should) do in response. But if middle-skill job losses arise instead from demand shifts in industries that are overly susceptible to monetary shocks and/or interest rates, then central bankers should not necessarily view middle-skill job losses as benign. Taking a longer view, if obtaining a higher-paying job is a not a realistic alternative for the typical middle-skill worker, then many middle-skill workers are likely to respond to long-run polarization by leaving the labor force entirely rather than by accepting lower-paying jobs.

Our empirical work is based primarily on data from the Current Population Survey (CPS) and can be divided into three parts. To gain some historical perspective for Figure 1, the next section investigates the postwar cyclicalities of employment and unemployment for workers in different skill classes. Because of regular changes in the government’s occupational classification system, this project requires the construction of a new dataset based on historical publications of the Bureau of Labor Statistics (BLS) and the Census Bureau.⁸ In section 3, the focus shifts to CPS microdata, available beginning in January 1976. The microdata allow us to disaggregate middle-skill employment and unemployment by industry, to measure job-finding and job-separation rates for separate industry/occupation cells, and to see where middle-skill workers go when they exit middle-skill employment. We therefore get some idea of these workers’ outside alternatives as well as the importance of industry effects for their employment movements. The microdata analysis suggests that transitions of unemployed workers to nonparticipation have increased during the past three decades, so section 4 investigates the relationship between nonparticipation and polarization among prime-age and older men. The regression model we use is based on one that Acemoglu and Autor (2011) use to study wages when workers with different skills can transition to jobs that perform tasks of higher or lower complexity, as they do in the polarization model.

Our bottom line is that industry effects, not outside alternatives, appear to be the most important factors driving high-frequency middle-skill employment fluctuations in re-

⁷Without idiosyncratic uncertainty, every middle-skill worker behaves in the same way when an aggregate productivity shock arrives; that is, if one worker separates, they all do.

⁸Our dataset is based on one used in Jaimovich and Siu (2013) that goes back to 1967. The theoretical section of that paper specifies an endogenous-separations model that links polarization to the recent phenomena of jobless recoveries. We discuss that model below.

cent years. The cyclical nature of employment for middle-skill production workers and operatives has been remarkably stable since the late 1940s, suggesting that changes in employment trends have had minor effects on the cyclical reallocation decisions of middle-skill workers. Employment for the “cognitive” middle-skill workers holding sales and clerical jobs has become more cyclical over time, but the same is true for high-skill workers, suggesting that this increase is driven by a factor specific to cognitive employment rather than to middle-skill occupations.⁹ The microdata show that within the middle-skill category, industry is a strong predictor of both job-flow patterns and occupational reallocation. Additionally, education appears to be a big hurdle for unemployed middle-skill workers hoping to move to high-skill jobs, as few such workers with only a high-school degree make this transition. Unemployed workers also appear reluctant to take low-skill jobs, though they would appear to satisfy the limited educational requirements of these jobs. In general, the limited ability for unemployed middle-skill workers to take high-skill jobs and their reluctance to take low-skill jobs suggests that nonparticipation is a highly relevant outside alternative for middle-skill workers. Though it is hard to follow individual workers through spells of nonparticipation, the task-based regression model estimated in section 4 shows that the demographic groups specializing in middle-skill work near the beginning of the microdata sample have a much higher probability of dropping out of the labor force by the end. Section 5 concludes the paper with a discussion of how our findings relate both to potential extensions of the search model and to the jobless recoveries the U.S. economy has recently experienced.

2 Historical Evidence on Occupational Employment and Unemployment

2.1 Constructing the historical dataset

Official BLS data on employment and unemployment by occupation are available beginning in January 1983 (for employment levels) and January 2000 (for unemployment rates). Earlier data are not published (or even available on machine-readable form) because the government’s occupational designations change about every 10 years, and BLS officials are not comfortable extending series corresponding to current occupations any further back in time. Yet it is possible to construct longer time series by using contemporary reports published by BLS and Census. Soon after the war ended, Census Bureau reports on the labor force included occupational employment levels for the first month of each quarter. Monthly data became available in January 1958, which we aggregate to the quarterly level.¹⁰ The published data

⁹The finding that high-skill employment has become more cyclical is consistent with Castro and Coen-Pirani (2008), who note an increase in the cyclical nature of employed hours among college-educated workers in CPS data.

¹⁰Because occupational employment data are available only for January, April, July, and October before January 1958, the quarterly dataset we construct is likely to have different seasonal properties before that

use the contemporary occupational classification systems, which must be aggregated into broader occupational categories in order to construct a consistent dataset. Table 1 shows the categories we use, which were originally suggested by Jaimovich and Siu (2013). For a polarization study, a key distinction among these categories is whether they encompass routine or nonroutine tasks. The *nonroutine cognitive* group is filled by high-skill workers doing managerial, technical, or professional tasks that are difficult to automate or offshore. *Routine cognitive* workers generally perform sales and clerical work. Like high-skill work, these jobs utilize brains more than brawn, but the tasks these workers perform are generally routine, so they are classified as middle-skill. The other middle-skill group are the *routine manual* workers who operate machines in factories, do construction work, or perform other types of production work. The final group consists of *nonroutine manual* workers, who engage in low-skill service tasks, such as food preparation and custodial work. While these tasks do not require much formal training, their nonroutine and location-specific nature makes them difficult to offshore or automate. The columns of Table 1 list the official major occupation groups defined by the government in selected decades. Individually, these groups undergo significant changes. A factory worker might be classified as an “operative or kindred worker” using the 1940 designations but a “production” worker using the 2010 designations. But the table shows that the shifting subgroups still can be aggregated consistently into four consistent categories over long periods. This is not to say that “seams” in the data never arise when these classification systems change. The two most important seams occur in 1971 and 1983 with the introduction of the 1970 and 1980 classification systems, respectively. The data appendix describes our method for dealing with these seams and our checks to ensure that the resulting data series are consistent over time.

Some useful background for the formal analysis of the four broad categories is gained by first consulting some disaggregated data from early in the sample. As noted earlier, the government’s occupational system was fairly stable from the end of World War II until the early 1970s. It turns out that the original publications we consult allow us to group occupational data into nine less aggregated categories from the end of the war (1945q3) through the final quarter of 1966. These data are plotted in Figure 2. The first two panels (for Proprietors, Managers, and Officials; and Professional and Semi-Professional workers) correspond to the high-skill category in Table 1 and display relatively modest business-cycle variation. The second panel shows that employment for professional workers in particular rises steadily in the immediate postwar era with the recessions of the 1950s and early 1960s having little apparent impact. More dramatic is the rapid rise of employment for proprietors and managers in the very early part of the sample displayed in the first panel. The postwar

time. We take account of that fact whenever we seasonally adjust the data. The availability of only one month per quarter in the earliest data also increases measurement error, but it is unclear how that error will bias comparisons of employment cyclicity across different occupational groups.

reallocation of labor was of intense interest to economic officials of the time, and a special Census report on postwar labor mobility noted that about 500,000 workers in the proprietors and managers group in August 1946 had been employed in some other occupational group 12 months earlier. By investigating the class-of-worker variable in the CPS, the officials determined that “a good many of the persons transferring to the proprietors, managers, and officials group were former wage and salary workers who went into business for themselves” (U.S. Census Bureau 1947, p. 2). Self-employment status is obviously relevant for employment cyclicity, a point we return to below.

The next two panels are for the Sales and Clerical workers that comprise the routine cognitive group. Cyclicity for these middle-skill workers is more pronounced than for high-skill workers, but routine cognitive employment varies much less over the cycle than it does for Craftsmen and Operatives, which are part of the routine manual group. These middle-skill workers experience much more significant employment declines in recessions than do the middle-skill workers in the routine cognitive class.¹¹ Finally, the last panel depicts employment for low-skill service workers, where employment cyclicity is low.

It would obviously be useful to measure the self-employment shares and industrial compositions of the four broad occupational groups we create by aggregating the data in Figure 2. Though a high-frequency decomposition is not possible, microdata from decadal Censuses gives us a sense of the relevant trends. Figure 3 depicts the share of employment held by self-employed workers and those in manufacturing and construction for each of the groups, beginning in 1910.¹² Not surprisingly, the self-employment share for the high-skill group is the largest, but this share falls rapidly until about 1970, at which point the decline becomes more gradual. The next panel depicts shares of manufacturing employment. In the routine manual category more than four in ten workers are employed in manufacturing until the 1990 Census. Though the share declines thereafter, manufacturing representation remains much higher for routine manual workers than for other group, even recently. Panel C depicts the share of construction employment, and here again the share for routine manual workers is highest. Additionally, a recent rise in the the construction share for routine manual workers has offset the decline in the manufacturing share, so that between 50 and 60 percent of routine manual workers have been employed in either manufacturing or construction for the past 100 years.

Figure 4 plots log employment levels for the four broad groups beginning in 1947q3, along with Hodrick–Prescott (HP) trends ($\lambda = 100,000$). The data are not seasonally adjusted,

¹¹Operatives in particular suffer a significant drop in employment during the recession of 1957–58.

¹²This analysis is possible because the IPUMS project has added an occupational code to the microdata that uses the 1950 classification system (OCC1950). This system can easily be aggregated into the four broad occupational groups. From 1910 through 2000, the data in Figure 3 are generated by decadal censuses. Beginning in 2001, the data come from the American Community Survey.

and we choose a starting date slightly after the end of the war so that immediate postwar reallocation has had a chance to sort itself out.¹³ The main takeaway from the figure is that routine manual employment has traditionally been far more cyclical than employment in the other groups, a fact we document more formally below. Routine manual employment also has a far more pronounced seasonal cycle.¹⁴ The steady increase in high-skill work chronicled in Goldin and Katz (2008) is also apparent in the figure. As the technological demands of industry grew, a progressively higher share of employment is accounted for by high-skill occupations; the total share of employment in high-skill occupations rises from about 20 percent in the late 1940s to about 40 percent today. Another takeaway of Figure 4 is that we must be careful using HP trends when analyzing these data. These trends are well known to have “endpoint problems,” that is, to be unduly influenced by movements at the end of the sample period. The big declines in employment for routine cognitive workers (Panel B) and routine manual workers (Panel C) during the Great Recession appear to have pulled down the HP trend in recent years, even with a relatively high smoothing parameter (100,000).

Starting in 1957, we are also able to obtain unemployment levels disaggregated by the last job that the unemployed person held, which remains the standard convention for the way that BLS measures unemployment by occupation. These data are available on a seasonally adjusted quarterly basis through 1981, and the data appendix explains how we use the CPS microdata to calculate unemployment rates for the four groups afterwards. The resulting unemployment rates appear in Figure 5. A main lesson from this graph is that while average levels differ, unemployment rates move together strongly over time. Perhaps this fact should not be too surprising, as unemployment rates published by BLS that are disaggregated in other ways (for example, by education) also reveal a common cyclical pattern.

There are, however, some important changes in relative unemployment rates over time. In the late 1950s and early 1960s, unemployment in the routine manual group was much higher than for other workers, a disparity that figured prominently in debates on the nature of unemployment at the time. Looking back on this period, Walter Heller, the chair of the Council of Economic Advisers under President Kennedy, wrote that many economists were skeptical of the proposed Kennedy tax cuts because they believed that most unemployment was structural:

The structural-unemployment thesis—the proposition that there had been a great *increase* in hard-core unemployment which would not yield to demand stimulus—found supporters from every point on the political spectrum. Automation sup-

¹³Note that movements in employment before January 1958 appear more irregular, probably due to the fact that we have only one month per quarter during that time.

¹⁴Regarding seasonality, the seasonal properties of the nonroutine manual workers in Panel D appear to change in 1983, when the 1980 occupation codes are introduced. Whenever we seasonally adjust the data, we therefore allow a different seasonal cycle to begin at that point.

posedly had us by the throat. There was a great mismatch between the skills demanded and the skills that could be supplied....Anyone who couldn't see that, we were told, was blind. The ravages of automation were obvious: displaced elevator operators, oil refinery workers, telephone operators, even assembly-line workers (Heller, 1966, p. 63, emphasis in original).

With the onset of the great boom of the 1960s, however, routine manual unemployment begins to decline, falling below the rate for the low-skill nonroutine manual workers by the middle of the decade. By and large, this has been the pattern in all business cycles ever since, including the booms of the 1980s and 1990s. Routine manual unemployment rises sharply in recessions, by an absolute amount that exceeds the increases for other groups. But declines in routine manual unemployment during recoveries are relatively large as well.

2.2 Basic cyclical patterns in employment and unemployment

We now assess the high-frequency relationship of both employment levels and unemployment rates with real GDP. To do so, we detrend all three series with HP filters that share the same smoothing parameter ($\lambda = 100,000$).¹⁵ We then construct rolling 10-year correlations of detrended GDP with detrended employment and detrended unemployment, respectively. Results appear in Figure 6. Looking first at the GDP–unemployment correlations, the common pattern of unemployment rates noted earlier suggests that all four unemployment rates should be negatively correlated with GDP, and the rolling correlations bear this out. For each of the four groups, the GDP–unemployment correlations are strongly negative and near -1 throughout the sample period. Yet the GDP–employment correlations vary widely. For middle-skill routine manual workers in Panel C, this correlation is uniformly large and positive, as suggested by the strong cyclical pattern for routine manual employment in Figure 4. But the GDP–employment correlation for the low-skill nonroutine manual workers in Panel D hovers near zero, indicating acyclical employment levels. Somewhere between these two extremes lies the nonroutine and routine cognitive workers in Panels A and B. For both groups, the GDP–employment correlations rise over time. At the beginning of the sample, detrended employment for both types of cognitive workers are somewhat negative, indicating that detrended employment is actually above trend in recessions. But over time, these correlations become positive, with this increase occurring earlier for the routine cognitive (middle-skill) employment category than for the nonroutine cognitive (high-skill).

Given the common correlations of unemployment with GDP, how can we explain the disparate correlations of GDP with employment levels? In a macroeconomic model with no population growth, one type of occupation, and no participation margin, the size of the labor

¹⁵Recall that this parameter is also used for the employment trends graphed in Figure 4.

force cannot change, so increases in unemployment must result in employment declines. Yet more complicated models could easily admit changes in unemployment that coincide with intensified movements of workers across occupations or changes in the flows in and out of the labor force, so the correlations of GDP with occupation-specific unemployment could differ. To some extent, the empirical work later in the paper investigates these additional margins. Unfortunately, the lack of microdata for the immediate postwar period, as well as limitations in the microdata available since 1976, prevent us from constructing a complete empirical model able to explain all the rolling correlations in Figure 6. We can, however, make some informed conjectures.

First, it is natural to suspect that the decline in self-employment noted earlier for high-skill workers is an important factor behind the rising cyclicity of high-skill employment: owners of firms rarely lay themselves off. The problem with this explanation is that the timing is not quite right. Recall from Figure 3 that the self-employment share for high-skill workers falls steeply throughout the 20th century, until about 1970, when the decline becomes more gradual. However, the rolling GDP–employment correlation for high-skill *declines* early in the sample even as the self-employment share is falling. The self-employment explanation also does not explain the rising correlation for sales and clerical workers in the routine cognitive category, where both the level and change in self-employment is minimal.¹⁶

We think a better explanation for the rising cyclicity of both types of cognitive employment involves a secular increase in the “cognitive input” required in the direct production of output over time. Think of two types of tasks that a high-skill worker could do: engage in overhead tasks like management and long-range planning, or alternatively perform tasks that directly produce output, such as writing a computer program or performing a medical operation. The long-run increase in demand for high-skill labor, which is chronicled by Goldin and Katz (2008) and which can reinforce the polarization of the labor market, has probably not occurred because the economy needs more managers. Rather, the skill requirements for tasks that directly produce output have grown.¹⁷ If so, then we would expect a rising percentage of high-skill workers to be involved in output production rather than in overhead labor. Going further, if we draw a distinction between overhead labor and direct production, then two questions become relevant when assessing the cyclicity of a group. First, how many of

¹⁶Castro and Coen-Pirani (2008) discuss the rising cyclicity of college-educated workers in general rather than high-skill occupations in particular. In their model, the decline in college-worker cyclicity is caused by a decrease of capital-skill complementarity occurring sometime in the mid- to late-1980s. They write that a calibrated model with this feature can explain about 60 percent of the increased relative volatility of college labor.

¹⁷Below, we reference a task-based model by Acemoglu and Autor (2011) that is useful for thinking about these issues. The model posits a continuum of tasks of increasing complexity and workers with different levels of skills needed to perform those tasks. In the model, high-skill workers specialize in the most complex tasks, so an increase in the complexity of the economy necessarily improves labor-market outcomes for high-skill workers.

these workers are directly involved in the production of output, making them vulnerable to a decline in aggregate demand. Second, how volatile is the demand for the type of output that these workers produce? Our conjecture is that routine manual workers have traditionally been deeply involved in the production of output—in the automobile industry, these workers are attaching fenders onto cars, not designing the cars. Additionally, the data show that routine manual workers tend to be employed in manufacturing and construction, where demand is relatively cyclical. Many workers in the other manual occupation, the low-skill nonroutine manual group, are probably also engaged directly in output-production tasks. But these tasks are linked more closely to consumption (restaurant meals, personal services, etc.), and consumption is less cyclical than investment. And of course many nonroutine manual workers are employed in overhead-type tasks that do not vary much over the cycle, such as custodial services.

Given that employment cyclicality differs across occupational groups, and occasionally over time, then why is the cyclicality of unemployment rates so uniform? Though we cannot decompose the unemployment rates into job-finding or job-separation rates, microdata after 1976 suggests that uniform movements in job-finding rates contribute significantly to the uniform movements in unemployment.¹⁸ Hall (2005) and others have shown that even in models with exogenous separation rates, small fluctuations in productivity can have big effects on job-finding rates when real wages are rigid.¹⁹ Consequently, with an accurate cyclical model for movements in and out of the labor force and for reallocation across occupations, it may be possible to embed an overhead/direct-production notion into a coherent explanation for the entire constellation of correlations displayed in Figure 6. While we are a long way from that explanation, we would also stress that our objective in this paper is more humble: to assess the relevance of outside options and industry composition for middle-skill workers in recessions. To some extent, that task is made easier by the stability of employment and unemployment correlations for the most cyclical group, routine manual workers.

2.3 A levels-based dynamic factor model of recent recessions and recoveries

The moving correlations paint a broad picture of cyclicality in the labor market, but a discussion of middle-skill employment during recent recessions requires a regression model.

¹⁸We will investigate job-finding and job-separation rates more formally below. Note that we are using “job-finding” as shorthand for the rate at which unemployed workers transition into employment. The job-finding rate and the UE rate differ when workers can enter jobs from outside the labor force. Similarly, what we call the job-separation rate is really flow from employment to unemployment (the EU rate).

¹⁹A seminal paper by Shimer (2005) shows that in a canonical search model with flexible, Nash-bargained real wages, productivity increases are essentially absorbed into higher wages. The wage increase offsets the effect of higher productivity on the incentives of firms to post additional vacancies, so the job finding rate does not rise much. If real wages are rigid, then higher productivity can cause a substantial increase in both vacancies and the job-finding rate and a substantial decrease in the unemployment rate.

In this section, we adopt the common strategy of estimating a dynamic factor model (DFM), in which a single business cycle impacts various subcomponents of the economy in potentially different ways. Our DFM is based on detrended log levels data and has the following form:

$$F_t = \rho F_{t-1} + \theta GDP_t + \nu_t \quad (1)$$

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

Here, F_t is a single unobserved common factor that depends on GDP and is constrained to be an AR(1) process with error term ν_t .²⁰ The four employment series E_{it} depend on the common business-cycle factor F_t , group-specific constants α_i , manufacturing output Mfg_t , real residential construction $ResConstr_t$, and an error term e_{it} .²¹ 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.²² To assess potential changes in the cyclicity of detrended employment over time, the DFM is estimated over three sample periods: an early sample that runs from 1947q3 to 1985q4, a later sample that extends from 1986q1 to 2013q4, and a third full sample that uses all available data.

Figure 7 shows how employment in the different occupations would be expected to change during the past three recessions and early recoveries. Each panel begins one year before the start of the relevant recession, which is denoted with a vertical line. The panels end four years after the recessions begin. The lines in the panels correspond to either actual data (black dashed line) or to dynamic predictions from one of the three models distinguished by their estimation periods. The dynamic predictions begin with the onset of each recession and are carried through the subsequent four years.²³ Turning first to routine manual employment panels in the third column, we see that the DFM generates very similar predictions no matter what sample period generates their coefficients. The lone minor exception is the dynamic prediction from the later model in the 2001 episode, which shows less of a drop than that experienced in actual data or predicted by the other models. By and large, however, the predictions for the routine manual group not only track one another closely, they track

²⁰The variance of ν_t is normalized to equal one.

²¹Manufacturing output comes from the Federal Reserve Board's industrial production release and residential construction comes from the national accounts. Both series are residuals from HP trends with $\lambda = 100,000$, as are the other variables in the model.

²²The individual errors ϵ_{it} are allowed to follow AR(1) processes. 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. 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).

²³The predictions are dynamic because they use no actual employment data once they begin, only actual real-side data such as GDP.

actual data well, too. Thus, the stability of the GDP–employment relationship we noted in the rolling correlations for routine manual workers carries over to more specific predictions for employment around recessions.

Other panels of the table also corroborate the correlational patterns. For high-skill workers in the first column, during the 1990–91 recession there is a significant miss for the predictions generated by the earlier data or the full sample, highlighting lack of much high-skill cyclicality in the earlier data. Estimating the DFM on the later model improves matters somewhat, but the miss remains large. The same pattern is displayed to a lesser extent during the 1990–91 recession for the routine cognitive workers in the second column. The earlier models predict more of an employment response for this group, consistent with the fact that cyclicality for routine cognitive workers began to rise sooner than it did for the high-skill group. Indeed, when the DFM is estimated on the later sample, it is able to explain the 1990–91 drop in routine cognitive employment fairly well. Finally, the fourth column shows that there is never much movement in the employment for the low-skill workers in the nonroutine manual group, either in the predictions of the DFM or in actual data.

2.4 A simpler model in differences

As discussed earlier, estimating a model with HP-detrended levels data has a potential drawback near the endpoints of the sample. As noted when discussing Figure 4, even with a strict smoothing parameter the HP filter suffers the common problem of being unduly influenced by endpoints. In light of this drawback, we estimate a model of occupational employment in log differences rather than log levels to evaluate behavior during and after the Great Recession. As before, we specify early, late, and full-sample periods over which to estimate the differenced models. Unlike the levels-based DFMs, however, we simply fold the GDP data into the observed equations, so that the model consists of four independent regressions rather than four observed equations linked by a common factor.²⁴ We also end the full and later samples in 2007q3 rather than 2013q4, so that the dynamic forecasts constructed at the onset of the Great Recession are true out-of-sample forecasts as well as dynamic forecasts. Finally, in order to generate comparisons with log levels data, the predicted log levels of employment are constructed by summing the dynamically forecasted log differences.²⁵

The results of the differenced model are presented in Figure 8. This figure is generally

²⁴The differenced model is therefore four independent regressions without a state equation.

²⁵Each regression projects the log difference of an employment series on the contemporaneous difference and two lags of log GDP, log manufacturing output, and log residential construction. Lags of the dependent variable are not included. We were concerned about seasonally adjusting the data before entering it into the differenced model, so we left the model in not-seasonally-adjusted form. To account for seasonality, we included seasonal dummies in the regressions, with interactions allowing for changes in the seasonal cycle in 1959 (when the quarterly data can be constructed by averaging monthly data) and in 1983 (because the use of the 1980 occupation codes appears to affect seasonality in the low-skill nonroutine manual group).

consistent with the main lessons of the previous DFM, but it shows that the HP-endpoint problem could be an issue for routine cognitive workers. Looking back at the levels-based DFM in Figure 7, the last panel in the second column indicates that all of the models do a passable job of explaining routine cognitive workers during and after the Great Recession. But of course the recession pulls the HP trend for these workers much lower. For the differenced model, which does not require a trend, Panel B shows that even a regression estimated on the 1986q-2007q3 sample misses the large, recent employment drop for routine cognitive workers. There are less substantial misses for the both high-skill nonroutine cognitive workers in Panel A and routine manual group in Panel C.

All told, results from both the levels and differenced models indicate that recent employment movements in the four occupational groups are close to what we would expect given movements in GDP, with the possible exception of routine cognitive workers during the most recent recession. The results therefore bear directly on whether recent recoveries should be considered “jobless.” Papers, including Galí, Smets, and Wouters (2012), have pointed out that once we take the behavior of GDP into account, lackluster employment growth coming out of the recessions should not be too surprising, so a better label for recent recoveries is “slow,” not jobless. The results from the two models estimated above indicate that this reasoning remains generally valid even when we consider employment movements for individual occupational groups rather than employment for the entire economy. More relevant for our purposes, though, is the finding that middle-skill job losses in particular have generally been in line with previous relationships, especially for routine manual workers. Although a lot of routine manual jobs were lost in the Great Recession, we should not necessarily view these losses as evidence that routine manual workers are more willing to separate from their jobs than they were in the past. To get a better sense of the outside alternatives available to middle-skill workers, we need to disaggregate occupational employment by industry and investigate the occupational alternatives available to middle-skill workers when they change jobs. We take up those topics next.

3 High-Frequency Movements of Middle-Skill Workers: Evidence from CPS Microdata

CPS microdata allow us to analyze worker flows over the business cycle and disaggregate employment by industry as well as occupation. The individual occupation codes in the microdata can also be used to construct more occupational classifications than are possible in the historical database on employment and unemployment. In this section, we make use of the high-, middle-, and low-skill occupational groups that underlie the 10 occupational groups in Figure 1. These 10 occupational classifications grew out of the standardization exercise conducted by Meyer and Osborne (2005), who created a classification system that generates

consistent occupational categories that can be applied to dataset such as the CPS microdata over many decades. While the Meyer–Osborne system has the advantage of consistency, it also has hundreds of entries—too many for a focused polarization study. Autor (2010) aggregates the Meyer–Osborne occupations into 10 coarser occupations, which appear in Figure 1. These occupations can be further grouped into the four broad categories we have used so far, defined by the cognitive/manual and routine/nonroutine disaggregation.²⁶

3.1 Employment growth by occupation and industry

Figure 9 depicts quarterly averages of monthly employment levels for the four broad groups broken down by industry. For each industry–occupation group, employment in the first quarter of 1976 is normalized to equal 100 so that differences in employment growth rates are easily seen. The first two panels, corresponding to manufacturing and construction, omit the low-skill service workers in the nonroutine manual group because very few low-skill workers are employed in those two sectors.

For manufacturing, Panel A reveals some early differences in both trend growth and employment cyclicalities for high- and middle-skill workers. High-skill manufacturing employment grew fairly steadily from the late 1970s to the late 1980s and was little affected by the two recessions of the early 1980s. The same is not true for middle-skill workers, either in the routine manual or routine cognitive groups. The 1982 recession in particular had a large effect on routine manual workers, though this category recovers partially during the ensuing recovery, as does the number of routine cognitive workers. Throughout the 1980s boom, however, employment growth in both of these categories trails that of high-skill workers in the manufacturing sector. These trends, of course, are what we would predict given the forces that drive polarization. Automation replaced the middle-skill workers performing routine tasks while high-skill workers were added, reflecting the complementary nature of high-skill work with capital.²⁷ Indeed, during the 1980s, the number of high-skill workers added to the manufacturing sector offset reductions in middle-skill workers nearly one for one, so that overall manufacturing employment fluctuated around a fairly stable level. Manufacturing output, by contrast, grew by nearly 40 percent from November 1982 to June 1990, reflecting the productivity gains that automation and other changes to factory processes engendered.

Manufacturing employment patterns changed with the onset of the 1990–91 recession. This downturn affected high-skill workers to a much greater degree than the recessions of

²⁶Specifically, all of the high-skill occupations (top row of Figure 1) go into the high-skill/nonroutine-cognitive class, the Sales and Office workers go into the middle-skill/routine-cognitive class, the Production Workers and Operators go into the middle-skill/routine-manual class, and all occupations in the bottom row go into the low-skill/nonroutine-manual class.

²⁷The complementarity between high-skill labor and capital throughout the economy is a key message of Goldin and Katz (2008).

the early 1980s did. High-skill manufacturing employment continued to decline during the weak early recovery from the 1990–91 recession, though it would eventually recover during the high-pressure labor market of the late 1990s. In the 2000s, however, high-skill employment in manufacturing began to trend lower, matching the negative trend in middle-skill manufacturing employment that had been going on for many years. At the same time, U.S. manufacturing output stagnated, growing by only 12 percent from December 2001 to January 2014.

One interpretation of the post-1990 changes in manufacturing employment patterns involves differences in the ways that automation and international trade affect middle- and high-skill labor in the factory sector. Papers such as Autor, Dorn, and Hanson (2013) and Pierce and Schott (2013) have documented the rapid increase in manufactured imports, particularly from China, that began close to the time that high-skill factory employment begins its downward trend in Figure 9. The increase in Chinese imports obviously reflects that country’s comparative advantage in routine labor, so we would expect a negative impact on routine manufacturing employment in the United States, similar to the effects of automation in the 1980s. But why might Chinese imports have a negative effect of high-skill employment in manufacturing? It may be that exploiting China’s comparative advantage in routine work requires some high-skill tasks to be relocated to China along with the middle-skill ones. For example, high-skill labor may be required to manage foreign workers, troubleshoot complex machinery, or to perform other hands-on tasks that cannot be done thousands of miles away in the United States. Alternatively, the effect of Chinese trade on high-skill U.S. manufacturing employment may reflect a measurement issue. Consider a group of high-skill workers doing design work inside a manufacturing plant in the 1980s. When production moves to China in the 1990s or 2000s, these workers may be reorganized into a different firm that is classified in a different industry (for example, inside business services rather than manufacturing). Whatever effect trade has on high-skill employment in U.S. manufacturing, its effect on middle-skill employment should be negative. And indeed the negative trend in routine employment in manufacturing appears to accelerate in the later part of the sample, when international trade becomes more important.

In construction, however, the relationship between high- and middle-skill employment trends is much different. Panel B shows that routine manual employment and high-skill employment grow nearly in lockstep throughout the sample period until the onset of the Great Recession in late 2007. Routine manual employment then collapses, with a significant but somewhat smaller decline in high-skill construction employment a short time later. Understanding both the *theoretical* forces behind polarization and the *empirical* classification of workers in BLS data helps explain this pattern. Workers swinging hammers and operating jackhammers on construction sites are often grouped in the middle-skill category in studies of

polarization, but this has more to do with the government’s occupational classifications than with these workers’ underlying vulnerability to polarization trends. Specific tasks performed on construction sites, though often requiring little education or skill, are much more difficult to automate or offshore than assembly-line work in factories. Yet Table 1 shows that the government did not include a major occupational group for “Construction and Extraction Occupations” until fairly recently. This means that while many construction workers are often labeled middle-skill—as are those in our historical analysis—they enjoy much better long-run employment prospects than middle-skill workers in the factory sector. Even so, because these construction workers are employed in a cyclical sector, their employment declines sharply in recessions, just as it does for middle-skill workers in manufacturing.

The final panel of Figure 9 displays data for industries other than manufacturing and construction. Two takeaways emerge. First, the steep and steady growth rate for high-skill workers in this panel, relative to the two middle-skill groups, illustrates that polarization does not solely arise from the relative decline of manufacturing, where many routine workers are employed.²⁸ Even outside of the goods-producing sectors, production processes are using high-skill labor more intensely. Consequently, while the absolute number of middle-skill workers is rising in Panel C, their long-run growth rates are much lower than that of high-skill workers. A second takeaway is that employment in goods producing sectors is relatively acyclical. Changes in the occupational classification systems complicate the analysis around the years 1983 and 2003, but an overall comparison of the two middle-skill groups in this panel with their counterparts in Panels A and B reveals more stable employment over the cycle.²⁹ One qualification is that sales and office workers in the routine cognitive group experience a very large employment decline in Panel C during the Great Recession. This decline, which had yet to reverse itself as of early 2014, is consistent with the differenced model of employment growth discussed in the previous section, which indicated that Great Recession employment losses for the routine cognitive workers were larger than the historical norm.

Regarding the general lessons for middle-skill workers, the main takeaway of Figure 9 is that industry designations matter for both trends and cyclical movements in employment. Consequently, in the remaining empirical work, we disaggregate middle-skill employment by industry (manufacturing, construction, and other) but group routine cognitive and routine manual workers together. This disaggregation is especially useful when considering movements of middle-skill workers through the unemployment pool, a task we take up next.

²⁸The handbook chapter of Acemoglu and Autor (2011) stresses this point.

²⁹A classification seam is apparent for the low-skill nonroutine manual group; employment jumps up discretely in 2003, when the 2000 codes becomes relevant.

3.2 Flows through unemployment

Figure 10 graphs the job-finding (U-to-E) rates and job-separation (E-to-U) rates for middle-skill workers disaggregated by industry. These rates are adjusted both for the changing demographic composition within groups over time and for time aggregation.³⁰ The top panel shows that job-finding rates are strongly synchronized across the three groups both in their average levels and in their cyclicalities. By contrast, the job-separation rates have very different means, with construction workers most likely to flow through the unemployment pool. Combined with the relatively high job-finding rate for middle-construction workers, this high separation rate indicates that employment inside the middle-construction category is much more fluid with respect to the unemployment margin than employment elsewhere. More relevant for our study of unemployment cyclicalities is the pattern in separation rates around recessions. The lower panel shows that while middle-other and middle-manufacturing workers have similar average separation rates, separations for the manufacturing workers “spikes” in recessions to a much greater degree than the middle-other group. A greater responsiveness of E-to-U movements is consistent with the higher degree of cyclicalities in middle-skill unemployment that we saw when investigating employment levels earlier.

What types of jobs are available for middle-skill workers moving through unemployment? The answer to this question is relevant for search theories with endogenous separation margins, because it elucidates workers’ outside options as they contemplate separating into unemployment during recessions. To provide some background for this analysis, the top panel of Figure 11 plots the destinations of middle-skill workers who move to new employers *without* experiencing a (measured) spell of unemployment in between. These so-called EE flows are available beginning in 1994, after the major redesign of the CPS, and are smoothed in the figure with a four-quarter trailing moving average. For all three types of middle-skill workers, the most common destination after an EE flow is to the same type of job that the worker held before. Even movements within overall middle-skill category are rare, though middle-manufacturing workers transition to middle-other jobs about 25 percent of the time while middle-construction workers do so somewhat less frequently. Also important is the fact that for all three types of middle-skill workers, transitions to high-skill jobs are uncommon. Middle-other workers are able to accomplish this feat more frequently than the other two groups, about as often as these workers transition to low-skill jobs.

With these patterns in mind, consider the UE flows plotted in Panel B. These data can be generated from 1976 on and are also smoothed with a four-quarter moving average. A key lesson from this panel is the same as one from the panel above: middle-skill workers usually

³⁰As pointed out originally by Shimer (2005) time-aggregation issues can arise because the CPS is a point-in-time survey and does not measure continuous flow rates. Thus, a change in the finding rate makes it more or less likely that a worker leaving a job in month t will still be unemployed when the survey for month $t + 1$ is taken.

transition back to the middle-skill jobs in the same sector they left. There are some interesting differences relative to the EE flows, however. First, for unemployed middle-manufacturing workers, the probability transitioning back to middle-manufacturing declines throughout the sample period. This negative trend is consistent with the fact that only in manufacturing was middle-skill employment declining in absolute terms, as seen earlier in Figure 9. Because the UE probabilities must sum to one, this negative trend must be offset in some way, and a rising probability of UE flows from middle-manufacturing to the middle-other group is largely responsible for doing so. Another lesson from the UE flows is that the probability that a middle-skill worker leaves unemployment by taking a high-skill job is lower than the corresponding EE probability. While the precise numbers are hard to make out in the graph, the share of middle-other flows to high via an EE flow averages 13.8 percent during the period of data availability (1994q1-2013q4). The share of middle-other UE flows ending in a high-skill job over the same period is 9.7 percent.³¹

3.3 A discrete-choice model of unemployment transitions

The probabilities plotted in Figure 11 show that a large majority of middle-skill reallocation is to other middle-skill jobs, and that the likelihood of getting a high-skill job is not necessarily enhanced by moving through the unemployment pool. As descriptions of the outside alternatives available to a middle-skill worker in recessions, however, the raw means in Figure 11 are incomplete. To start with, the probabilities take no account of worker characteristics such as age or education, which vary at low frequencies (for example, because of the aging of the labor force) and perhaps over the business cycle as well. Additionally, the UE probabilities are calculated *conditional* on exiting unemployment for another job. Yet many unemployment spells end in a transition to nonparticipation rather than employment. Finally, to learn about the reallocation decisions of middle-skill workers in recessions, it would be useful to measure middle-skill reallocation with reference to a specific cyclical indicator.

We therefore estimate multinomial logits for transitions out of unemployment in order to address some of these issues. 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 normalized as the baseline choice, the unconditional probabilities of

³¹For middle-manufacturing and middle-other workers, the probabilities of transitioning to a high-skill job in the 1994q1-2013q4 period are 9.4 percent via an EE flow and 5.3 percent via a UE flow. The corresponding probabilities for middle-construction workers are 6.9 and 4.1 percent, respectively.

transitioning from unemployment to either E or N are:

$$\Pr(E_{ij,t+1}|U_{ijt}) = \frac{\exp(\Gamma'_{iE}X_{ijt})}{1 + \exp(\Gamma'_{iE}X_{ijt}) + \exp(\Gamma'_{iN}X_{ijt})} \quad \text{and} \quad (3)$$

$$\Pr(N_{ij,t+1}|U_{ijt}) = \frac{\exp(\Gamma'_{iN}X_{ijt})}{1 + \exp(\Gamma'_{iE}X_{ijt}) + \exp(\Gamma'_{iN}X_{ijt})}, \quad (4)$$

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. We allow different types of employment exits in the actual model, though we group exits to middle-other, middle-manufacturing, and middle-construction employment into a single middle-skill category. Demographic characteristics of workers are included in the X vector, as the CPS data can generate controls for educational attainment, gender, marital status, and age. Because the CPS includes a question on how long the unemployed worker has been looking for a job, we also enter reported unemployment duration in the model.³²

To capture the effect of the business cycle, the X vector also includes an “average” job-finding rate for all workers in the economy. This rate is constructed by first calculating job-finding rates for five industry–skill groups: the three middle-skill groups plus high- and low-skill workers. These finding rates are then adjusted to hold constant the average demographic characteristics of workers in that group over time. They are also corrected for time-aggregation, which arises from the point-in-time nature of the CPS. The three middle-skill finding rates constructed in this way were previously depicted in the top panel of Figure 10, which showed a strong degree of comovement among them. Were high- and low-skill finding rates to be plotted in this panel as well, this comovement would survive, as the high- and low-skill finding rates would essentially lie on top of the finding rates for middle-other and middle-manufacturing workers. We therefore calculate the average economywide finding rate as essentially the first principal component of this system of five series.³³ This common factor allows us to ask how the destinations of unemployed middle-skill workers change when jobs become easier or more difficult to find.

Figure 12 displays our initial results. Each middle-skill group has its own logit model,

³²Age is specified as a cubic polynomial in the worker’s true age minus 35 years, so that all three age terms equal zero when the worker is 35 years old. Dummies are entered for nonwhite, female, and married. The three included education categories are less-than-high-school, some college, and college graduate, as high-school graduate is the omitted educational category. 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 (zero duration is omitted). We also exclude workers who are more than 70 years old from the estimation sample and include quarterly dummies (the first quarter is omitted).

³³More accurately, the common finding rate entered in the model is the common factor from a dynamic factor model with no covariates, other than quarterly dummies in the individual observation equations. This common factor is constrained to follow an AR(1) process. The factor generated by this model is essentially identical to what emerges from a simple principal components analysis.

which is estimated over two sample periods, 1976:q1–1985:q4 and 1986:q1–2013:q4. The figure depicts unemployment transition probabilities implied by these logits for a baseline worker who is 35 years old, male, unmarried, and white, and who has a high-school diploma but no further education. We also zero out the duration dummies, so the bars correspond to transition probabilities for unemployed workers with zero to one weeks of unemployment duration. The panels on the left side of the figure correspond to the early sample period while those on the right are from the later sample. Within each panel, the darker bars are the relevant probabilities when the common finding factor is one standard deviation above its mean, while the lighter bars correspond to probabilities when this factor is one standard deviation below its mean.

The first thing to notice about the probabilities in Figure 12 is that, as suggested by the raw UE flows, middle-skill workers tend to stay in middle-skill jobs. A middle-skill baseline worker with the specific set of covariates listed above has a very small probability of transitioning either to high- or low-skill employment. This lack of mobility holds true regardless of whether this worker previously held a middle-other, middle-manufacturing, or middle-construction job, and it remains true in both the early and later sample periods. The probabilities of transitioning to nonparticipation are higher than moving to a different skill class, but even these probabilities do not vary much over the cycle. An increase in the overall job finding rate raises the probability that the baseline middle-skill worker who has just become unemployed will return to middle-skill work in the next month. But because no other transitions are significantly affected, the only other effect of a tighter labor market is to reduce the probability that the worker remains unemployed.

The acyclical probabilities of moving to a high-skill job in Figure 12 are even lower than the raw UE means plotted in the lower panel of Figure 11. The reason for discrepancy is that the baseline worker generating the probabilities from the logits is assumed to have a high-school education, and high-school workers rarely transition to high-skill employment. To show this, Figure 13 calculates the same probabilities after “turning on” the college dummy. The probability of exiting unemployment for high-skill work rises significantly for all three groups, even though all the unemployed workers in the estimation sample previously held middle-skill jobs. The other qualitative features of the probabilities remain the same, however; in particular the lack of much cyclical movement in transitions to nonparticipation.

As noted above, the multinomial logits used to construct Figure 12 are estimated with duration dummies as covariates, and the probabilities pictured there assumed that duration equals one week or less. Figure 14 illustrates how transition probabilities for the baseline high-skill workers change as unemployment duration rises, assuming that the overall finding-rate factor is fixed at its mean. In each panel, the solid lines depict probabilities estimated over the early sample (1976:q1–1985:q4) and the dashed lines correspond to the later sample

(1986:q1–2013:q4). Each panel in this figure replicates the common observation that job-finding probability for unemployed workers displays negative duration dependence: the longer an unemployed worker has been looking for a job, the smaller the probability that he will leave unemployment in the next month. But the panels also show something that is less well known—this negative duration dependence exists only for reemployment to the worker’s most recent skill class. For middle-skill workers with high-school educations, this lack of duration dependence may not be surprising, given the low probabilities that these workers ever transition to jobs in different skill categories. But we find it in unreported work on unemployed high- and low-skill workers as well. The finding is consistent with recent work on unemployment and recalls by Fujita and Moscarini (2013), who use data from the Survey of Income and Program Participation (SIPP), which can reveal whether an unemployed worker is eventually recalled to the same employer. The authors find not only that recalls are common, but that negative duration dependence exists only for recalls to the same firm.³⁴ Consequently, the same-skill duration dependence depicted in Figure 14 may well be generated by recalls to the unemployed worker’s previous employer.

Another interesting result in Figure 14 concerns the probability of exiting for nonparticipation. We might expect positive duration dependence for nonparticipation exits. The longer a worker remains unemployed, the more discouraged that worker is likely to get, and the higher the probability that job search would be abandoned, at least for a time. Yet there is no evidence of positive duration dependence in Figure 14, a pattern that Fujita and Moscarini (2013) also find in the SIPP data. A comparison of nonparticipation probabilities across sample periods, however, shows that in the more recent sample period, unemployed middle-skill workers who have been out of work for some time are more likely to give up searching than workers with similar durations in the 1976–1985 period. In other words, though exits to nonparticipation at zero duration never seem sensitive to the business cycle, nonparticipation exits are more likely at positive duration in the later sample period, regardless of the stage of the cycle.

Putting the pieces together, the results of this section suggest a consistent picture of middle-skill opportunities. The most important finding is that middle-skill workers appear to have limited options (or limited desires) to transition to jobs in different skill classes, unless they have a college degree and can move to a high-skill job. Unemployed middle-manufacturing workers facing absolute declines in employment do move more often to middle-skill work outside of the goods-producing sector, suggesting some overlap in the skill requirements in those two middle-skill areas. But overall, employment options for middle-skill workers with only a high-school diploma appear to be significantly constrained. In light of

³⁴Specifically, SIPP data show that more than one quarter of all workers who separate from their jobs return to work for their previous employer. When one considers only those separating workers who remain unemployed (that is, the workers who do not leave the labor force) the returning fraction rises to 40 percent.

this finding, the patterns relating to nonparticipation exits in the latter part of the sample take on additional importance. Do increasing numbers of middle-skill workers now realize that they need more education to compete in today’s job market, and therefore choose to leave the labor market temporarily to obtain those skills? Or do these workers simply give up and choose not to work at all? We take up this question in the next section.

4 Declining Middle-Skill Job Opportunities and Male Nonparticipation

While the nonparticipation results in the previous section are intriguing, in many ways the CPS is not well-suited for a study of nonparticipation. The concepts of “nonparticipation” and “unemployment” are distinct to economists but not necessarily to the public, as many CPS respondents who do not have jobs cycle through these two states in matched CPS data (Elsby, Hobijn, and Şahin 2013). Due to the short-panel nature of the CPS, it is difficult to learn whether someone exiting to nonparticipation has left the labor force for good. And studying the relationship between nonparticipation and occupational polarization is especially difficult since the CPS does not ask respondents about their previous occupations when they are classified as nonparticipants. When unemployed respondents exit for nonparticipation in our logits, it is hard to know what happens to them. And we cannot study the past occupations of all nonparticipants in the CPS, because their occupational data is not available if they never worked or searched for jobs while in the sample.³⁵

These empirical complications are related to deeper theoretical issues that arise when taking a “task-based” model to the data. This type of model is set forth in the handbook chapter of Acemoglu and Autor (2011) as an alternative to the canonical model of the labor market, in which certain amounts of high- and low-skill labor combine with capital to produce output.³⁶ The model of Acemoglu and Autor (henceforth AA) essentially places a layer of tasks between skills and output. Workers apply their (exogenous) skills to tasks of varying complexity, and the completion of these tasks results in the production of output. A set of high-skill workers is better able to perform very complex tasks than workers with intermediate or low amounts of skill. Given some regularity conditions, comparative-advantage arguments imply that the set of tasks of continuously varying complexity will be allocated monotonically to workers with discretely varying amounts of skill. High-skill workers will perform the most complex tasks and earn the highest wages. Middle-skill workers will perform the tasks of

³⁵There are also potential issues related to the CPS redesign in 1994. As noted earlier, our logits enter a post-1994 dummy to account for the redesign and this dummy is turned on when calculating the later-sample transition probabilities. Turning the dummy off reduces (but does not eliminate) the gap in nonparticipation-exit probabilities in Figure 14, and the effect is largest for middle-other and middle-manufacturing workers. Additionally, in a model estimated on high-skill workers, we also find a moderate increase in these exit probabilities when moving from the early to the later sample.

³⁶For a direct comparison of the canonical model with the task-based model, see Acemoglu and Autor (2012).

intermediate complexity and earn less, while the low-skill workers perform the simplest tasks and earn the least.

The economy's allocation of skills to tasks can change in at least two ways. One of these ways, technical change, arises when the abilities of one skill class rise relative to the abilities of the other groups. For example, an increase in the abilities of college-educated workers may result in them undertaking tasks previously allocated to middle-skill workers. One example is a college-educated worker, aided by a high complementary with computers, doing an office administrative job much more efficiently than a high-school worker who previously performed it. Another way in which the skill-task allocation can change is through automation or trade, which AA argue gives rise to polarization. Specifically, AA posit that tasks in the middle of the task-complexity distribution are more routine than the tasks at the ends, so the workers performing the intermediate tasks are better candidates to be displaced by automation or by foreign workers.³⁷ When this displacement occurs, middle-skill workers must apply their skills to new tasks elsewhere in the task-complexity distribution.

Where do these displaced middle-skill workers go? AA explain that it is theoretically possible for middle-skill workers to start performing some tasks that previously had been allocated to high-skill workers. Because the displacement of middle-skill workers lowers middle-skill wages, the displaced workers may be able to “underbid” high-skill workers for the least-complex tasks they had previously performed. However, AA show that with some realistic assumptions about structure of comparative advantage across skills, most of the middle-skill workers will “slide downhill” and start doing low-skill jobs. Displaced autoworkers may be willing to perform surgeries or teach physics classes at bargain rates, but their inadequate ability to perform those tasks, relative to the abilities of high-skill workers, will keep them from being hired for them. The displaced autoworkers will wind up mowing lawns and doing custodial work instead.³⁸

AA point out that shifts in the allocation of tasks complicates the study of long-run forces like technical change and polarization. Consider the college graduates who use computers to perform office administrative tasks more productively than the high-school workers who performed them in the past. The wages now paid to office administrators may rise thanks to the higher production of the new workers, but that does not mean that labor market outcomes for the workers who had traditionally performed administrative jobs have improved. A long-run wage equation with a simple “office-administration” dummy on the right-hand-side would therefore be misleading. A similar problem arises when polarization occurs. Polar-

³⁷Automation is easy to introduce into the model via an additional productive factor, capital, which has a strong comparative advantage for performing routine tasks.

³⁸This “middle-to-low” pattern of reallocation is hard-coded into the empirical model of Autor and Dorn (2013), which uses a task-based model to investigate the spatial pattern of occupational reallocation across U.S. commuting zones since 1980.

ization worsens the labor market outcomes of the mid-level workers it displaces, but once these workers take new jobs, they are classified as “low-skill” workers in empirical data, and information about their original comparative advantage for tasks in the middle of the skill distribution is easily lost. AA’s task-based model does not have a participation margin, but it is easy to see how this complication matters for nonparticipation. Thanks to automation, an autoworker may be displaced and take a job as a janitor. He gradually realizes that he does not like to work at night for low wages, so he drops out of the labor force for good. Without having followed that worker from his auto-industry job to the janitorial job, and then to nonparticipation, we would not realize that polarization was the true source of the labor-force exit.

4.1 A potential regression specification

The problems in taking the task-based model to the data could be solved if we observed each worker’s innate level of skill, and thus his initial underlying comparative advantage in the task-complexity distribution, rather than just the task or occupation he most recently performed. In their empirical work, AA assume that this underlying comparative advantage can be proxied by the task-distribution of that worker’s demographic group in some base year. For example, if most males in Michigan aged 25–30 with high-school diplomas worked in middle-skill jobs in the base year, we would assume that the comparative advantage for this group throughout the sample is in middle-skill tasks. If we later found that men of the same age (25–30) with the same education and living in the same part of the country were paid low wages or had higher rates of nonparticipation, we could lay some of the blame on polarization. This assignment would not require us to follow workers from this group into lower paying jobs. Indeed, the task-based model implies that once we know the “true” comparative advantage of a worker, his specific occupational assignment is no longer empirically useful.

In their paper AA study wage changes for demographic groups using decadal Census data and choosing 1959 as a base year. For our study of nonparticipation we use quarterly data for males in the CPS. Let P_i^H be the baseline share of high-skill employment for demographic group i , where the baseline period is defined to run from the first quarter of 1976, when microlevel CPS data become available, through 1981:Q4. Following Acemoglu and Autor (2011), we demarcate our demographic groups on the basis of gender, education, age, and geographic area. We differ from them in that we include only men in the sample, we use five-year age groups rather than 10-year groups, and use Census division as the geographic classification rather than Census region.³⁹ Finally, we estimate separate equations for prime-age men (ages 25–54) and older men (55+). The shares for middle- and low-skill employment,

³⁹Both our estimation samples and theirs excludes workers younger than 25 years old. Their wage equation excludes workers older than 64 years old; our participation equation includes these older workers.

P_i^M and P_i^L , are defined analogously. By construction $P_i^H + P_i^M + P_i^L = 1$.

A possible regression specification using these data is:

$$N_{it} = \beta_t^H \cdot (P_i^H \times \phi_t) + \beta_t^M \cdot (P_i^M \times \phi_t) + \beta_t^L \cdot (P_i^L \times \phi_t) + \gamma^e \cdot \phi_e + \gamma^a \cdot \phi^a + \gamma^g \cdot \phi_g + e_{it}, \quad (5)$$

where N_{it} is the nonparticipation share for demographic group i in time quarter t ; ϕ_e , ϕ_a , and ϕ_g are dummies for education, age, and geographic groups; and ϕ_t are time dummies. In the top row, the interactions of the time dummies ϕ_t with the baseline occupational probabilities (the P s) mean that the occupation-specific β s will trace out the combined effects of skill-specific productivity shifts, automation, and other developments on workers with different baseline comparative advantages. Estimating the sample over 1982:Q1 through 2013:Q3, we expect the sequence of β_t^M s to be increasing for groups with large P^M s: That is, the demographic groups with comparative advantage in mid-level tasks are likely to experience falling participation rates over time as the demand for mid-level tasks declines.

4.2 Long differences

The regression specification above does have an important disadvantage, however. Given the supply of skills and demand for tasks, it assumes that the baseline period reflects self-selection into tasks based on comparative advantage. Consequently, *any* change over time in a group's participation rate is a function of initial occupational shares.⁴⁰ To explore whether a correlation between occupational shares and participation behavior is reasonable, we perform some long-difference analysis. For each demographic group, we calculate the average rate of participation from 1976:Q1 to 1979:Q3 and from 2010:Q1 to 2013:Q3. We then correlate the long difference between these averages with the initial group-specific occupational shares as a check on the logic behind our model.

The regressions in table 2 are cross-sectional since, as before, there is only one long difference per demographic group, and they also include dummies that capture the main effects of age, education, and geography.⁴¹ Column 1 of Panel A includes the baseline high-skill share along with the main demographic controls for the sample of prime-age men. The coefficient on high-skill share is positive (27.96) and strongly significant. The correlation says the a 10-percentage-point increase in the initial high-skill share is associated with an increase in the high-skill participation rate of nearly 2.8 percentage points. The correlation in Column 2 says that a 10-percentage-point increase in the initial middle-skill share is associated with a decrease in middle-skill participation of 2.6 percentage points. In the third column, the

⁴⁰While the regression does include the main effects of education, age, and Census division, these enter as simple level effects and are not interacted with any time-varying variable.

⁴¹All regressions in table 2 are weighted by population.

coefficient on low-skill share is small and insignificant.

The last column enters all the baseline skill shares at the same time and drops the constant, which is required since the shares sum to one.⁴² Essentially, this column runs a horse race to determine the type of employment (high, middle, or low) that are most closely correlated with changes in participation over long periods of time. The middle-skill share emerges as the most robust predictor. The implication is that the high-skill share is positively correlated with participation changes in the earlier analysis because the high-skill share is negatively correlated with the middle-skill share (as the shares must sum to one). Panel B of table 2 repeats the analysis for older men, ages 55 and older, and the same pattern emerges. We see a significant decline in labor force participation for older men specialized in mid-level tasks, just as we did for prime-age men.

4.3 Panel regressions

The analysis of the previous subsection establishes that participation changes are correlated with baseline skill shares over a long time frame. We now relate high-frequency movements in participation rates to these shares in a panel setting. The regression specification has sufficient flexibility so that participation rates can change in ways that do not necessarily depend on skill shares. To do so, we replace the high-skill interactions with the time dummies in the original specification with a quadratic trend, giving

$$N_{it} = \varphi_1 \text{Trend}_t + \varphi_2 \text{Trend}_t^2 + \beta_t^M \cdot (P_i^M \times \phi_t) + \beta_t^L \cdot (P_i^L \times \phi_t) + \gamma^e \cdot \phi_e + \gamma^a \cdot \phi^a + \gamma^g \cdot \phi_g + e_{it}. \quad (6)$$

Omitting the high-skill interactions to make room for the quadratic trend is informed by the results of the previous section, which indicated that the effect of baseline high-skill shares on participation changes is small.⁴³

Estimating the panel regressions allows us to construct counterfactual labor force participation rates to see how the trends in coefficients matter for the overall participation rates, which are presented in Figure 15. The heavy black line in each panel is the participation rate reported by the BLS for either prime-age men (Panel A) or older men (Panel B). The red line in each panel fixes the influence of baseline middle-skill shares to be constant throughout the sample, with this constant equal to the average of the β_t^M interactions from 1982:Q1 through 1985:Q4. That is, given the labor market’s ability to allocate skills to tasks, the

⁴²The loss of a constant term means that the skill-share coefficients in the last column of the table are not strictly comparable to estimates from the previous three regressions.

⁴³We also performed a robustness check in which we left the high-skill interactions in the panel regression and dropped the trend. Our results regarding the effects of baseline middle- and low-skill shares on participation changes do not change when this is done. All the regressions, including the one presented in the paper, include quarterly dummies.

red line measures what the participation rate would have been if the demand for middle-skill tasks had not declined. The blue line is constructed in a similar way by holding the effect of low-skill participation opportunities constant at its 1982–1985 average. Finally, the green line is the prediction for the participation rate when both the baseline middle-skill and low-skill job opportunities are fixed.

The striking message of Panel A is that falling demand for middle-skill tasks appears to completely explain the much-described decline in the prime-age male participation rate. The red line in this graph has no downward trend, indicating that participation would have remained stable had the coefficients on the middle-skill interactions not changed. The green line holds both the low- and middle-skill opportunities constant; again, the implied participation rate does not decline much.

The regression includes a quadratic trend, so if participation was declining among all demographic groups, without regard to the groups' initial middle-skill shares, the regression estimates could have reflected a secular decline in participation with large coefficients on the trend and small coefficients on the baseline skill interactions. What actually happens, though, is that participation falls the most among the groups with large baseline middle-skill shares, so the coefficients on the middle-skill interactions become increasingly large in absolute value over time. The coefficients on the trend terms are smaller, indicated by the modest bend in the green line when both the middle- and low-skill coefficients are fixed.

The lower panel of figure 15 performs the analogous experiment for older men. Declining job opportunities for middle-skill workers had a quantitatively significant effect on male participation rates. A comparison of the red line, which holds middle-skill opportunities constant, and the black line, which depicts the published participation rate, suggests that the participation rate of older males would be about nine percentage points higher if middle-skill opportunities did not change over time. This is an even larger effect than was registered for prime-age males.

There appears to be some fallout from the Great Recession. Older workers with a comparative advantage in middle-skill tasks may have chosen to leave the labor force more quickly than they otherwise would have. The question is whether these older workers did so in order to spend time going to school or otherwise increase their levels of human capital, as the recessions-as-reallocations theory would suggest. A simpler story is that these workers ended their participation in the labor force early because the returns to continued participation had fallen.

All told, declining middle-skill opportunities appear to have had significant effects on male participation rates over the past several decades, consistent with declining relative wages for workers that have a comparative advantage in middle-skill work. But the movements in participation that are driven by polarization appear to work out slowly over time. Higher-

frequency movements in participation appear confined to low-skill prime-age men during the housing boom and older men near the Great Recession. But these movements are more likely driven by changes in the return to participation at any point in time and not by workers taking advantage of recessions to increase their skills.⁴⁴

5 Conclusions and Implications for Jobless Recoveries

This paper began by highlighting the large employment losses among middle-skill workers during the Great Recession. A relevant question for students of the business cycle is the degree to which these losses reflect the effect of temporary productivity fluctuations on the endogenous and efficient separations in matches with poor long-run outlooks. How does the empirical work presented above shed light on this question? First, the analysis of historical occupational data indicates that middle-skill workers typically experience large employment losses in recessions due to the industries that employ them, and recent middle-skill movements are for the most part consistent with past cyclical behavior. Industry also has a direct effect on middle-skill job flows, as illustrated in CPS microdata. Consequently, recent middle-skill job losses should not be interpreted as *prima facie* evidence that efficient separations from middle-skill jobs were accelerated by long-run trends away from middle-skill work, unless this type of acceleration has accompanied most postwar recessions in the United States.

Second, analyzing the reallocation of unemployed middle-skill workers since 1976 suggests that many face a significant barrier in getting a better job: education. Unemployed middle-skill workers with a high-school degree are much more likely to transition to nonparticipation than they are to high-skill work. Additionally, unemployed middle-skill workers appear reluctant to take low-skill service jobs, though they presumably fulfill the limited educational requirements of these jobs. Studying month-to-month movements though nonparticipation in CPS is fraught with pitfalls, but the apparent inability of less-educated middle-skill workers to take high-skill jobs, combined with an unwillingness to take low-skill jobs, suggests that long-run polarization trends should negatively affect labor force participation among workers with mid-level skills. We provide some indirect evidence involving changes in participation rates among male demographic groups to support this view.

These results imply that much of what we need to know about middle-skill employment reallocation can be understood with simple comparisons of education levels and real wages for different industry-skill groups, which is done in Figure 16. The chart on the left shows that since 1990, about 85 percent of high-skill workers had taken at least some college courses.⁴⁵

⁴⁴Regarding participation movements around the housing boom, Charles, Hurst, and Notowidigdo (2012) suggest that construction opportunities during the housing boom drew noncollege men into employment, “masking” the effect of the ongoing decline in opportunities in manufacturing employment.

⁴⁵That is, using the standard classifications, most high-skill workers were listed as either “some college” or “college graduates.”

This fraction falls to about half for middle-other workers and about one-third for middle-manufacturing and middle-construction workers. The educational makeup of these latter two groups is similar to those holding low-skill jobs; the two middle-skill groups have about the same number of less-than-high-school workers but somewhat more high-school graduates. The panel on the right of Figure 16 graphs average inflation-adjusted wages for the same five groups, using the data from the outgoing rotation groups of the CPS.⁴⁶ The three middle-skill groups earn significantly more than the low-skill workers but significantly less than the high-skill group.⁴⁷ With these patterns in mind, consider a middle-skill worker whose job is threatened by automation or trade. Moving up to a high-skill job is likely to take significantly more education, but moving down to a low-skill job is likely to entail a big cut in pay. If getting more education is infeasible and working for much lower wages is not attractive, then this worker will be more likely to drop out of the labor force over time. The educational requirements needed to obtain high-skill work are directly relevant to models of middle-skill reallocation, because they should figure into the decisions of middle-skill workers contemplating job separations.

The educational requirement also bears on some recent work that connects the polarization of the labor market to the joblessness of recent recoveries. Using a model of efficient separations, Jaimovich and Siu (2013) argue that large numbers of middle-skill workers, aware of long-run polarization trends, have efficiently separated during recent recessions in hopes of obtaining a high-skill job. These workers cannot simply search in the high-skill market, however. The separating middle-skill workers must enter a so-called switching market, which operates much like any other market in a Diamond-Mortensen-Pissarides (DMP) world. Vacancies are posted in this market and matches are made, and soon after a switching match is formed the worker obtains the skills necessary to hold a high-skill job. To reflect the long stretches of joblessness that middle-skill workers have to endure, Jaimovich and Siu (2013) posit that the exogenous vacancy-creation costs in the switching market are high. These high costs reduce the number of vacancies posted and consequently the number of switching matches, thereby forcing middle-skill workers who have left their jobs to wait a long time before becoming skilled—hence the joblessness of the recoveries. The results of this paper suggest two questions for a model of this type. First, on a qualitative level, how well is the real-world *education sector* approximated by a standard DMP job market with high vacancy-posting costs? In many cases, people can go to school part-time, so they would not need to quit middle-skill positions in order to improve their labor market outcomes. If so, then the

⁴⁶Elsby, Shin, and Solon (2013) use the same data to construct a similar plot that is not disaggregated by occupation. We discuss this paper below.

⁴⁷The wages in this panel are deflated by the personal consumption deflator. We have also performed a demographic adjustment so that changes in real wages for a specific group are not affected by changes in the demographic makeup of that group.

motive for endogenous separations in recessions is reduced. Second, the smaller the likelihood of obtaining a switching match—that is, the more jobless the recovery—the less attractive an initial separation from a middle-skill job will be. In order to generate large numbers of endogenous separations with long switching times, the model would appear to need a large gap in high-skill vs. middle-skill wages or large recessionary declines in aggregate productivity. While there may be other ways for the model to generate a large number of separations, the fundamental tension between the joblessness of the recovery and severity of the recession makes a precise calibration of the model necessary. We would argue that a better way to think of recessionary middle-skill employment movements is to first recognize the cyclicity of the industries that employ middle-skill workers. Combined with rigid real wages, this high cyclicity may generate large amounts of not-necessarily-efficient separations that do not depend on search motives.

Wage rigidity plays a role in other theories of jobless recoveries (Shimer 2012; Galí, Smets, and Wouters 2012). In these models, recoveries are slowed because real wages are stuck too high, due to either exogenous rigidities or changes in wage markups. A look at Figure 16 suggests that wages appear more rigid following the Great Recession than they did during the severe recessions of the early 1980s. Noting this pattern in aggregate data, Elsby, Shin, and Solon (2013) posit that the low rate of nominal price inflation after the Great Recession might have kept real wages artificially high, due to the difficulty of making nominal wages cuts. But they also note that wage changes in the microdata do not provide much support for that view. Our finding of a significant nonparticipation response to middle-skill employment opportunities could be relevant here, because rising nonparticipation reduces the need to cut real wages in order to clear the labor market. While a rigorous examination of middle-skill wage cyclicity is beyond the scope of this paper, it would appear to be an interesting topic for future research.

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.
- Acemoglu, Daron, and David Autor. 2012. “What Does Human Capital Do? A Review of Goldin and Katz’s *The Race between Education and Technology*.” *Journal of Economic Literature* 50(2): 426–463.
- 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. 2013. “The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market.” *American Economic Review* 103(5): 1553–1597.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103(6): 2121–68.
- 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.
- Barsky, Robert B., Christopher L. House, and Miles S. Kimball. 2007. “Sticky-Price Models and Durable Goods.” *American Economic Review* 97(3): 984–998.
- Blanchard, Olivier J., and Jordi Galí. 2010. “Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment.” *American Economic Journal: Macroeconomics* 2(2): 1–30.
- Bloom, Nicholas. 2009. “The Impact of Uncertainty Shocks.” *Econometrica* 77(3): 623–685.
- 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.
- 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.

- Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin. 2013. “On the Importance of the Participation Margin for Labor Market Fluctuations.” Federal Reserve Bank of San Francisco Working Paper No. 2013-05. Available at <http://www.frbsf.org/publications/economics/papers/2013/wp2013-05.pdf>.
- Elsby, Michael W.L., Donggyun Shin, and Gary Solon. 2013. “Wage Adjustment in the Great Recession.” National Bureau of Economic Research Working Paper No. 19478. Available at <http://www.nber.org/papers/w19478>.
- Fisher, Jonas D. M. 2006. “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks.” *Journal of Political Economy* 114(3): 413–451.
- Fujita, Shigeru, and Giuseppe Moscarini. 2013. “Recall and Unemployment.” National Bureau of Economic Research Working Paper No. 19640. Available at <http://www.nber.org/papers/w19640>.
- Galí, Jordi, Frank Smets, and Rafael Wouters. 2012. “Slow Recoveries: A Structural Interpretation.” *Journal of Money, Credit and Banking* 44: 9–30.
- Goldin, Claudia D., and Lawrence F. Katz. 2008. *The Race Between Education and Technology*. Cambridge, Mass.: Belknap Press of Harvard University Press.
- Hall, Robert E. 2005. “Employment Fluctuations with Equilibrium Wage Stickiness.” *American Economic Review* 95(1): 50–65.
- Jaimovich, Nir, and Henry E. Siu. 2013. “The Trend is the Cycle: Job Polarization and Jobless Recoveries.” Duke University Working Paper. Available at <https://sites.google.com/site/nirjaimovich/papers/JPJR.pdf>.
- 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>.
- Mortensen, Dale T., and Christopher A. Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies* 61(3): 397–415.
- Pierce, Justin R., and Peter K. Schott. 2013. “The Surprisingly Swift Decline of U.S. Manufacturing Employment.” U.S. Census Bureau Center for Economic Studies Working Paper 13-59. Available at <ftp://ftp2.census.gov/ces/wp/2013/CES-WP-13-59.pdf>.
- Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*. Second edition. Cambridge, Mass.: MIT Press.

- Pissarides, Christopher A. 2009. “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?” *Econometrica* 77(5): 1339–1369.
- Shimer, Robert. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies.” *American Economic Review* 95(1): 25–49.
- Shimer, Robert. 2012. “Wage Rigidities and Jobless Recoveries.” *Journal of Monetary Economics* 59, Supplement(0): S65 – S77.
- 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.
- U.S. Census Bureau. 1947. *Current Population Reports: Labor Force, Series P-50, No. 1. Industrial and Occupational Shifts of Employed Workers: August, 1945, to August, 1946.* April.

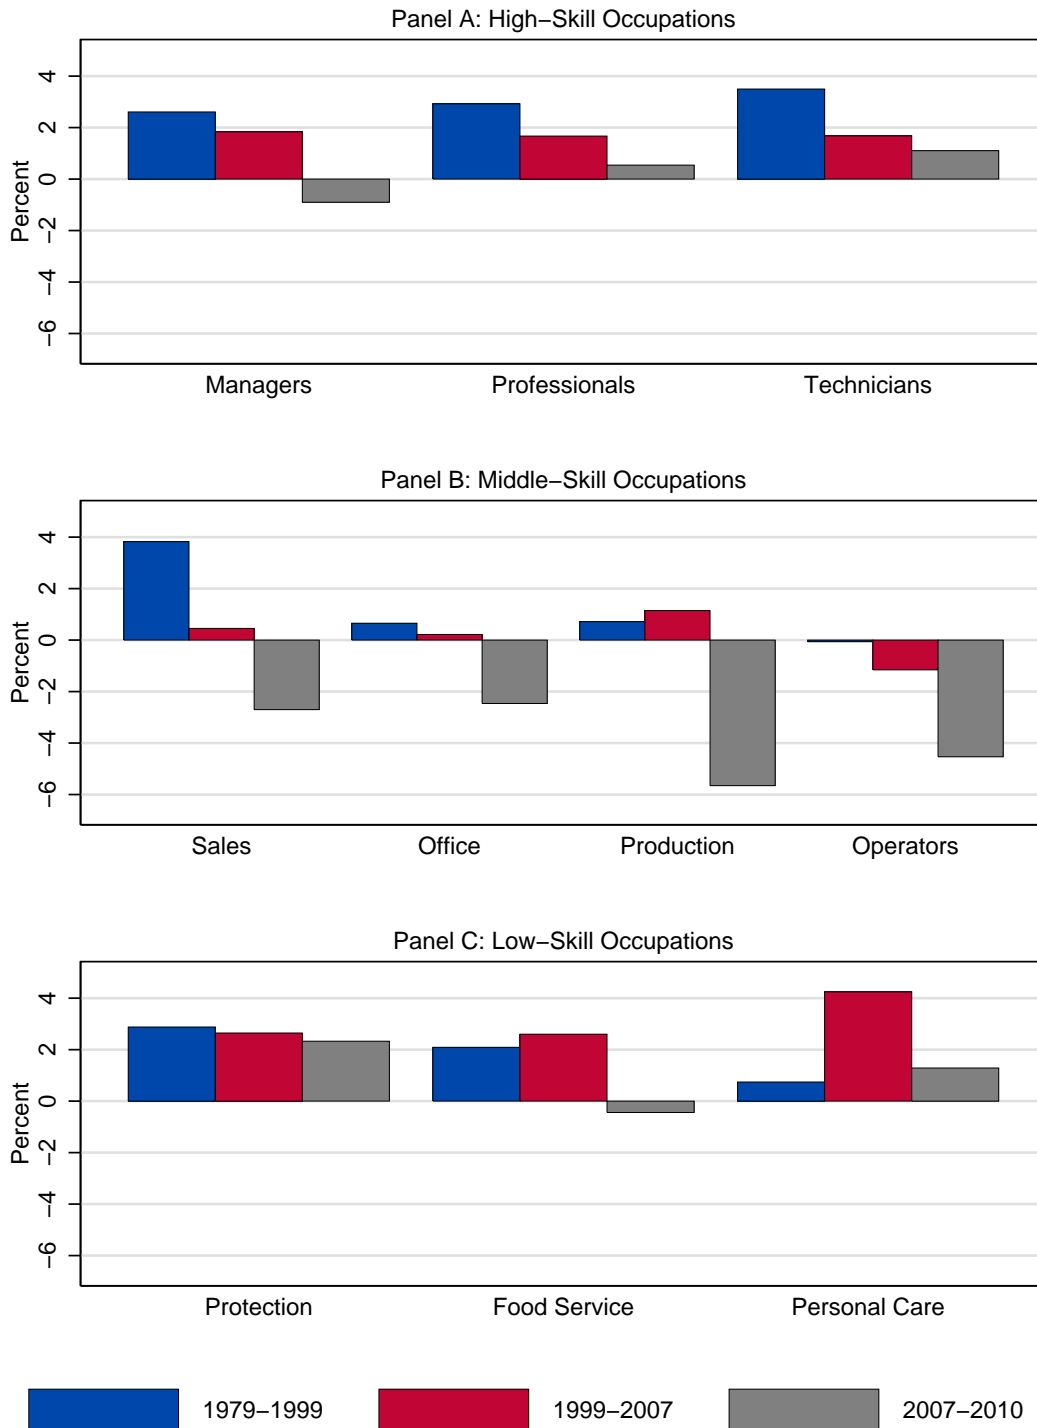


Figure 1. ANNUALIZED EMPLOYMENT GROWTH BY OCCUPATION IN VARIOUS PERIODS. Note: This graph is based on one that originally appeared in Autor (2010).

Theory-Based Occupational Classification	Selected Census-Based Major Occupation Groups Used in Current Population Survey			
	1940 Groups	1970 Groups	1980 Groups	2010 Groups
Nonroutine Cognitive (High Skill)	Professional & Semi-Professional Workers Proprietors, Managers & Officials, excluding Farm	Professional, Technical & Kindred Workers Managers & Administrators, excluding Farm	Professional Specialty Technicians & Related Support Executive, Administrative & Managerial	Professional & Related Occupations Management, Business & Financial Operations Occupations
Routine Cognitive (Middle Skill)	Clerical & Kindred Workers Salesmen and Saleswomen	Clerical & Kindred Workers Sales Workers	Administrative Support, including Clerical Sales Occupations	Office & Administrative Support Occupations Sales & Related Occupations
Routine Manual (Middle Skill)	Craftsmen, Foremen & Kindred Workers Operatives & Kindred Workers Laborers, excluding Farm & Mine	Craftsmen & Kindred Workers Operatives, excluding Transport Transportation Equipment Operatives Laborers, excluding Farm	Precision Production, Craft & Repair Operators, Fabricators & Laborers	Production Occupations Transportation & Material Moving Occupations Installation, Maintenance & Repair Occupations Construction & Extraction Occupations
Nonroutine Manual (Low Skill)	Domestic Service Workers Service Workers, excluding Domestic	Service Workers	Service Occupations	Service Occupations

Table 1. CONSISTENT OCCUPATIONAL GROUPS IN THE CURRENT POPULATION SURVEY. Note: The four theory-based occupational classifications were originally suggested by Jaimovich and Siu (2013).

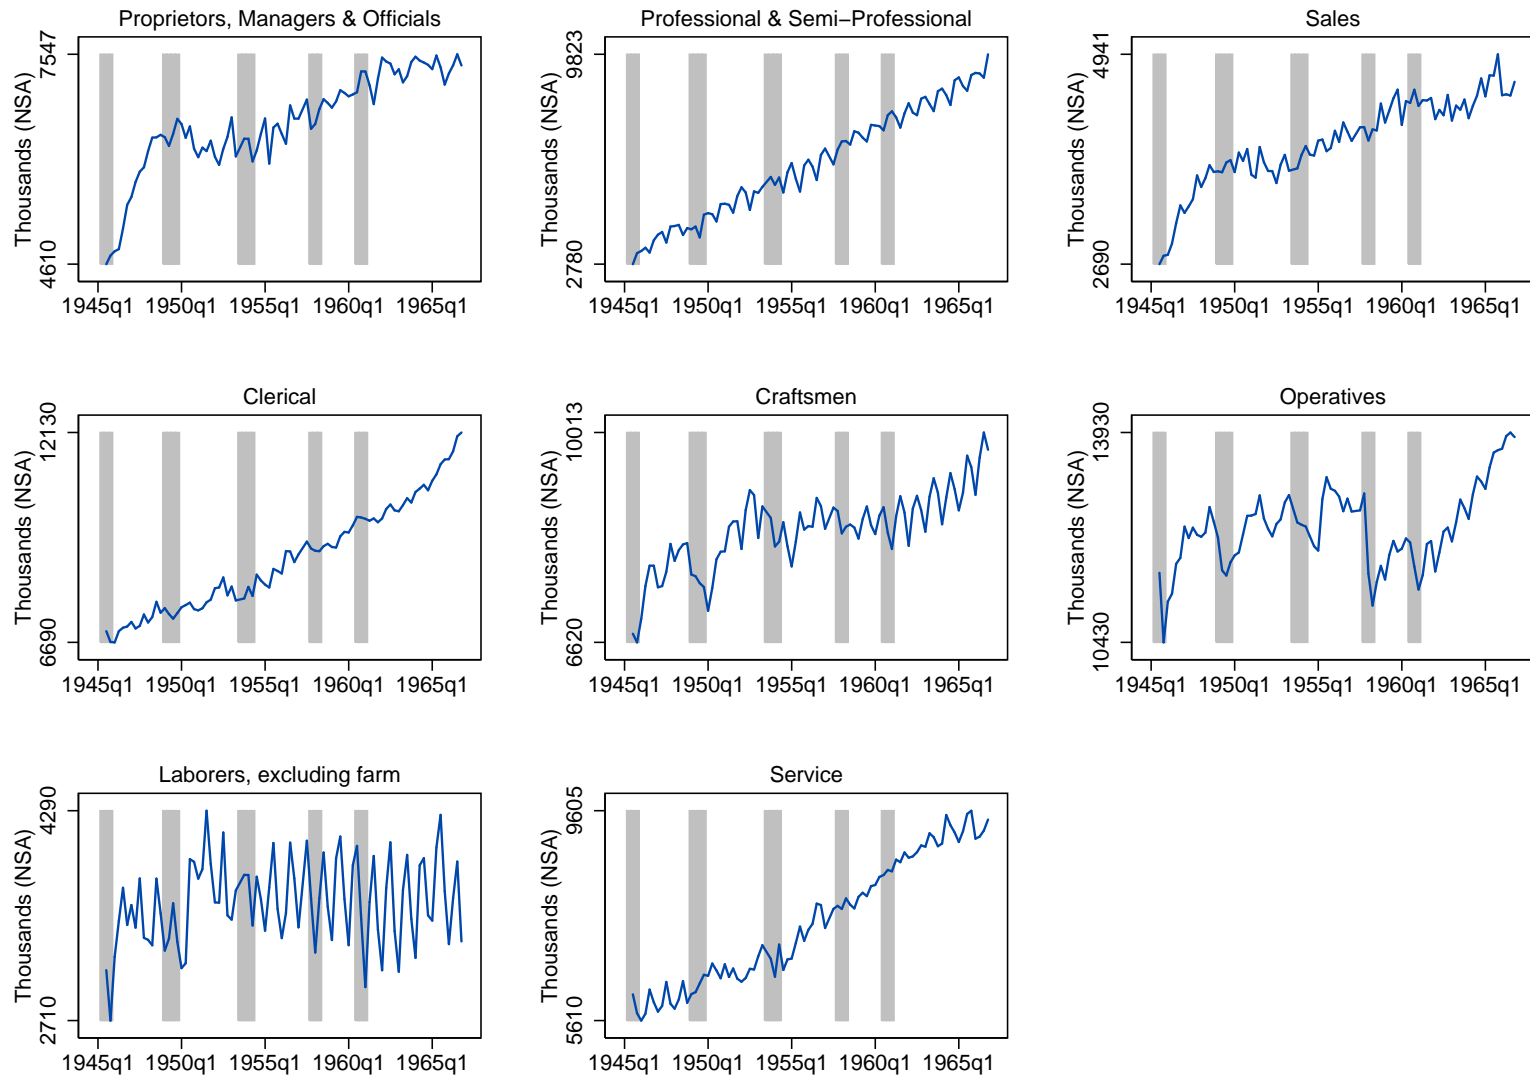


Figure 2. SOME EARLY OCCUPATIONAL DATA.

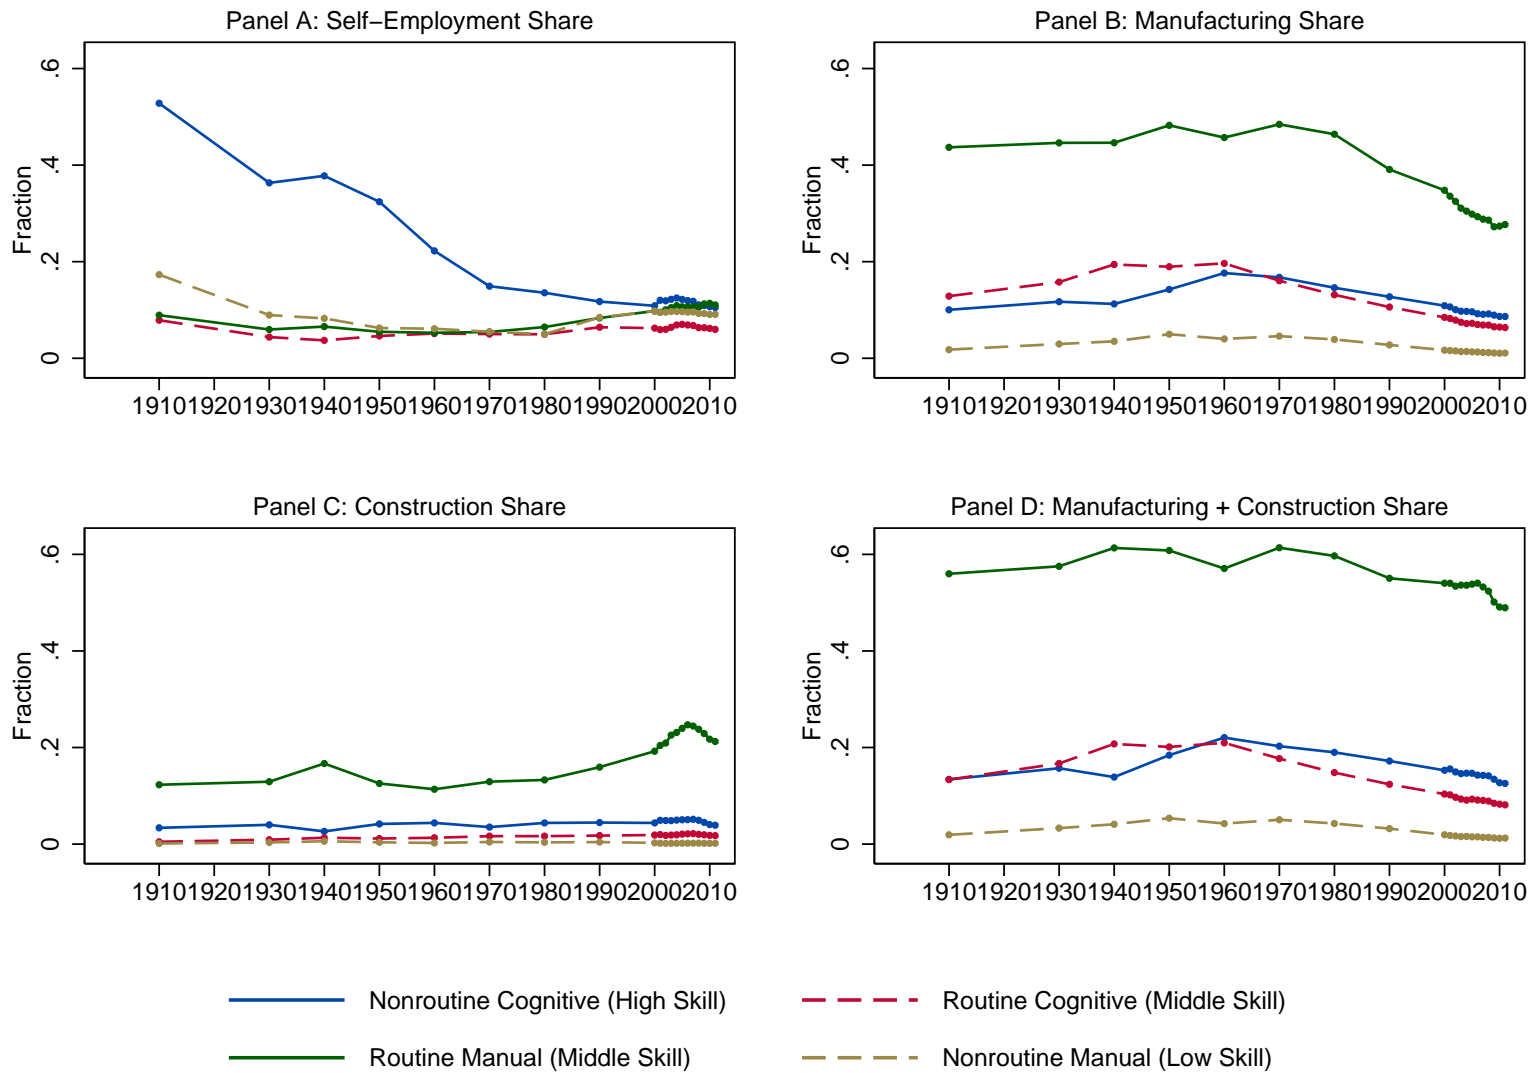


Figure 3. SELECTED CHARACTERISTICS OF OCCUPATIONAL CLASSIFICATIONS. Source: IPUMS.

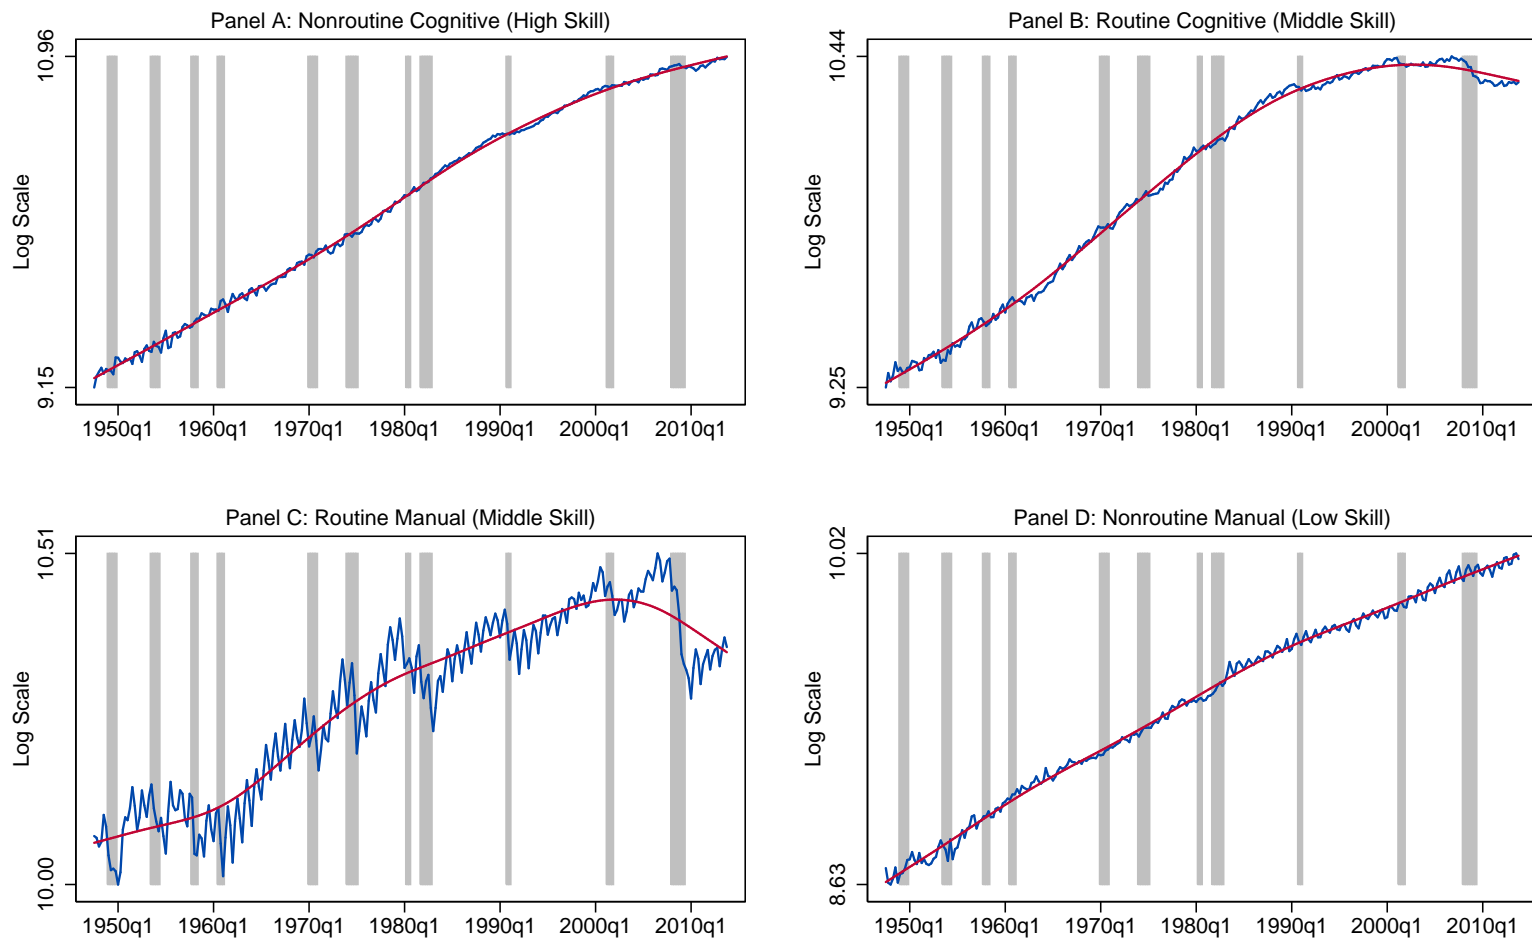


Figure 4. EMPLOYMENT FOR FOUR OCCUPATIONS: 1947Q3-2013Q4: LOG LEVELS AND HP TRENDS ($\lambda = 100,000$).

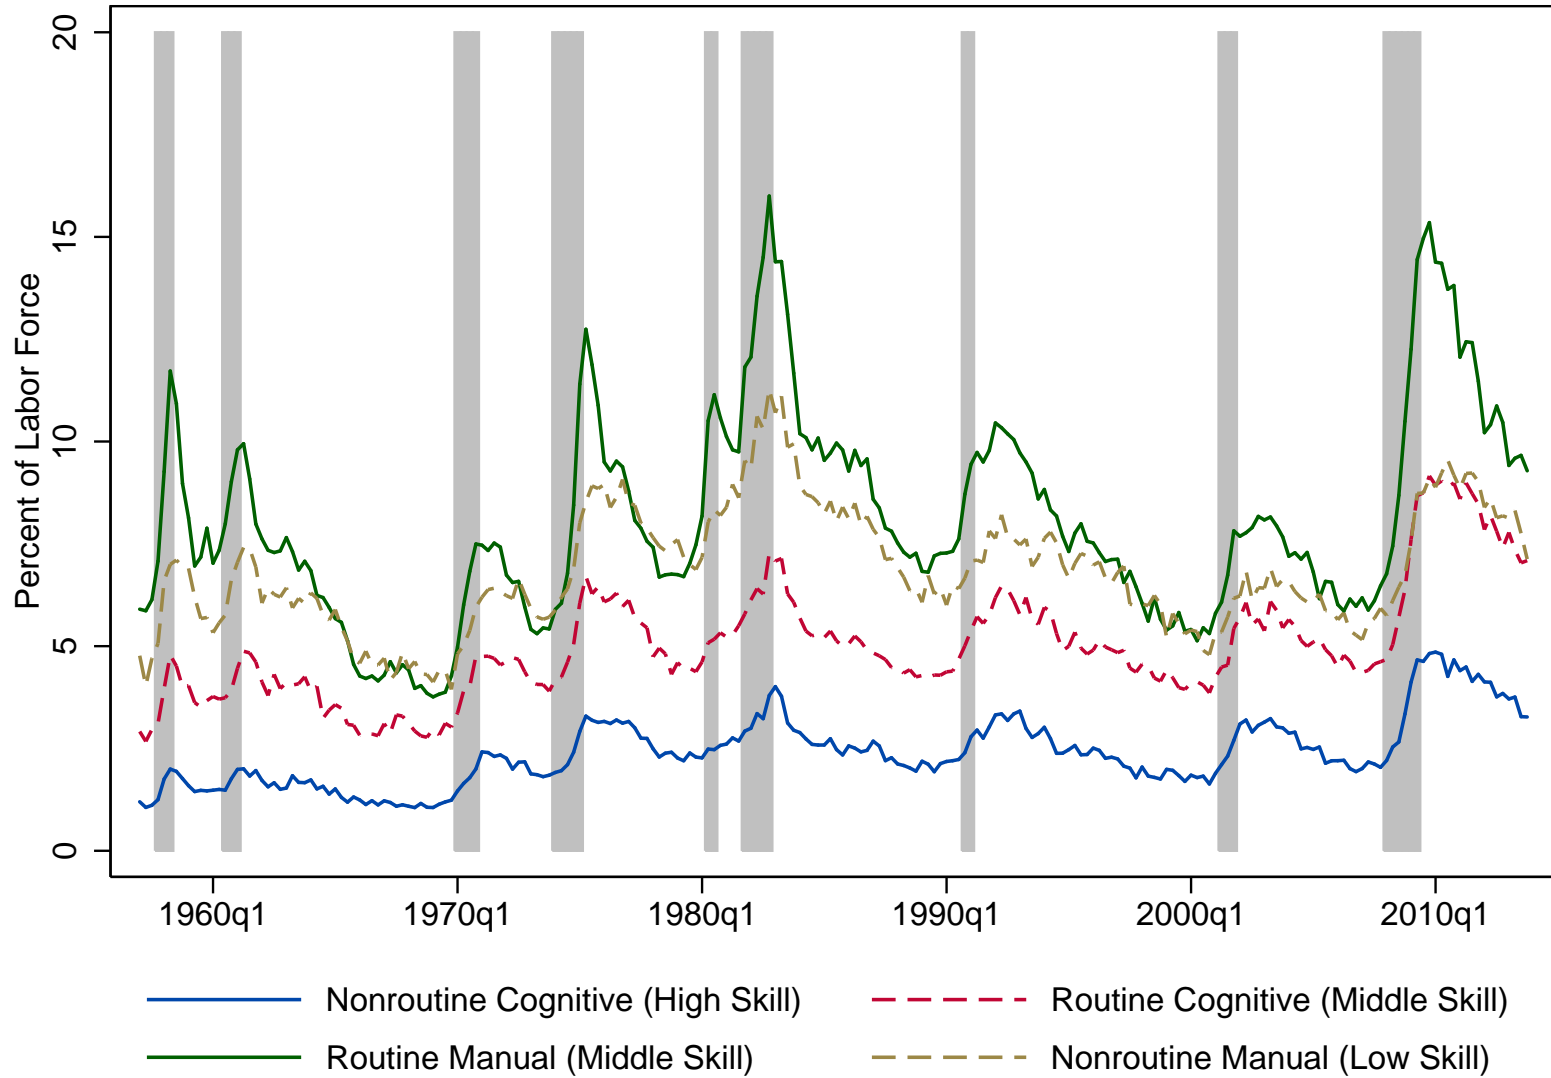


Figure 5. CONSISTENT UNEMPLOYMENT RATES: 1957Q1-2013Q4.

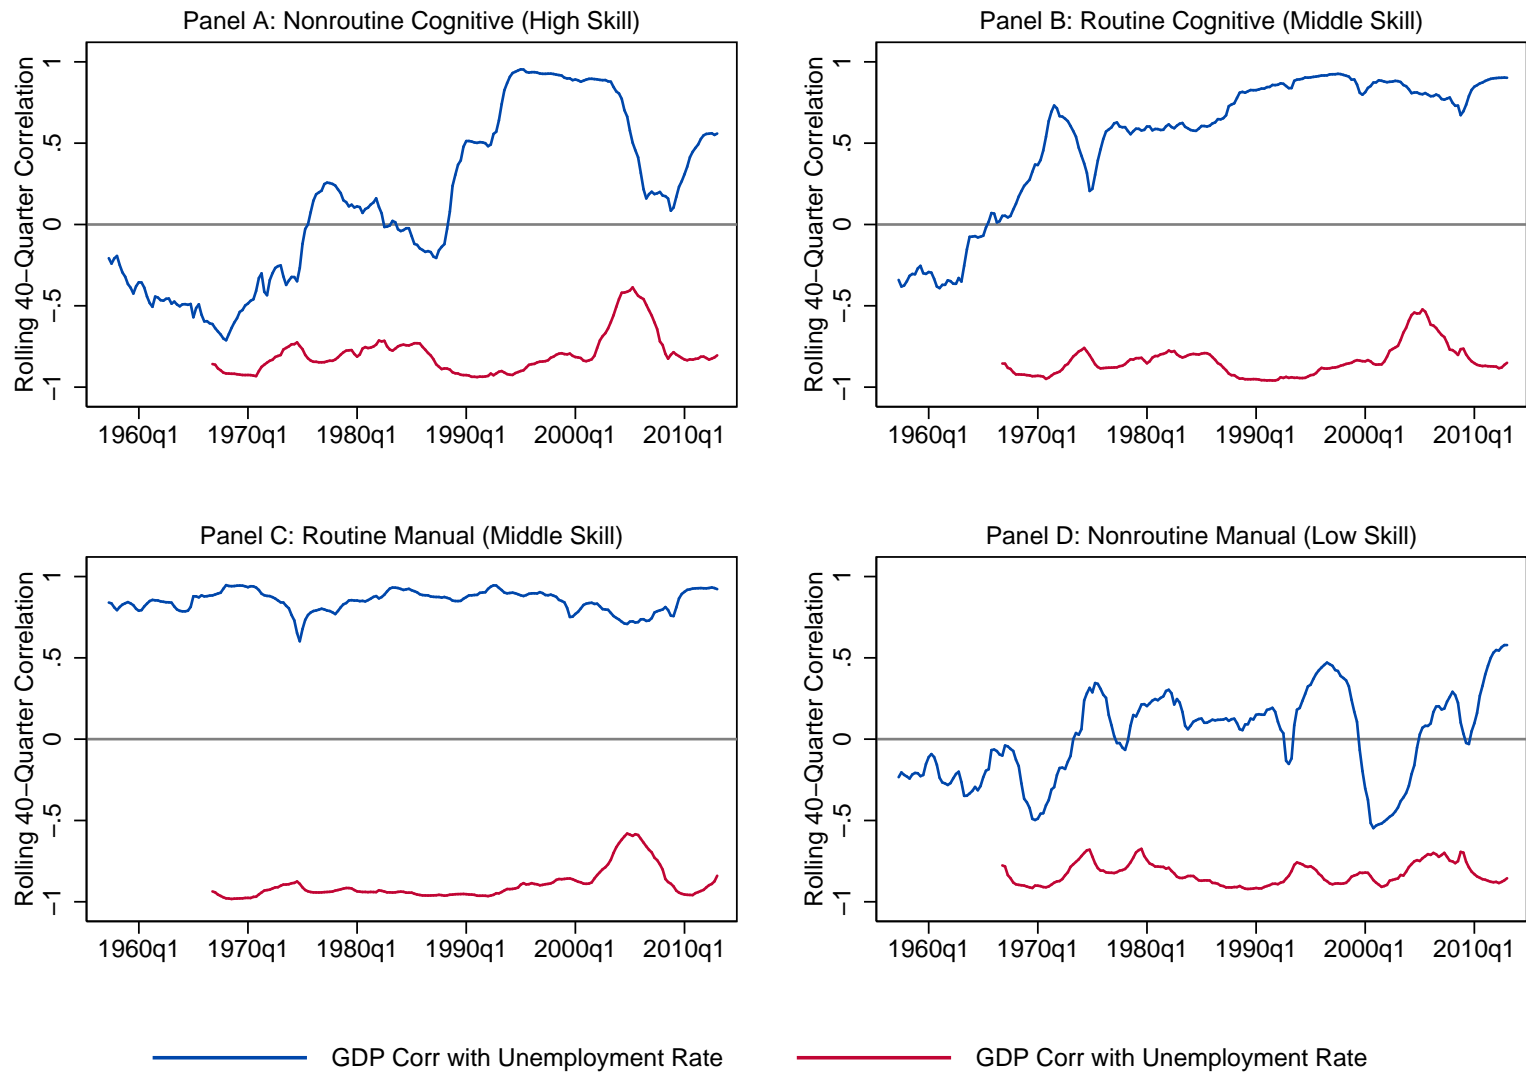




Figure 6. ROLLING CORRELATIONS AMONG OCCUPATION-SPECIFIC EMPLOYMENT LEVELS AND UNEMPLOYMENT RATES AND REAL GDP.


Panel A: 1990–91 Recession



Panel B: 2001 Recession



Panel C: 2008–09 Recession



--- Actual

— Dynamic Forecast: Later Model (1986q1–2013q4)

— Dynamic Forecast: Early Model (1947q3–1985q4)

- · - Dynamic Forecast: Full-Sample Model

Figure 7. DYNAMIC FORECASTS FROM LEVELS-BASED DYNAMIC FACTOR MODEL.

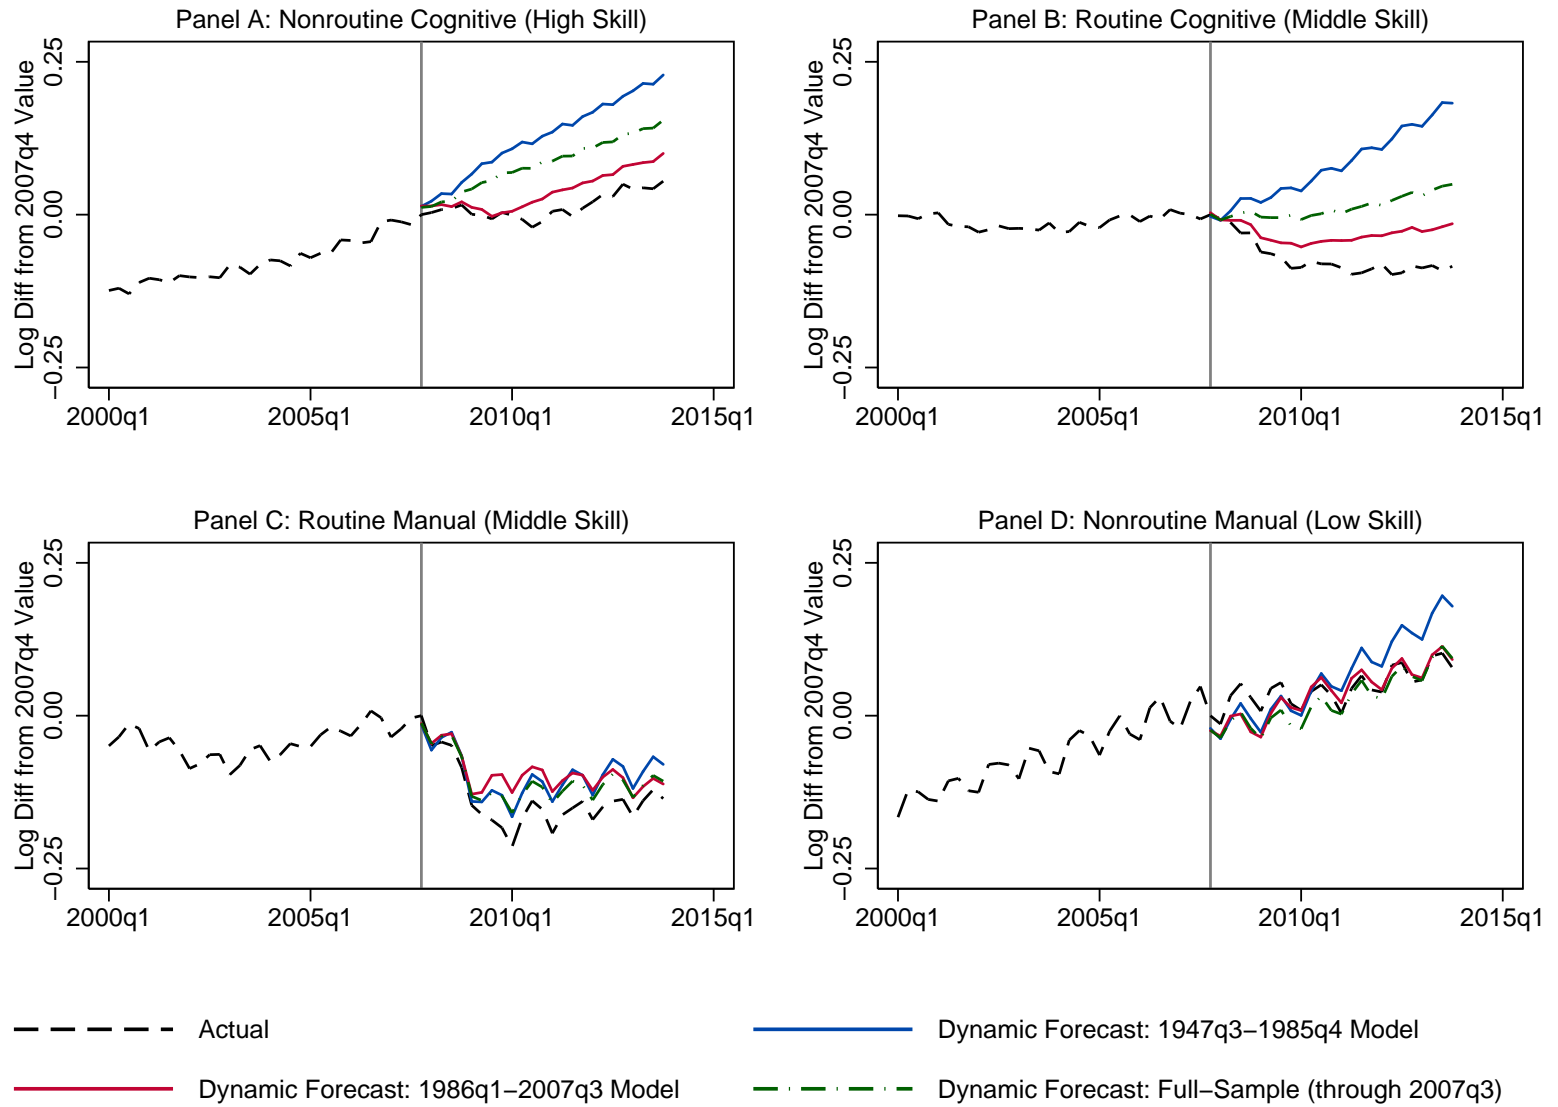


Figure 8. DYNAMIC FORECASTS FROM DIFFERENCED MODEL OF OCCUPATIONAL EMPLOYMENT GROWTH.

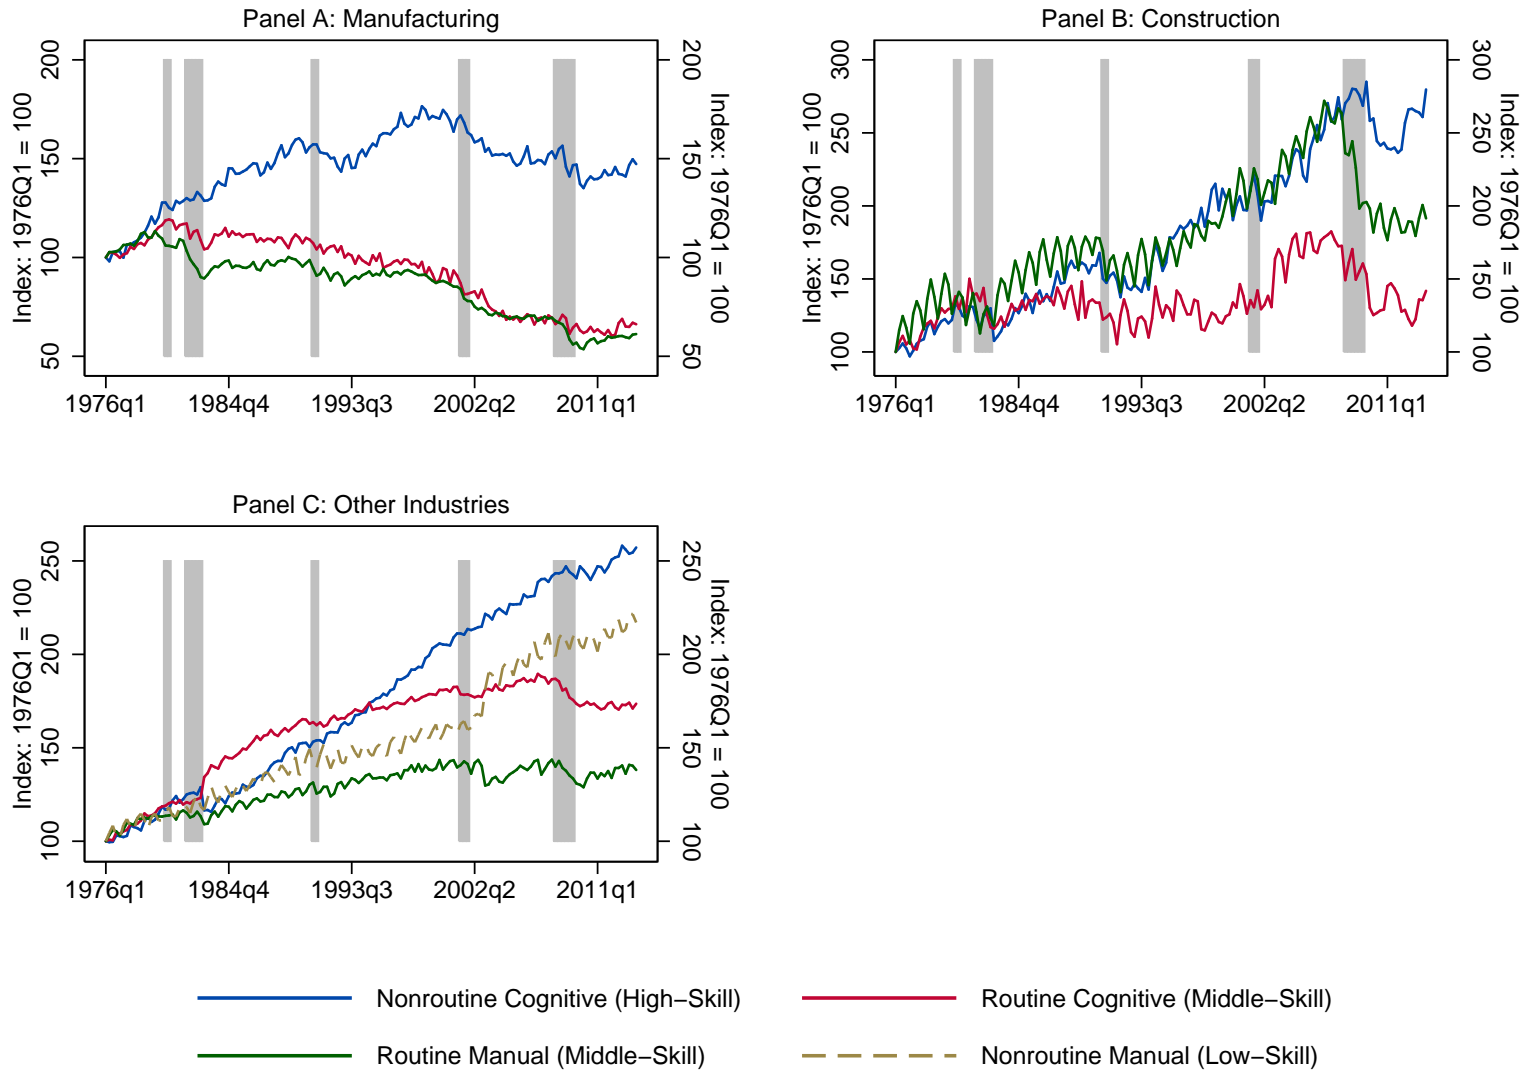


Figure 9. EMPLOYMENT GROWTH BY OCCUPATIONAL GROUP AND INDUSTRY.

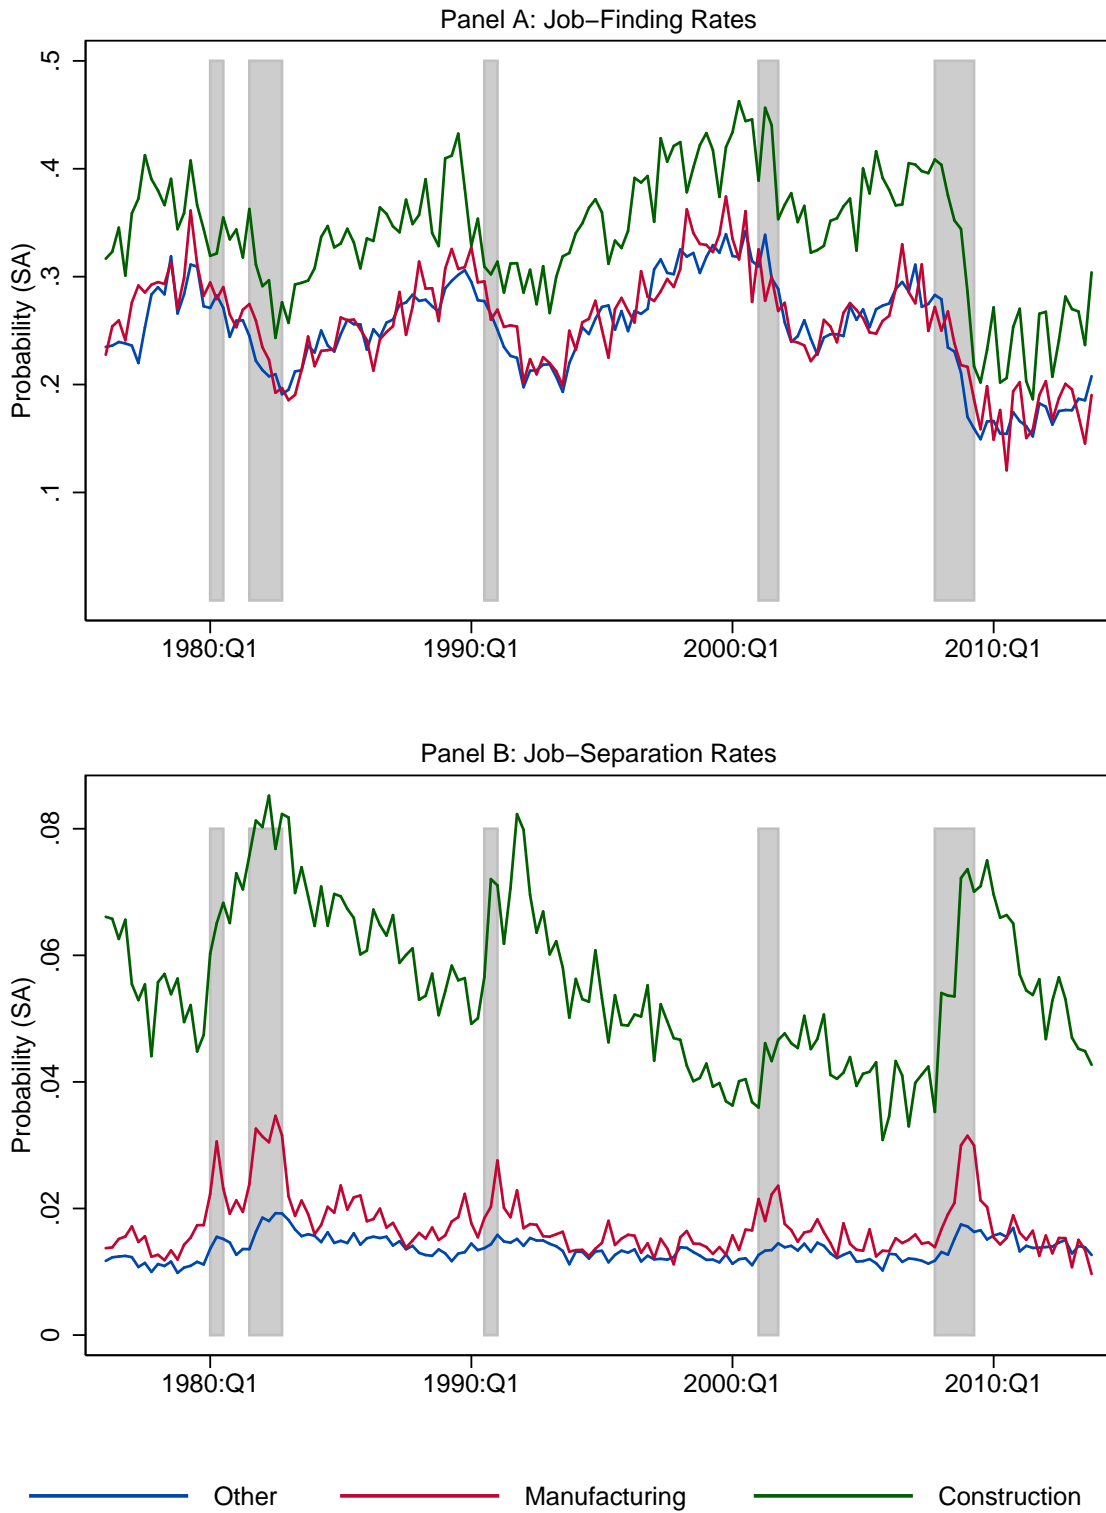




Figure 10. JOB-FINDING (U-TO-E) AND JOB-SEPARATION (E-TO-U) RATES FOR MIDDLE-SKILL WORKERS IN THREE INDUSTRIES.

Panel A: EE Flows (1994–2013)



Panel B: UE Flows (1976–2013)



— To High
 — To Mid–Other
 — To Mid–Mfg.
 — To Mid–Const.
 — To Low

Figure 11. COMPOSITION OF JOB-FINDING (UE) AND JOB-SEPARATION (EU) RATES.

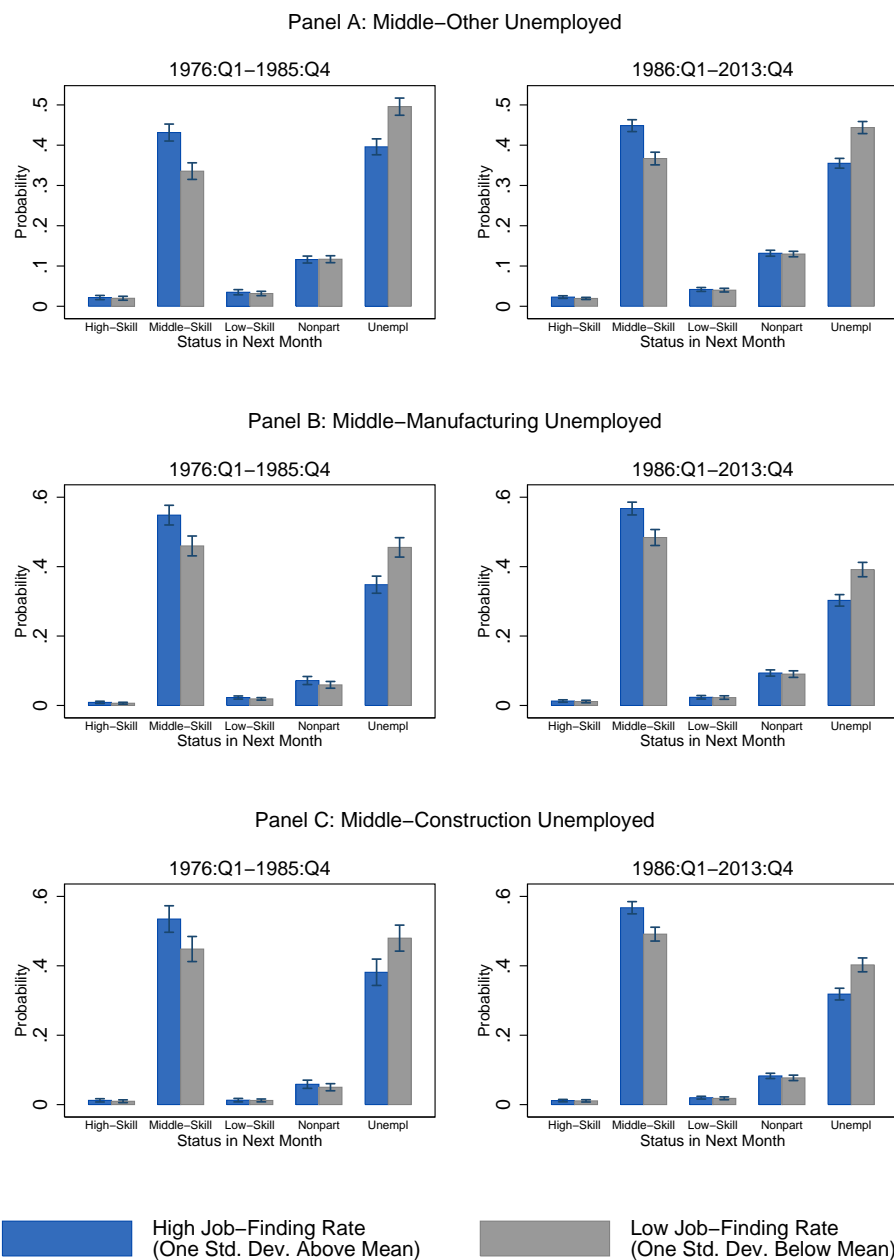


Figure 12. TRANSITION PROBABILITIES FOR A BASELINE MIDDLE-SKILL WORKER WITH HIGH-SCHOOL DIPLOMA. Probabilities are estimated from multinomial logits of job-finding rates. Separate logits are run in the two sample periods. A “high” job finding rate corresponds to a common factor of the five industry-skill job-finding rates that is one standard deviation above its mean, while a “low” job finding rate is one standard deviation below. The five industry-skill job-finding rates correspond to the following classifications: high, middle-other, middle-manufacturing, middle-construction, and low. Regressors in the multinomial logit include demographic and duration dummies defined so that the baseline worker is an unmarried 35-year-old white male with a high-school diploma who reports zero to one weeks of unemployment duration. Standard errors clustered by quarter.

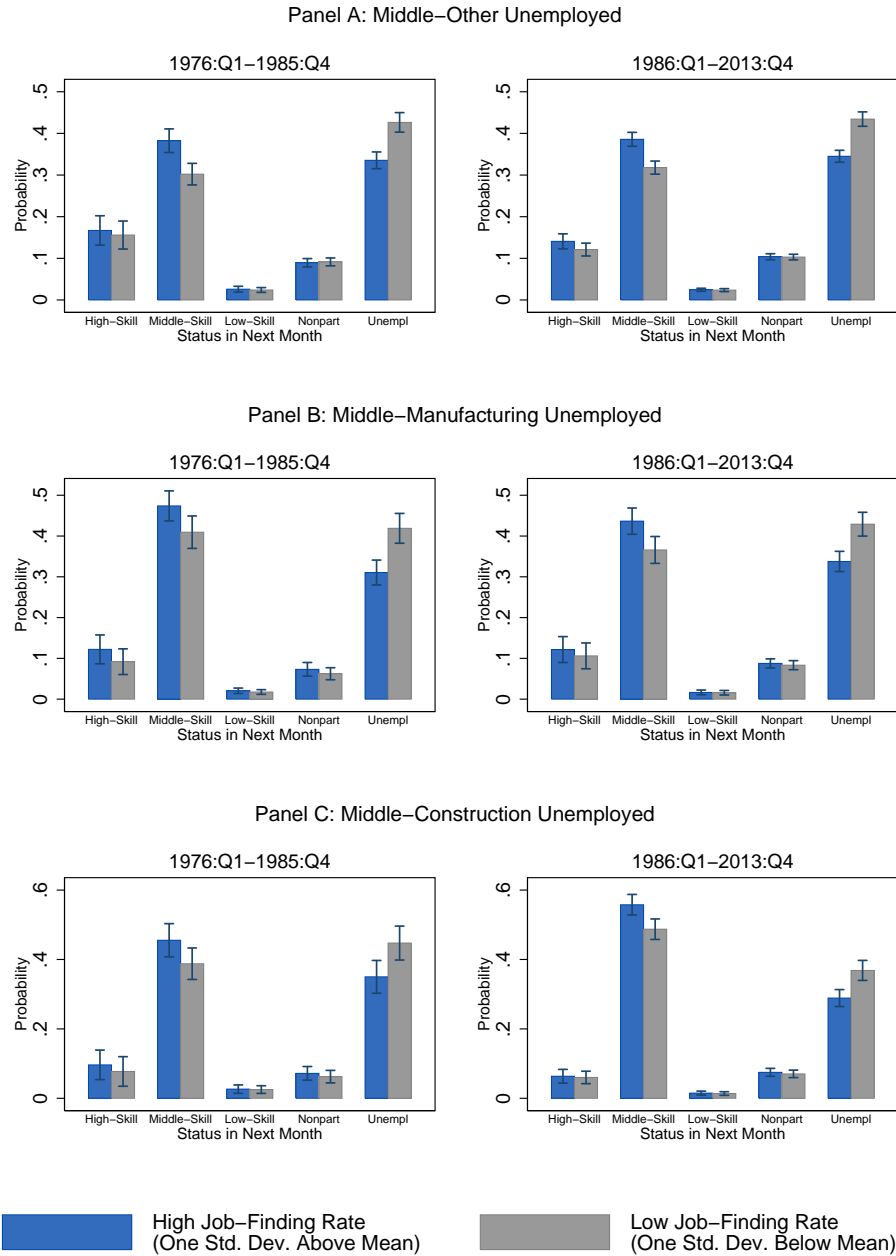


Figure 13. TRANSITION PROBABILITIES FOR A BASELINE MIDDLE-SKILL WORKER WITH COLLEGE DEGREE. See notes to Figure 12 for specification of model. Probabilities here correspond to a worker with a college degree rather than a high-school diploma.

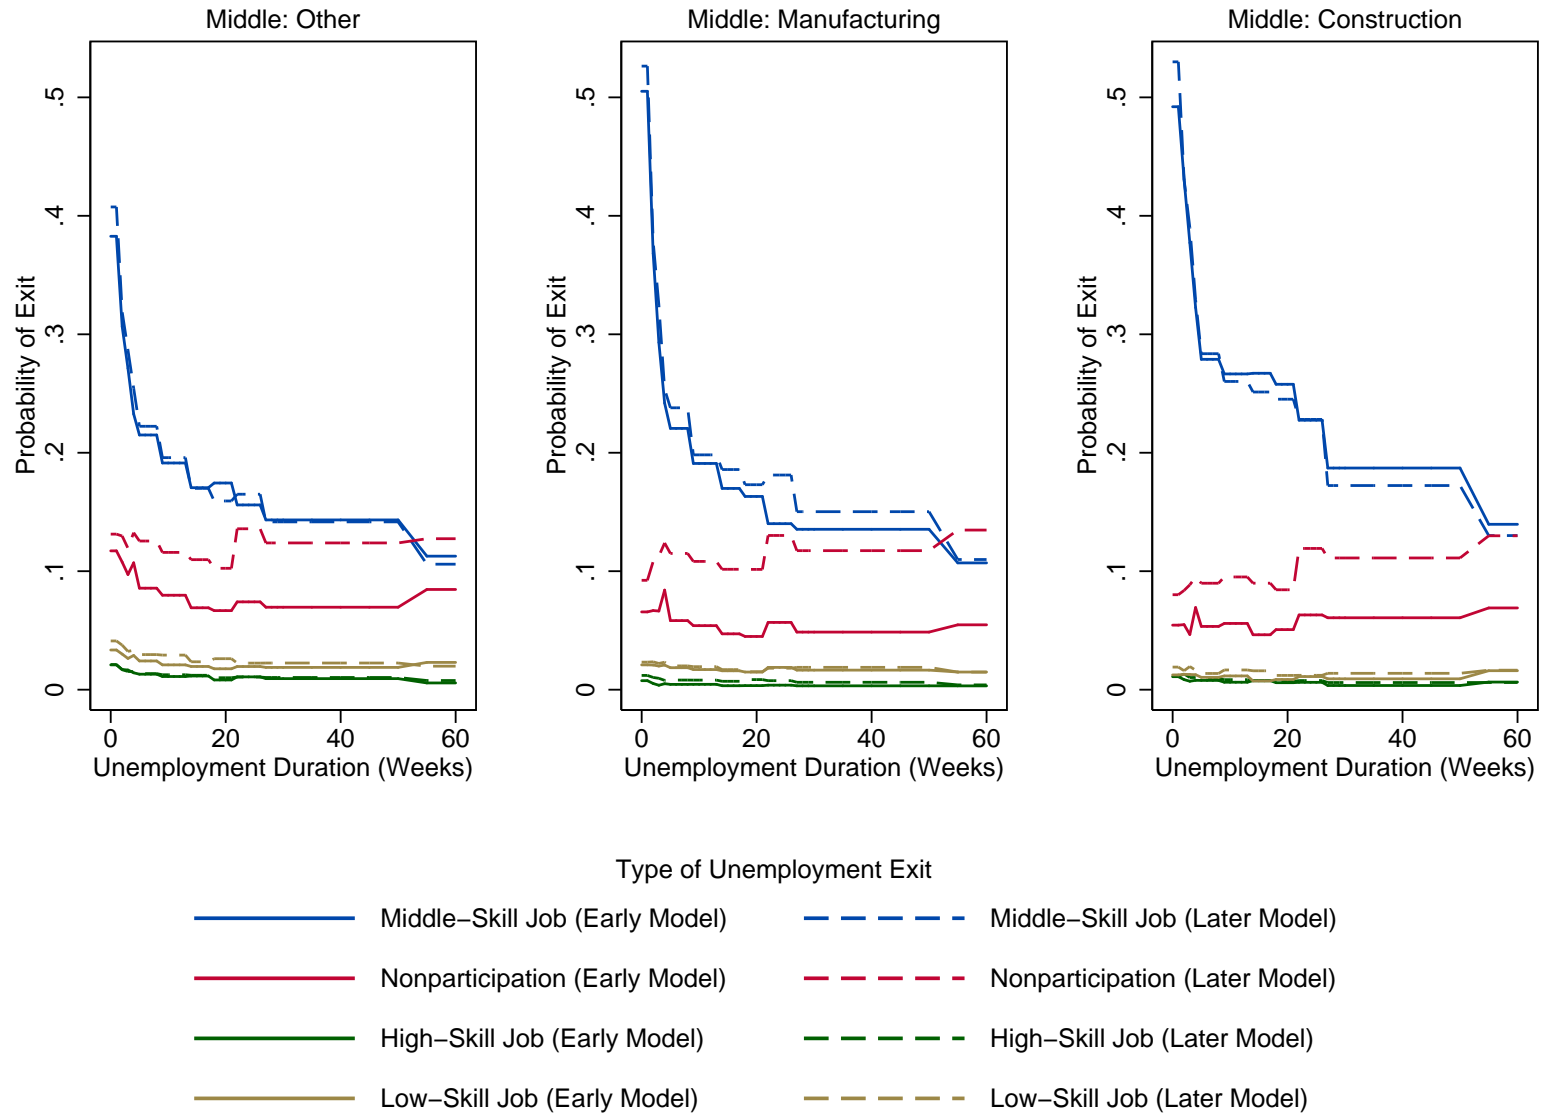


Figure 14. EFFECT OF UNEMPLOYMENT DURATION ON MIDDLE-SKILL UNEMPLOYMENT EXITS IN TWO SAMPLE PERIODS. Note: Probabilities are for a baseline worker with a high-school education, as described in the notes to Figure 12.

Panel A: Prime Age Men					
	(1)	(2)	(3)	(4)	(5)
High-Skill Share		27.25*** (6.89)			3.76 (2.07)
Middle-Skill Share			-26.16*** (6.58)		-25.10*** (5.25)
Low-Skill Share				4.63 (15.31)	-10.50 (14.42)
Constant	-4.05*** (0.81)	-23.31*** (5.10)	2.45 (1.69)	-4.26*** (1.04)	
Observations	216	216	216	216	216
Panel B: Older Men					
	(1)	(2)	(3)	(4)	(5)
High-Skill Share		36.89*** (10.10)			2.94 (2.59)
Middle-Skill Share			-34.36*** (7.51)		-25.65*** (7.61)
Low-Skill Share				47.04*** (13.69)	14.96 (16.02)
Constant	-4.60*** (1.35)	-32.08*** (7.89)	4.37* (1.96)	-4.31*** (1.10)	
Observations	144	144	144	144	144

Robust (White) standard errors in parentheses. Age, education, and Census Division dummies are included in all regressions.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2. LONG-DIFFERENCE REGRESSIONS FOR MALE LABOR-FORCE PARTICIPATION RATES. Note: The constant is dropped from the regressions in Column 5.

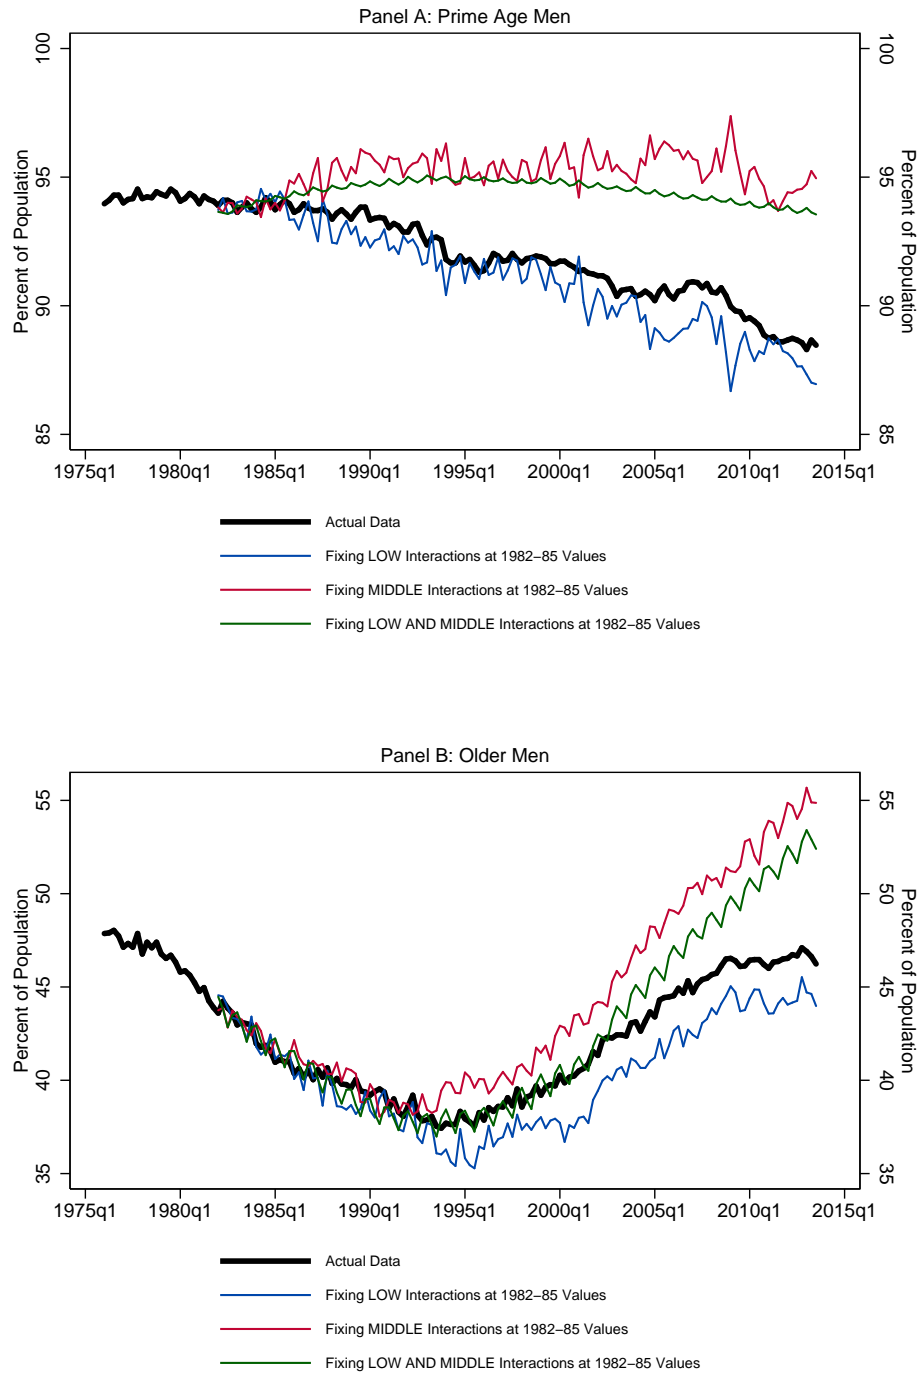


Figure 15. COUNTERFACTUAL LABOR FORCE PARTICIPATION RATES FOR MEN. Note: For counterfactual participation rates, coefficient estimates on skill-time interactions in the regression model are held constant at their average 1982-1985 values.

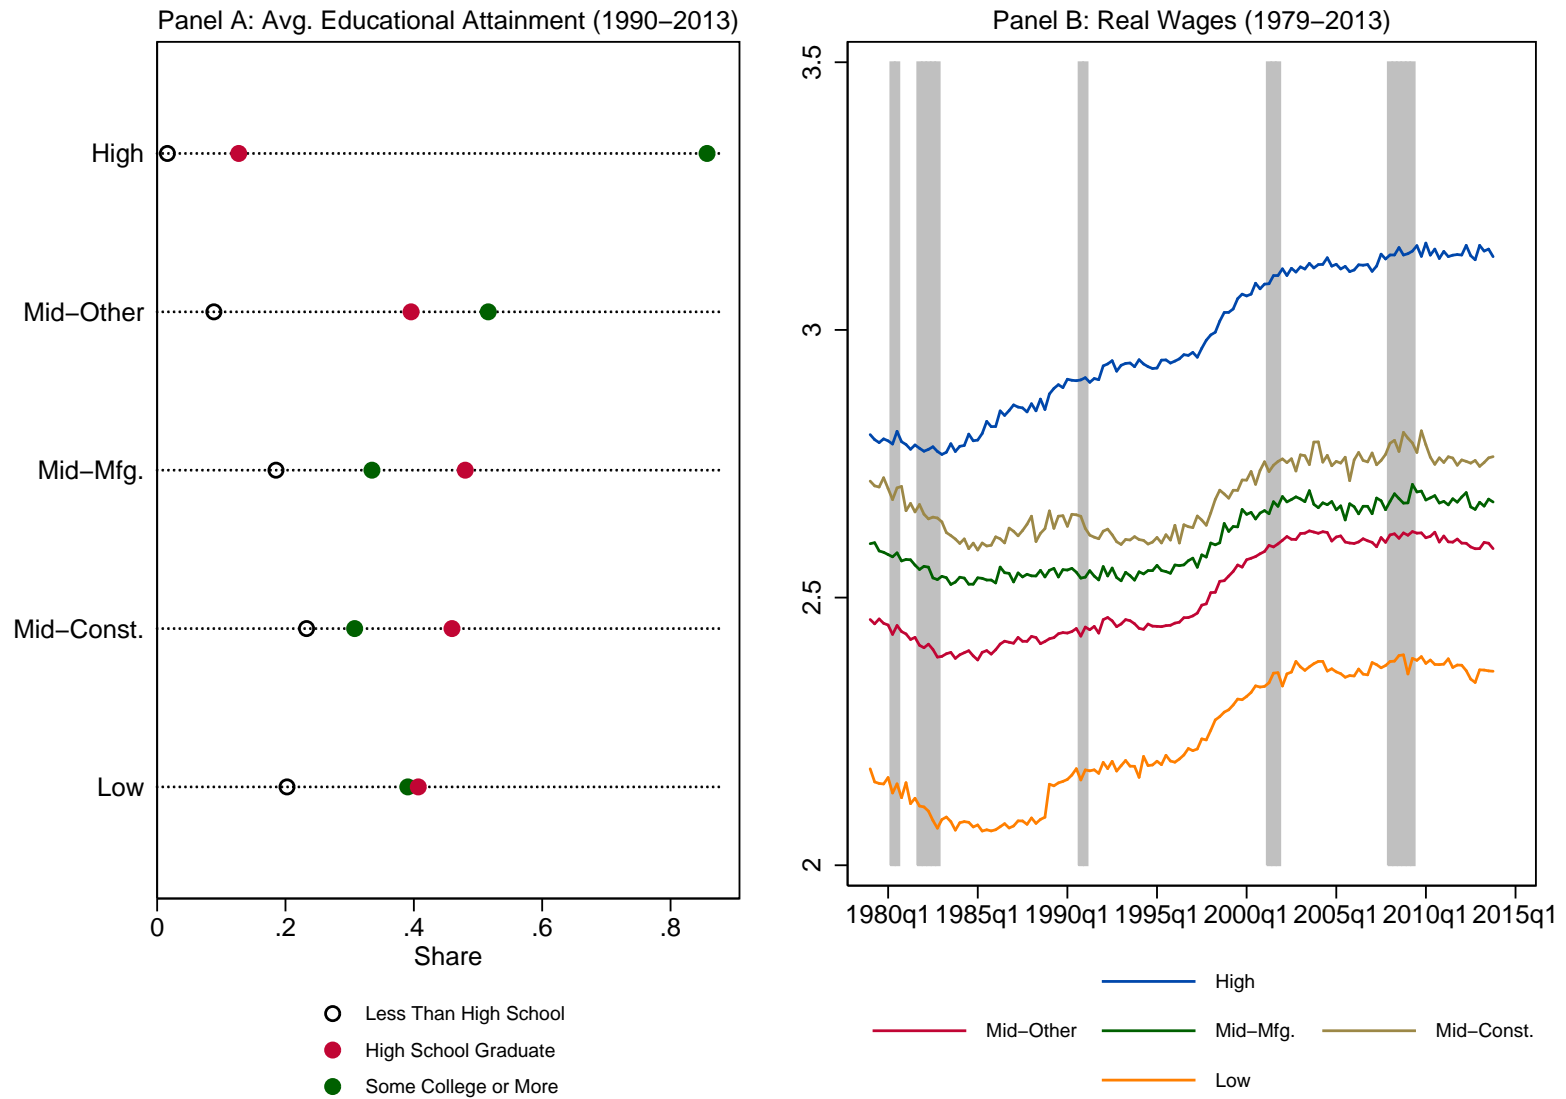


Figure 16. EDUCATIONAL ATTAINMENT AND REAL WAGES FOR FIVE INDUSTRY/OCCUPATION GROUPS.