

Job Flows, Jobless Recoveries, and the Great Moderation

June 2008

R. Jason Faberman

Federal Reserve Bank of Philadelphia
Ten Independence Mall, Philadelphia, PA 19106
Jason.Faberman@phil.frb.org
Phone: (215) 574-6440
Fax: (215) 574-4303

Abstract:

This paper uses new data on job creation and job destruction to find evidence of a link between the jobless recoveries of the last two recessions and the recent decline in aggregate volatility known as the Great Moderation. I find that the last two recessions are characterized by jobless recoveries that came about through contrasting margins of employment adjustment—a relatively slow decline in job destruction in 1991-92 and persistently low job creation in 2002-03. In manufacturing, I find that these patterns followed a secular decline in the magnitude of job flows and an abrupt decline in their volatility. A structural VAR analysis suggests that these patterns are driven by a decline in the volatilities of the underlying structural shocks *in addition to* a shift in the response of job flows to these shocks. The shift in structural responses is broadly consistent with the change in job flow patterns observed during the jobless recoveries.

Keywords: job flows, jobless recoveries, business cycles, aggregate volatility, Great Moderation

JEL Codes: E24, E32, J36

Acknowledgments

I thank Steve Davis, John Haltiwanger, Scott Schuh, Keith Sill, Tom Stark, Dan Sullivan, and seminar participants at the Federal Reserve Board of Governors, the Bureau of Labor Statistics, Lehigh University, the Federal Reserve Bank of Philadelphia, and the Midwest Macro Meetings for helpful comments. Any remaining errors are my own. The views expressed are solely those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper is available free of charge at www.philadelphiafed.org/econ/wps/.

1. Introduction

It is clear that aggregate economic activity has become considerably less volatile over the last quarter-century. Most economists refer to this phenomenon, first reported by Kim and Nelson (1999) and McConnell and Perez-Quiros (2000), as the “Great Moderation.” A wealth of research has emerged on, among other things, its sources (whether it is the result of good policies, or simply good luck) and its dynamics (whether the decline is a break in the time series or part of a long-run trend).¹

It is also clear that the economy in recent years has seen a divergence in the cyclical patterns of output and employment. In particular, following the 1990-91 and 2001 recessions, employment began to grow well after output recovered, leading many to label the post-recession periods as “jobless recoveries.” These recoveries have also garnered significant attention, with researchers putting forth several explanations.² While both the Great Moderation and the jobless recoveries reflect pronounced changes in the cyclical dynamics of economic activity that roughly occur concurrently, there is nearly no work to date that attempts to link the two.³

In this paper, I present new evidence on the behavior of job creation and job destruction as it relates to the Great Moderation and the last two economic downturns. I draw upon both new and existing data on gross job flows to highlight several new stylized facts. First, while the *net* employment growth patterns appear similar during the last two downturns, the *gross* job flow patterns are quite different. That is, the jobless

¹ Notable examples of this research include Blanchard and Simon (2001), Stock and Watson (2002), Ahmed, Levin, and Wilson (2004), Ramey and Vine (2006), and Justiniano and Primiceri (2007).

² See Groshen and Potter (2003), Schreft and Singh (2003), Aaronson, Rissman, and Sullivan (2004), Koenders and Rogerson (2005), and Bachmann (2007).

³ A notable exception is Engemann and Owyang (2007), who use a Bayesian approach to study an econometric link between jobless recoveries and the Great Moderation.

recovery following the 2001 recession occurs because of a large, persistent decline in the job creation rate. In contrast, the 1990-91 downturn shows little change in the job creation rate and has an above-average job destruction rate during the recovery period. This contrast appears to occur across a broad range of industries, both within and outside of manufacturing. Second, like many other measures of real economic activity, job creation and job destruction rates within manufacturing experienced a marked decline in their volatility during the mid-1980s, the period usually identified with the Great Moderation. The declines, however, are asymmetric. The volatility of job destruction falls much more than the volatility of job creation, causing the relative volatility of destruction to creation to fall considerably. This asymmetry is notable particularly in light of the conventional wisdom that has emerged from the early work of Davis and Haltiwanger (1990, 1992) and Davis, Haltiwanger, and Schuh (1996), who argued that cyclical employment fluctuations were driven primarily by movements in the job destruction rate.⁴ My evidence suggests that this argument may not apply to future business cycles. In addition, I find that the *magnitudes* of job creation and destruction within manufacturing also experience a secular decline over this period. In contrast to the volatility evidence (and previous research on the Great Moderation), this decline exhibits a steady trend that begins in the early 1960s, rather than a sharp drop in the 1980s. If one thinks of the magnitudes of job creation and job destruction as measures of the cross-sectional dispersion of micro-level employment fluctuations, this finding suggests that this dispersion has declined by about one-third over the last 45 years.

⁴ Foote (1998) provides evidence that the Davis-Haltiwanger argument may apply only to manufacturing, though in this study I show that cyclical movements in job destruction are large, even outside of manufacturing.

While interesting in their own right, these findings raise more questions than they answer. In particular, they raise the possibility that the jobless recoveries of the last two downturns may be the result of structural changes in the labor market (perhaps related to the Great Moderation) rather than anomalies of the business cycle. I explore this possibility with a structural vector autoregression (VAR) approach developed by Davis and Haltiwanger (1999) that decomposes movements in job creation and job destruction into structural responses to either aggregate or allocative shocks. I split my time series on manufacturing job flows into pre- and post-1984 samples and run the structural VAR separately for each period.

While some results are sensitive to the identifying assumptions used in the structural VAR, several notable findings emerge. First, the volatilities of both aggregate and allocative shocks are lower in the post-1984 period. The standard deviation of aggregate shocks falls between 55 and 60 percent, while the standard deviation of allocative shocks falls between 31 and 54 percent, depending on the identifying assumptions. These declines either decrease the relative volatility of aggregate to allocative shocks or leave it unchanged, again, depending on the identifying assumptions.

I also find that the responses of job creation and job destruction to aggregate shocks are considerably different in the two periods. The 1947-1983 period shows responses to a negative aggregate shock that are relatively brief – a large spike in job destruction dissipates within several quarters, and a smaller decline in job creation is quickly followed by a subsequent increase. The 1984-2006 period, however, exhibits responses of both job creation and job destruction that are much more persistent, and unlike the early period, there is no rebound in job creation. The responses to an allocative

shock show few differences between the two periods, and all impulse responses are qualitatively similar regardless of the identifying assumptions.

Thus, the structural VAR results suggest that, within manufacturing, the size of both aggregate and allocative shocks declined substantially in the post-1984 period. The results also show a structural change in how employment responds to aggregate shocks. This change is consistent with the dynamics observed during the two jobless recoveries: job creation and job destruction have more persistent responses to an aggregate shock, and job creation no longer exhibits a “rebound” effect following an aggregate shock, leading to a longer period of negative employment growth. When coupled with the evidence of declining job flow magnitudes and volatilities, and an increase in the relative volatility of job creation to job creation, the findings suggest a logical link between the Great Moderation and the jobless recoveries. In the labor market, the Great Moderation occurs in the form of smaller shocks and altered responses to these shocks. The jobless recoveries are a consequence of the altered job flow responses to an aggregate shock.

Evidence on other facets of the labor market also suggests a change in the cyclical behavior of employment fluctuations. For example, inflows into the unemployment pool have experienced a secular decline and a decrease in volatility starting in the mid-1980s.⁵ The decline in volatility is driven by a virtual disappearance of cyclical movements in temporary layoffs, which, in turn, occurs coincidentally with a rise in production worker hours within the temporary help industry. A shift in firm behavior from using temporary layoffs to temporary help as their primary mechanism for employment adjustment may be the source of the altered response of employment to aggregate shocks, though further research is needed. Regardless, the evidence here clearly shows that the labor market

⁵ See Davis, Faberman, Haltiwanger, Jarmin, and Miranda (2008).

during the Great Moderation period is characterized by a change in its structural responses to shocks—in addition to a decline aggregate volatility—that has fundamentally altered its dynamic behavior over the business cycle.

The next section reviews previous research on jobless recoveries and the Great Moderation. Section 3 presents the evidence on job flows during the last two economic downturns. Section 4 presents the evidence of the declining magnitudes and volatilities of job flows within manufacturing over a much longer horizon. Section 5 presents the structural VAR analysis of the manufacturing job flows, as well as the supporting evidence on unemployment and temporary help employment, while Section 6 concludes.

2. Background

Several recent papers have explored the topic of jobless recoveries, while a much larger literature exists on the Great Moderation. Koenders and Rogerson (2005) present a model where jobless recoveries are the outcome of prolonged economic expansions. They occur because the inefficiencies built up during these expansions require an extended “shakeout” period. Bachmann has an alternative model where jobless recoveries are instead the result of relatively mild recessions. In these cases, firms choose to adjust hours rather than employees during the recession, leading to less hiring during the recovery period. Empirically, Groshen and Potter (2003) argue that the evidence on jobless recoveries may be the result of a structural reallocation of employment across industries. In my evidence below, I show that the job flow evidence during the jobless recoveries lends some credence to both the Koenders-Rogerson and Bachmann arguments, but neither of their models can fully characterize the employment adjustments observed during both jobless recoveries. I also show that job flows exhibited similar

patterns across industries during both jobless recoveries, suggesting that that the recoveries were likely not the result of an across-industry reallocation of employment.

Much of the research on the Great Moderation has focused on the volatility of output and its components. It has built on the work of Kim and Nelson (1999) and McConnell and Perez-Quiros (2000), who first identified the decline in volatility and dated its occurrence as a structural break in the time series that occurred around 1984. Some researchers have looked at whether the decline in volatility was due to “good luck”, in the form of smaller aggregate shocks, or improved policy (Stock and Watson, 2002; Ahmed, Levin and Wilson, 2004). Others have argued that the decline is more of a trend decline than a break in the time series (Blanchard and Simon, 2001). Others still have looked at whether the decline in volatility is the result of more structural changes (Ramey and Vine, 2006; Justiniano and Primiceri, 2007). While the research on the topic is broad, there is little consensus on either the causes or the consequences of the Great Moderation.

The research on the Great Moderation has also focused mostly on output, though there are some recent studies that focus on the labor market. Galí and Gambetti (2007) study the volatility and comovement of hours, output, and labor productivity. They highlight the fact that the correlations between hours, output, and productivity fundamentally changed during the Great Moderation period. Stroh (2007) takes a similar approach through a “volatility accounting” exercise that focuses more on the production than the final demand perspective of the economy.

A focus on potential structural changes in addition to changes in aggregate volatility is logical when studying the labor market, given some recent empirical evidence. For example, Aaronson, Rissman, and Sullivan (2004), among others, note how

the cyclical nature of temporary layoffs all but disappeared during the 1980s. Figure 1 reinforces this point using unemployment data from the Current Population Survey. The recessions prior to 1984 show large spikes in both temporary and permanent layoffs, as well as increases in labor force entrants who start as unemployed. The recessions following 1984 show increases in unemployment due to permanent layoffs and labor force entry, but almost no cyclical movement in temporary layoffs. This occurs simultaneously with a steady rise in temporary help employment, as documented by Schreft and Singh (2003). Figure 2 further shows that the rise in production worker hours in the temporary help industry occurs concurrent with a decline in production worker hours in manufacturing. While the evidence in Figure 2 is only suggestive, it is consistent with the findings of Dey, Houseman, and Polivka (2007), who use occupational data to document a steady rise in the amount of production and other manual occupations in the temporary help industry. Finally, as Davis et al. (2008) document, there has been a secular decline in the flow of workers into unemployment. Thus, there is at least suggestive evidence of a structural shift in the labor market in how firms adjust their labor. It will be precisely these types of changes in the structural response of labor to shocks that motivate my structural VAR analysis below.

3. Job Flows and Jobless Recoveries

Evidence on gross job flows can greatly enhance our understanding of the last two jobless recovery periods. I use data from the relatively new Business Employment Dynamics (BED) data of the Bureau of Labor Statistics (BLS) to study job creation and job destruction patterns during the 1990-2006 period. Before describing these patterns, it is useful to describe the measures and data used. Throughout this paper, I measure job

creation and job destruction based on the concepts employed by the BLS and Davis, Haltiwanger, and Schuh (1996). *Job creation* is the sum of all jobs gained at expanding and opening establishments, and *job destruction* is the sum of all jobs lost at contracting and closing establishments. While these are collectively referred to as flow measures, they actually reflect a summation of net employment changes at the establishment level, measured between the third months of each quarter. Consequently, an establishment that has an equal number of hires and separations over a quarter will not have any jobs counted as either created or destroyed. *Net employment growth* (aggregated across establishments) is simply the difference between job creation and job destruction. *Job reallocation*, a measure of the cross-sectional dispersion of employment changes, is the sum of job creation and job destruction. Throughout the paper, I express these measures as a rate by dividing each by the average of the current and previous period's employment level.⁶

The BED is a longitudinal set of establishment data built from state unemployment insurance (UI) agency records. The data are quarterly, and since they are based on administrative records, they are a virtual census of establishments.⁷ The BLS publishes statistics on job creation and job destruction by industry, geography, and establishment size for all private establishments but only does so from the third quarter of 1992 forward. Administrative data exist back to 1990, but changes in how the administrative data were collected, particularly in how multi-establishment firms

⁶ Davis, Haltiwanger, and Schuh (1996) show that this produces a symmetric growth rate that is bounded between -200 and 200 percent. It also produces a rate that allows a symmetric treatment of employment changes at opening and closing establishments.

⁷ The self-employed and private households are the primary exclusions from the BED. For more on the BED data, its scope, its creation, and its measurement of job flows, see Pivetz, Searson, and Spletzer (2001) and Spletzer et al. (2004).

reported, made it impossible to link records across quarters using the standard BLS linkage algorithm. Consequently, one would overstate job flows estimated from the earlier data because of a large upward bias in the number of opening and closing establishments. If one were to go beyond the standard BLS algorithm, however, the nature of the reporting change makes it possible to identify many of the missed longitudinal linkages that it caused. For example, one can identify the continuation of many large establishments by identifying the simultaneous “closing” and “opening” of establishments within a narrow criteria of characteristics (e.g., county and four-digit industry) and above a particular employment threshold. I employ such a technique to identify missing linkages in the early administrative data, and I detail my methodology in the Appendix.

My extended BED data series spans the second quarter of 1990 through the third quarter of 2006. I calculate job flow estimates for all three-digit NAICS industries in the private sector.⁸ The data range in coverage from 5.0 million establishments representing 89.3 million employees in March 1990 to 6.7 million establishments representing 114.4 million employees in September 2006. On average, establishment expansions and contractions at continuous establishments make up about 80 percent of quarterly job creation and destruction, with the remainder accounted for by employment changes at opening and closing establishments.

Figure 3 shows the estimated job creation and job destruction rates for the private sector from 1990 to 2006. The benefit of the extended BED series is clear as it depicts

⁸ I do not create more disaggregated calculations for disclosure reasons. For these same reasons, I also combine Rail Transportation (NAICS 482) with Transit and Ground Passenger Transportation (NAICS 485) and Central Banks (NAICS 521) with Credit Intermediation (NAICS 522). Note that three-digit NAICS industries are generally comparable to two-digit SIC industries, with a major difference being a greater disaggregation of the service sector.

sharp movements in job destruction during the 1990-91 recession. This recession, which the NBER dates as lasting from 1990:4 to 1991:1, is followed by negative or slow employment through 1992:3. Throughout the recession and this “jobless recovery”, there is little notable change in the job creation rate, which hovers between 7.9 and 8.3 percent. The job destruction rate, however, spikes up sharply to 9.6 percent during the recession. It falls almost as quickly but remains relatively high (above 8 percent) for several quarters after the recovery. During the expansion from 1992 to 2000, the job creation and destruction rates deviate little from about 8.2 percent and 7.6 percent, respectively. The NBER dates the next recession as lasting from 2001:2 through 2001:4, though job creation starts to fall and job destruction begins to rise as early as 1999:4, over a year earlier. During the 2001 recession, job destruction again spikes upward but only rises to about 8.2 percent, while job creation continues to fall, to a low of 7.2 percent. Following the recession, a second jobless recovery occurs, though this time, job destruction returns to its pre-recession levels by early 2002 (and falls even further afterward), while job creation, following a brief recovery, continues to fall until it reaches a low of 7.0 percent in mid-2003. Net growth recovers by 2003:3 but with lower rates of job creation and job destruction.

Figure 3 illustrates clearly that, while both recessions experienced jobless recoveries characterized by prolonged employment losses, the underlying job flows during the subsequent recovery periods were quite different. The recovery during following the 1990-91 recession is characterized primarily by a flat job creation rate and a relatively slow decline in the job destruction rate, while the recovery following the 2001 recession is characterized by a large, persistent drop in job creation. Furthermore,

the decline in job creation during and after the 2001 recession runs counter to the conventional wisdom that cyclical employment fluctuations are driven primarily by movements in the job destruction rate.

While the comparison of job flows during the last two recessions is interesting in its own right, it is difficult to draw broader conclusions about the labor market from the findings. The data span only the last 16 years, so one cannot make comparisons to earlier periods without appealing to research that uses different data sources that usually cover only manufacturing. Research by Foote (1998) suggests that job flows in manufacturing may behave differently than in other industries. While the evidence in Figure 3 covers both manufacturing and nonmanufacturing industries, Foote's research still raises the concern that the job flow movements observed may be driven by only a few industries.

To address this issue, I pool job flow rates over this period for the 92 three-digit NAICS industries. I then regress the log of these rates on an industry fixed effect and an interaction of the fixed effect with a dummy, T_t , which is equal to 1 for one of the two downturns (i.e., either the quarters between 2001:1 and 2003:2 or the quarters between 1990:4 and 1992:3) and zero otherwise. I use the log of each job flow rate because there are wide variations in average job flow rates across industries. I run separate regressions for job creation and job destruction and for the 1990-92 and the 2001-03 downturns (for a total of 4 regressions). Formally, each regression is

$$(1) \quad \ln J_{it} = \mu_i + \delta_i T_t + e_{it},$$

where J_{it} represents either the job creation or the job destruction rate. Given this specification, the δ_i coefficients estimate the log point deviation in the job flow rate from its mean for industry i during the downturn period. Secular changes in the trends of job

creation or job destruction can bias the estimates of δ_i , however. To account for this, I also run each regression using the (log) deviation of each job flow rate from its industry-specific HP-filtered trend as the dependent variable.⁹

Figure 4 reports the results of this exercise. It presents the employment-weighted distribution of the δ_i coefficients across industries for each regression. The leftmost panels show the coefficient distributions for the unfiltered data and the rightmost panels show the coefficient distributions for the HP-filtered data. In each panel, the horizontal axis represents the log point change in each job flow rate. The top panels present the results for job creation. In the unfiltered data, there is a clear difference in the distribution of changes in job creation across industries. During the 2001-03 downturn, most industries had a decline in their job creation rate, while during 1990-92 downturn, most industries actually saw their job creation rates increase. The HP-filtered results show a much smaller difference between the distributions during the two downturns, as well as distributions that are more concentrated, but a clear difference in job creation patterns between the two periods remains. The results show that job creation in most industries did not deviate from its trend during the 1990-92 period, while job creation in most industries was below trend during the 2001-03 period. The distribution of changes during the 2001-03 period also exhibited greater dispersion, with the vast majority of industries showing a decline. The lower panels show the results for job destruction. In the unfiltered data, the 1990-92 downturn has a wide distribution of increases in the job destruction rate across nearly all industries. In contrast, the 2001-03 downturn has a wide distribution of increases *and* decreases in job destruction across industries, as well as a fat right tail,

⁹ For the HP filtering, I use a low-frequency smoothing parameter of $\lambda = 10^5$.

representing very large increases in job destruction concentrated within a handful of industries. Again, the results differ considerably when estimated relative to an HP trend. The majority of industries are concentrated together, showing a small increase in job destruction during the 1990-92 downturn. Industries during the 2001-03 downturn again have dispersed changes in job destruction relative to trend, though most show relative increases in job destruction.

Overall, the results suggest that the large, persistent decline observed in job creation during the 2001-03 downturn occurs across a wide variety of industries, even after accounting for a declining secular trend in job flows. This contrasts with the behavior of job creation during the 1990-92 downturn, when most industries experienced no change relative to trend. The persistently high rates of job destruction during the 1990-92 downturn also occur across nearly all industries. If anything, large increases in job destruction are concentrated within a few industries in the later downturn. Thus, it appears that the job flow patterns observed in Figure 3 are not driven by only a few industries.

4. Job Flows and the Great Moderation

4.A. Overview

The job flow evidence during the last two downturns, while interesting in its own right, is hard to put in a larger context of employment fluctuations over the business cycle because of the data's relatively short time series. Much attention has been paid recently to the decline in aggregate volatility, commonly referred to as the Great Moderation, that began in the mid-1980s. The job flow evidence during the jobless recoveries could reflect a potential shift in cyclical behavior. A natural question would be to ask whether the two

phenomena are related. In principle, jobless recoveries could have come about through a change in aggregate shocks, a change in employment's response to these shocks, or some combination of the two.

The ability of gross flow data to depict the underlying dynamics of employment fluctuations makes them well-suited to address this question. Unfortunately, it is impossible to do so without a longer time series than the BED data afford, and comparable data for the U.S. private-sector do not exist for earlier years. Comparable job flow data do exist, however, *for manufacturing*. Davis, Haltiwanger, and Schuh (1996) have tabulated estimates from the Longitudinal Research Database (LRD) of the Census Bureau using data that go back to 1972. In addition, Davis and Haltiwanger (1999) have combined these data with job flow estimates derived from the now-discontinued BLS Labor Turnover Survey (LTS). Their merged LRD-LTS data have manufacturing job flow estimates from 1947 through 1993.

While there are some methodological and conceptual differences between their merged series and the BED, the two essentially measure the same thing. So long as I account for their measurement differences, I can use the two data sources to create a time series that spans both the earlier and later periods. I do so using a methodology similar to the one used by Davis and Haltiwanger to merge the LRD and LTS data. They use the period where the two data series overlap to estimate predicted job flow estimates for the earlier data that are consistent with the job flow estimates from the later data. My approach differs in that I use a GMM approach that produces predicted job flow estimates for the earlier series that reconcile measurement differences between the two data sets while preserving key moments of the data.

4.B. Constructing a Longer Time Series of Job Flows

The GMM approach merges the manufacturing job flow estimates from the BED from 1990:2 through 2006:3 to the LRD-LTS job flow estimates produced by Davis and Haltiwanger for 1947:1 through 1998:4.¹⁰ Let C_t and D_t represent the job creation and job destruction rate estimates, respectively, at time t from the BED, and let POS_t , NEG_t , and NET_t represent the job creation, job destruction, and net growth rate estimates, respectively, from the LRD-LTS data. Finally, let NET_t^C be the net growth rate calculated from the Current Employment Statistics (CES) survey. I use the CES net growth rate as part of the splicing procedure because it provides a consistent measure of growth that covers the early and later periods. In addition, the CES draws its sample from the same administrative data used to create the BED, so the two data sources have nearly identical growth rates by construction.

The GMM estimation produces predicted estimates of the job creation and job destruction rates for the 1947:1-1990:1 period using the above data and parameters estimated by matching several key moments of the data. The predicted job flow rates come from the following specifications:

$$(2) \quad JC_t = \alpha_0 + \alpha_1 POS_t + \alpha_2 NET_t^C + u_t^c, \text{ and}$$

$$(3) \quad JD_t = \beta_0 + \beta_1 NEG_t + \beta_2 NET_t^C + u_t^d,$$

where the α 's and β 's are parameters estimated through GMM. The estimation uses an overidentified model of seven moments to estimate the six parameters. The fact that the most important moments (from a cyclical standpoint) are nonlinear, and often

¹⁰ I am grateful to Steve Davis and John Haltiwanger for providing an updated version of the LRD-LTS data and their suggestions for this splicing exercise.

interrelated, representations of the job creation and job destruction rates makes an overidentified model necessary.

I list the moments used, the period they cover, and the data source they come from in the top panel of Table 1. I chose these moments to preserve the cyclicity and comovement of job flows in the earlier period while ensuring that the estimated measures of magnitude and volatility are consistent across the early and later periods. Note that the excess reallocation rate, XR_t , is the sum of the job creation rate and the job destruction rate minus the absolute value of the net growth rate. Like the job reallocation rate, it is a measure of the cross-sectional dispersion of employment fluctuations, but unlike the job reallocation rate, it adjusts for reallocation due to changes in aggregate growth.¹¹

The model minimizes the distance between the actual values of these moments and the values predicted from the GMM estimates,

$$\min_{\theta} (M^s(\theta) - M^d)' W (M^s(\theta) - M^d).$$

In this specification, θ is the vector of α and β parameters from equations (2) and (3), $M^s(\theta)$ is the vector of estimated moments, and M^d is the vector of actual moments. I use the identity matrix for the weighting matrix, W . The estimation produces a unique solution for the estimated $\hat{\theta}$, and the resulting job creation and destruction series are robust to variations on the moments chosen, whether they be changes in the periods

¹¹ One might argue for using a variance of growth that allows for a structural break in the mid-1980s in the volatility of employment growth, but imposing such a break would impose a structure on the job flow estimates without any evidence *a priori* to suggest that the job flow data should exhibit such a break. As Figure 8 below and Figure A.1 in the appendix show, the job flow estimates do show a drop in volatility during the mid-1980s despite the fact that the spliced estimates in the earlier period exhibit slightly less volatility than the original LRD-LTS data.

covered or in the moments themselves.¹² I report the resulting parameter estimates in the bottom panel of Table 1. The final spliced series uses the predicted job flow estimates, based on equations (2) and (3), for 1947:1 – 1990:1 and uses the job flow estimates from the BED manufacturing data from 1990:2 forward. Figure A.1 in the appendix compares the final spliced series of job creation and job destruction to the original LRD-LTS estimates.

4.C. Evidence

Figure 5 presents the full time series of the job creation and job destruction rates for the spliced manufacturing data. Five findings immediately stand out in the figure. First, job destruction spikes upward in every recession, though the magnitude of these spikes declines with each cycle. Second, job creation appears to rebound sharply following all but the last two recessions. Third, there appears to be a declining secular trend in the magnitude of both job creation and job destruction. Fourth, the cyclical changes in job creation and job destruction appear to become less volatile over time. Finally, although job destruction remains high well after the end of the 2001 recession, job creation exhibits the same large, persistent decline observed for the private sector in Figure 3. In fact, job creation rates remain low well after the 2001 recession, and are the lowest observed over the entire 60-year period.

¹² I check robustness, in terms of how well the estimated job flow series preserve the cyclical patterns of the original series, using several variations on the chosen moments. These variations include moments calculated only from the LRD portion of the earlier series (i.e., 1972-1998), moments calculated only from the overlapping period (i.e., 1990-98), and an exactly identified model of six moments that excludes the variance of excess reallocation. The exactly identified model fails to produce a unique solution because of the highly nonlinear relations between moments. The other checks have unique solutions but produce job flow estimates that deviate dramatically from the original series. In contrast, the estimates from the specification I use closely match the behavior of the original job flow series.

Figures 6 through 8 explore the changes in magnitude and volatility further. Figure 6 shows the excess reallocation rate (defined above) for the spliced manufacturing series. The figure depicts the estimated rate as well as its HP trend.¹³ The trend highlights a secular decline in the magnitude of the reallocation rate. This decline is consistent with the decline in firm-level volatility found by Davis, Haltiwanger, Jarmin, and Miranda (2006). This decline is also notable because i) it appears to begin in the early 1960s, ii) it exhibits a trend, not a break, and iii) it is large—the excess reallocation rate declines by about one-third between 1961 and 2006. Remember that the reallocation measures can be thought of as the cross-sectional dispersion of employment fluctuations at a point in time. Note that this declining trend contradicts most of the previous research on the Great Moderation, which generally finds a break in aggregate volatility during the mid-1980s (Blanchard and Simon, 2001, are a notable exception).

Figures 7 and 8 show that this contrast with previous research is mostly due to the measure of volatility used. Figure 7 depicts the net growth rate in manufacturing over the postwar period. Consistent with the earlier research, employment growth has a notable decline in its volatility around 1984. Figure 8 explores the time series of volatility in more detail. The figure depicts the 10-year, centered, rolling standard deviations of job creation and job destruction. These measures also show a notably sharp decline, particularly for job destruction, in the mid-1980s. Other notable patterns also appear in Figure 8. The volatility of job flows falls steadily from the beginning of the sample period through the 1960s. The volatilities *increase* during the 1970s, with a sharp increase in the volatility of job destruction and a modest increase in the volatility of job

¹³ The HP trend uses a smoothing parameter of $\lambda = 10^5$.

creation. When the volatilities decline during the 1980s, job destruction has a much sharper decline than job creation.

Table 2 quantifies the changes in the magnitude, volatility, and comovement of job creation in manufacturing over the postwar period. I split the sample into three periods – 1947-1966, 1967-1983, and 1984-2006 – based on when significant changes in job flow volatility appear in the data. The table reports the mean and standard deviations for the job flow rates, the relative volatility of job destruction to job creation, and the correlations between job creation and job destruction, and job reallocation and net growth. For the latter two statistics, I report results for both the unfiltered and the HP-filtered data, since both the cyclical and secular job flow behavior appear to exhibit changes during the Great Moderation period. The statistics show that the magnitudes of job creation and destruction have fallen considerably throughout the postwar period. The standard deviations of nearly all variables have also fallen throughout the period, but to different degrees. As was evident in Figure 8, the standard deviation of job destruction fell further than the standard deviation of job creation. In fact, the standard deviation of job creation barely changes between the 1967-83 and 1984-2006 periods, but the standard deviation of job destruction falls by one-third. Consequently, the relative volatility declines dramatically, from 2.11 to 1.29 (this comes after an increase in the relative volatility between the 1947-66 and 1967-83 periods). The HP-filtered data, however, show essentially no change in the relative volatility. In the HP-filtered data, the volatilities of job creation and job destruction fall more dramatically than in the unfiltered data but do so in proportional amounts. The standard deviation of net growth also falls, declining in both cases by over 40 percent.

Changes in the absolute and relative volatilities of job flows altered their comovement patterns. The earlier periods show strong negative contemporaneous correlations between job creation and job destruction, and between job reallocation and net growth. These results are consistent with much of the previous research done on job flows in manufacturing (e.g., Davis and Haltiwanger, 1990, 1992; Davis, Haltiwanger, and Schuh, 1996). The 1984-2006 period, however, shows notable changes in the comovement patterns. The correlation between job creation and job destruction weakens considerably, going from strongly negative to slightly positive in the unfiltered data. The correlation between net growth and job reallocation declines (in magnitude) considerably in the unfiltered data but remains essentially unchanged in the HP-filtered data.

The long time series on manufacturing job flows, like several other aggregate measures of real activity, shows a large decline in volatility, consistent with previous evidence on the Great Moderation. At the same time, the evidence suggests a structural shift in the behavior of job flows during the Great Moderation – a decline in their relative volatility and comovement coincides with the emergence of jobless recoveries driven by movements in both the job creation and job destruction margins. It is worth noting that this is consistent with previous evidence that suggests a structural change in labor market dynamics. For example, Galí and Gambetti (2007) note a shift in the relationship between hours and labor productivity. Figures 1 and 2 in this paper show a change in the use of temporary layoffs and temporary help. In the next section, I complement this suggestive evidence with a structural vector autoregression analysis to identify the underlying shocks and characterize the structural responses of job creation and job destruction to these shocks.

5. A Structural VAR Analysis of Manufacturing Job Flows

5.A. Overview

My structural vector autoregression (VAR) analysis is based on a specification by Davis and Haltiwanger (1999). While the formulation is the same, the goal of the analysis is different. Davis and Haltiwanger studied the extent to which allocative shocks, relative to aggregate shocks, play a role in driving cyclical employment fluctuations. Here, I am not concerned with the contribution of the shocks, but with how their size and contributions may have changed during the post-1984 period and with potential changes in the job flow responses to these structural shocks.

5.B. Characterization and Identification

The structural VAR postulates that movements in job creation and job destruction are the result of two types of structural shocks: *allocative shocks*, which alter the distribution of economic activity, and *aggregate shocks*, which affect the level of economic activity over the business cycle. These shocks cannot be immediately recovered from a reduced-form VAR specification. Instead, one must make identifying assumptions that restrict the structural parameters to produce a range of plausible results.

Let $\mathbf{Y}_t = [JC_t, JD_t]'$ be a vector of the job creation and job destruction rates observed in the data. Also, let $\varepsilon_t = [\varepsilon_{at}, \varepsilon_{st}]'$ be a vector containing the structural disturbances that drive the movements in job creation and job destruction, with ε_{at} representing the structural disturbances due to aggregate shocks and ε_{st} representing the structural disturbances due to allocative shocks. Assume that the relationship between the innovations and the observed outcomes has a linear moving average representation,

$$(4) \quad \mathbf{Y}_t = \mathbf{B}(L)\varepsilon_t, \quad \mathbf{B}(0) = \mathbf{B}_0,$$

where $\mathbf{B}(L)$ is an infinite-order matrix lag polynomial. Without loss of generality, one can normalize the diagonal elements of the contemporaneous response matrix, \mathbf{B}_0 , to one.

Equation (4) depicts the structural representation of job flows and the structural shocks. The VAR, however, estimates the parameters of the a reduced-form model,

$$(5) \quad \mathbf{Y}_t = \mathbf{D}(L)\eta_t, \quad \mathbf{D}(0) = \mathbf{I},$$

where $\mathbf{D}(L)$ is the infinite-order lag polynomial implied by the estimated coefficients of the reduced-form VAR and $\eta_t = [\eta_{ct}, \eta_{dt}]'$ is the vector of reduced-form innovations, with η_{ct} representing innovations to job creation and η_{dt} representing innovations to job destruction. The structural and reduced-form representations are related by the following,

$$\eta_t = \mathbf{B}_0\varepsilon_t, \quad \text{and} \quad \mathbf{B}(L) = \mathbf{D}(L)\mathbf{B}_0.$$

Therefore, full knowledge of \mathbf{B}_0 allows one to recover the structural estimates and innovations from the reduced-form VAR.

Full knowledge of \mathbf{B}_0 is impossible from the data alone, so one must restrict the range of its parameter values based on a set of plausible identifying assumptions. Let b_{ij} denote the element in the i^{th} row and the j^{th} column of \mathbf{B}_0 , where $i = c$ (job creation) or d (job destruction), and $j = a$ (aggregate innovation) or s (allocative innovation). Davis and Haltiwanger (1999) discuss in depth the relevant spectrum of possible short-run and long-run restrictions and the implications these restrictions have for the estimated role of allocative shocks. I focus on a relatively tight set of restrictions from their identification approach to produce a range of plausible estimates, rather than attempt to take a stand on a particular specification for \mathbf{B}_0 . My identifying assumptions are:

$$(i) \quad b_{da} \leq -1,$$

$$(ii) \quad 0 < b_{cs} \leq 1,$$

$$(iii) \quad \text{corr}(\varepsilon_{at}, \varepsilon_{st}) = 0.$$

The first assumption postulates that aggregate disturbances cause job creation and job destruction to move in opposite directions. The boundary case of $b_{da} = -1$ is one where job creation and job destruction have symmetric (though opposite) responses to an aggregate shock. Cases where b_{da} is more negative imply asymmetrically larger responses for job destruction. The second assumption states that allocative disturbances cause job creation and job destruction move in the same direction – given the definition of an allocative disturbance, this must be true almost by construction. The boundary case of $b_{cs} = 1$ is the one where job creation and job destruction have a symmetric response to an allocative disturbance. The final assumption imposes zero covariance between the aggregate and allocative disturbances. This is a common, though critical, assumption for the structural VAR identification.¹⁴

The zero covariance restriction produces the following mapping between the variances and covariances of the reduced-form and structural innovations:

$$\sigma_c^2 = \sigma_a^2 + b_{cs}^2 \sigma_s^2,$$

$$\sigma_d^2 = b_{da}^2 \sigma_a^2 + \sigma_s^2,$$

$$\sigma_{cd} = b_{da} \sigma_a^2 + b_{cs} \sigma_s^2.$$

These three equations present four unknowns: the structural parameters b_{cs} and b_{da} , and the variances of the structural innovations, σ_a^2 and σ_s^2 . While this leaves the system under-identified, it does present a one-to-one mapping between b_{cs} and b_{da} .

¹⁴ Davis and Haltiwanger (1999) discuss in detail the plausibility of the orthogonality condition, as well as the potential interpretation of the sources of these shocks.

$$(6) \quad b_{cs} = \frac{\sigma_{cd} - b_{da}\sigma_c^2}{\sigma_d^2 - b_{da}\sigma_{cd}}.$$

Using a given value of b_{da} and the covariance matrix of the reduced-form innovations, one can use equation (6) to calculate the corresponding value of b_{cs} . This fully characterizes \mathbf{B}_0 and subsequently allows one to recover the variances of the structural innovations and as well the structural coefficients of the model.

5.C. Estimation and Results

I split the spliced manufacturing series into two periods: an early period that covers 1947:1 – 1983:4 and a later period that covers 1984:1 – 2006:3.¹⁵ I then run the reduced-form VAR estimation separately on each series, using a lag length equal to four. Figure 9 illustrates the implications of the identifying assumptions. The top panel depicts the mapping between b_{da} and b_{cs} for the early and later periods. The relationships differ because each period has a different reduced-form covariance matrix associated with it. The early period shows a nonlinear, decreasing relationship between the two parameters, while the later period shows a more linear, decreasing relationship. For the case of symmetric responses to allocative shocks ($b_{cs} = 1$), $b_{da} = -1.85$ in the early period and -1.94 in the later period. The reduced-form estimates imply that the case of symmetric aggregate shocks ($b_{da} = -1$) violates assumption (ii) in the early data, so I use the limiting case of $b_{da} = -1.05$ instead (the maximum for which assumption (ii) is satisfied in the early data), and to be consistent, I apply it to both periods. As one can see from Figure 9, this small shift has a minimal effect on the results.

¹⁵ I have also run the VAR on the data after HP filtering and on the unfiltered data using a linear and quadratic time trend. All approaches produce results that are qualitatively very similar.

The bottom panel shows the estimated standard deviations of the aggregate and allocative innovations for the early and later periods as a function of b_{da} . The panel shows stark differences, in both absolute and relative terms, in the estimated standard deviations of the early and later periods. The differences *between* periods appear robust to the identifying assumption for b_{da} . Results *within* each period are more sensitive. In the earlier period, the estimated volatility of allocative disturbances increases substantially with b_{da} , but the estimated volatility of aggregate disturbances does not, so identification affects their relative volatility considerably. If one assumes symmetric job flow responses to allocative shocks, then aggregate disturbances are more volatile, and if one assumes symmetric responses to aggregate shocks, then allocative disturbances are more volatile. In the later period, the volatilities of aggregate and allocative disturbances are not particularly sensitive to the identifying assumptions, and in all cases, the allocative disturbances are relatively more volatile. Mechanically, the decline in the comovement of job creation and job destruction between the early and later periods (observed in Table 2) plays a role in the differences in sensitivity observed across the two periods. In the early period, the covariation of job creation and destruction is strongly negative. As one can see from equation (6), as σ_{cd} becomes larger in magnitude, the relationship between b_{cs} and b_{da} becomes more nonlinear, and consequently the estimated volatility of the structural innovations becomes more sensitive to the identifying assumptions. In the later period, the estimates are much less sensitive to identifying assumptions because the correlation between job creation and job destruction is weaker.

Table 3 quantifies the differences between the early and later periods under the limiting cases of identifying assumptions. The top panel lists the standard deviations of

the structural innovations under the assumption that job flows have symmetric responses to allocative shocks ($b_{cs} = 1$), while the bottom panel lists the standard deviations under the assumption that they have (nearly) symmetric responses to aggregate shocks ($b_{da} = -1.05$). Under the first assumption, the standard deviation of aggregate disturbances falls by 60 percent, while the standard deviation of allocative disturbances falls by 31 percent. This leads the relative volatility of allocative to aggregate disturbances to rise from 0.63 to 1.07. In other words, the later period experiences a paradigm shift from one where aggregate disturbances are more important for cyclical employment fluctuations to one where the two disturbances play a roughly equal role. Under the second assumption, the standard deviation of aggregate disturbances falls by 55 percent, comparable to the case where $b_{cs} = 1$. The standard deviation of allocative disturbances falls by 54 percent, more than in the previous case, so that the relative volatility of the two disturbances (which is greater than one in both periods) remains essentially unchanged.

The last exercise estimates the impulse response functions of the job creation and job destruction rates to aggregate and allocative shocks. For each period, I estimate the impulse responses using the same limiting cases of $b_{cs} = 1$ and $b_{da} = -1.05$ as before. The results are in Figure 10. The impulse responses are to a unit standard deviation decrease in the aggregate disturbance and a unit standard deviation increase in the allocative disturbance. I use a unit change to study the responses of job creation and job destruction without the known declines in the volatility of the structural innovations complicating the analysis.

The top two panels of Figure 10 show the results under the assumption of symmetric responses to allocative shocks, while the bottom two panels show the results

under the assumption of symmetric responses to aggregate shocks. The point of the exercise is to see whether the responses of job creation and job destruction differ at all between the earlier and later periods of the time series. The set of assumptions used matters for the size of the job flow responses but not for their qualitative behavior. In addition, the responses of job creation and job destruction to an allocative shock differ little between the two periods. The same is not true of their responses to an aggregate shock. In the early period, a negative aggregate shock increases job destruction and decreases job creation, with a smaller effect on job creation by assumption. Following the shock, job destruction falls and job creation rises. The rise in job destruction dissipates within 4-5 quarters (depending on the assumption used). The fall in job creation dissipates after 3-4 quarters. Job creation then continues to rise, showing a slight positive rebound after the shock. The magnitudes of these initial responses in the later period are notably similar to the magnitudes in the earlier period. The patterns of the subsequent quarters, however, differ considerably. The decline in the job destruction response is smoother and slower, staying above trend for 8 to 12 quarters, depending on the identifying assumption used. The rise in the job creation response is also smoother and slower. It, too, lasts longer: eight or more quarters, depending on assumptions. In addition, the job creation response in the later period shows no “rebound” effect in its response to the aggregate shock.¹⁶

Overall, the structural VAR suggests several reasons for the change in manufacturing job flow behavior during the 1984-2006 period. First, as one might have

¹⁶ To get a sense of the robustness of these differences, I present the job flow responses in with their standard error bands included in Figure A.2 of the Appendix. The figure uses the assumption of $b_{cs} = 1$. As one can see, the differences in job flow patterns, particularly those for job creation, are significantly different from each other in the two periods.

expected, the size of the structural shocks has declined. The volatility of aggregate shocks fell between 55 and 60 percent, and the volatility of allocative shocks fell between 31 and 54 percent. Second, depending on the identifying assumptions used, the relative importance of these shocks in driving job flow movements may have changed as well. Finally, regardless of the identifying assumptions, the results show a shift in the responses of job creation and job destruction to an aggregate shock. The responses are more persistent in the later period, and the response of job creation no longer exhibits a positive rebound following its initial decline. Thus, the observed changes in behavior are attributable to a large decline in the volatility of the underlying structural shocks *and* to a shift in the structural responses to those shocks. In addition, the shift in structural responses appears consistent with the job flow patterns observed during the last two economic downturns, including their jobless recoveries, in the sense that employment adjustments occur much slower and more smoothly in the later period. Thus, there seems to be a strong case for a link between the Great Moderation and jobless recoveries.

6. Conclusions

This paper studies the cyclical behavior of job creation and job destruction before and during the Great Moderation period, which most date as starting in the 1980s. Previous research shows that the aggregate volatility of output and other real economic variables declined sharply during this period. In addition, the two recessions that have occurred during this period are characterized by “jobless recoveries,” where stagnant employment growth continues well after the recovery of output growth.

I appeal to new and existing data on job creation and job creation to find that while the jobless recoveries have similar net employment growth patterns, their

underlying job flows are quite different. Over a longer horizon, I find that both the magnitude and volatility of job flows declined considerably during the postwar period; magnitude exhibited a secular decline that began in the 1960s, while volatility exhibited a sharp drop in the mid-1980s. Comovement patterns changed as well. A structural VAR analysis decomposes the job flow movements into those attributable to aggregate shocks and those attributable to allocative shocks. The analysis estimates the volatility of these structural shocks as well as the implied job flow responses separately for the pre-1984 and post-1984 periods. The analysis suggests that the volatility of both structural shocks decline in the later period. In addition, there is a notable shift in the response of job flows to an aggregate disturbance. The responses of both job creation and job destruction become more persistent over time. In addition, job creation shows no positive rebound after the initial shock, as it does in the early period. This shift in responses is generally consistent with the shift in job flow behavior observed during the jobless recoveries. It provides at least suggestive evidence that the jobless recoveries observed following the last two downturns are linked to the Great Moderation period through a less volatile and more persistent response of job flows to an aggregate shock. It also suggests that the change in behavior reflects a structural change in the nature of employment adjustment, such as a shift in relative importance between temporary layoffs and temporary help employment as an adjustment mechanism, and is not just the result of decreased aggregate volatility.

Appendix

Linking the 1990-92 BED Data

Implementation of the “Multiple Worksite Report” (MWR) to state administrative records in the early 1990s caused serious complications for the BED linkage process. The BLS implemented the MWR so that multi-establishment firms could easily report the employment and payroll of their separate establishments. Prior to the MWR, firms in many states reported their multiple establishments as a single record. The MWR changed that, but when it did so, its restructuring of administrative records (which involved breaking out the single records into their individual establishments) was not fully recorded by every state. Consequently, the BLS did not have the necessary information to link what were in reality continuous units. This created large overstatements of opening and closing establishments in the first quarter of 1991 and the second quarter of 1992.

I correct for these overstatements by using unique characteristics of the MWR implementation process. First, since the MWR implementation occurs at the state level, I only need to focus on the affected states. Second, firm identifying codes do not change during the MWR implementation, only the codes for the individual reporting units change (this is not necessarily true of other administrative changes). Third, since these changes are theoretically only changes in paperwork, there should be no movement of employment across industries or locations, which sometimes occurs in the data during corporate mergers and other account restructurings. Finally, the administrative data have a fine level of geographic and industry detail (county level, and either 4-digit SIC or 6-digit NAICS, respectively). Large employment fluctuations at these levels of detail are relatively rare and thus easily identifiable in the data.

Given these characteristics, I use a three-step process. The first step calculates job flows using the standard BED methodology. From this, I take the subset of establishments identified as openings or closings. The second step uses a grid search for openings and closings with identical firm identifiers by county and detailed industry. I assume that these records are the result of the MWR implementation and match them. In the cases where there are multiple openings and multiple closings within the same cell (as opposed to one closing and one or more openings), I match probabilistically based on the employment level of each record. The final step recognizes that, in practice, some new records will have different industry codes than their predecessor. It takes the remaining unlinked records with identical firm identifiers and attempts to match within counties only. This last step produces less than 10 percent of the total matches I identify.

My approach is not without risks. First, it may produce false matches of truly opening and closing establishments. I am not too concerned with this possibility, since the false match would have to occur among opening and closing establishments within both the same firm and the same county, an occurrence that is extremely rare. Second, I may miss links that occur either within firm accounts and across counties or across entirely different firm account identifiers. Without predecessor or successor record information, I cannot identify these matches without increasing the chances of a false match among other records, so some small potential for missed links remains.

Table A.1 lists the results of my matching strategy for the quarters of interest. The matches significantly reduce employment changes at opening and closing establishments and slightly increase employment changes at continuing establishments (newly matched records often have legitimate changes in employment during these quarters).

References

Aaronson, Daniel, Rissman, Ellen R. and Sullivan, Daniel G. "Can Sectoral Reallocation Explain the Jobless Recovery?" *Economic Perspectives*, Federal Reserve Bank of Chicago (Second Quarter 2004), 28(2), pp. 36-49.

Ahmed, Shaghil, Levin, Andrew, and Wilson, Beth Anne Wilson. "Recent U.S. Macroeconomic Stability: Good Practices, Good Policies, or Good Luck?" *Review of Economics and Statistics* (August 2004) 86(3), pp. 824-832.

Bachmann, Ruediger. "Understanding Jobless Recoveries: A Tale of Two Margins," Yale University, mimeo (January 2007).

Blanchard, Oliver J., and Simon, John "The Long and Large Decline in U.S. Output Volatility." *Brookings Papers on Economic Activity, Vol. 1* (2001), pp. 135-164.

Davis, Steven J., Faberman, R. Jason, Haltiwanger, John C., Jarmin, Ron, and Miranda, Javier, 2006. "Business Volatility, Job Destruction, and Unemployment," University of Chicago working paper (June 2008).

Davis, Steven J. and Haltiwanger, John C. "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," in Oliver Jean Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual 1990*. Cambridge, MA: MIT Press (1990), pp. 123-168.

_____. "Gross Job Creation, Gross Job Destruction and Employment Reallocation." *Quarterly Journal of Economics* (August 1992), 107(3), pp. 819-63.

_____. "On the Driving Forces behind Cyclical Movements in Employment and Job Reallocation." *American Economic Review* (December 1999), 89(5), pp. 1234-1258.

Davis, Steven J., Haltiwanger, John C., Jarmin, Ron, and Miranda, Javier. "Volatility Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., *2006 NBER Macroeconomics Annual*. Cambridge, MA: National Bureau for Economic Research, (2006), pp. 107-156.

Davis, Steven J., Haltiwanger, John C. and Schuh, Scott. *Job creation and destruction*. Cambridge, MA: MIT Press, (1996).

Dey, Matthew, Houseman, Susan, and Polivka, Anne. "Manufacturers Outsourcing to Employment Services," Upjohn Institute Working Paper No. 07-132 (2007).

Engemann, Kristie M., and Owyang, Michael T. "Whatever Happened to the Business Cycle? A Bayesian Analysis of Jobless Recoveries," Federal Reserve Bank of St. Louis Working Paper 2007-013A (March 2007).

Foote, Christopher. "Trend Employment Growth and the Bunching of Job Creation and Destruction," *Quarterly Journal of Economics* (August 1998), 113(3), pp. 809-834.

Galí, Jordi, and Gambetti, Luca. "On the Sources of the Great Moderation," Universitat Pompeu Fabra, mimeo (June 2007).

Groschen, Erica L., and Potter, Simon. "Has Structural Change Contributed to a Jobless Recovery?" *Current Issues in Economics and Finance* (August 2003), 9(8).

Hall, Robert E., "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years," in Mark Gertler and Kenneth Rogoff, eds., *2005 NBER Macroeconomics Annual*. Cambridge, MA: National Bureau for Economic Research, (2005), pp. 101-137.

Justiniano, Alejandro, and Primiceri, Giorgio E.. "The Time-Varying Volatility of Macroeconomic Fluctuations." Northwestern University, mimeo (February 2007).

Kim, Chiang-Jin, and Nelson, Charles R. "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle." *Review of Economics and Statistics* (November 1999), 81(4), pp. 608-616.

Koenders, Kathryn, and Rogerson, Richard. "Organizational Dynamics over the Business Cycle: A View on Jobless Recoveries," *Federal Reserve Bank of St. Louis Review* (July/August 2005), 87(4): 555-79

McConnell, Margaret M., and Perez-Quiros, Gabriel. "Output Fluctuations in the United States: What has Changed Since the Early 1980's?" *American Economic Review*, (December 2000), 90(5), pp. 1464-1476.

Pivetz, Timothy R.; Searson, Michael A. and Spletzer, James R. "Measuring Job Flows and the Life Cycle of Establishments with BLS Longitudinal Establishment Microdata," *Monthly Labor Review*, (April 2001), 124(4), pp. 13-20 .

Ramey, Valerie A., and Vine, Daniel J. "Declining Volatility in the U.S. Automobile Industry," *American Economic Review* (December 2006), 96(5), pp. 1876-1889.

Schreft, Stacey L., and Singh, Aarti. "A Closer Look at Jobless Recoveries," *Federal Reserve Bank of Kansas City Economic Review* (Second Quarter, 2003), 88(2), pp. 45-73.

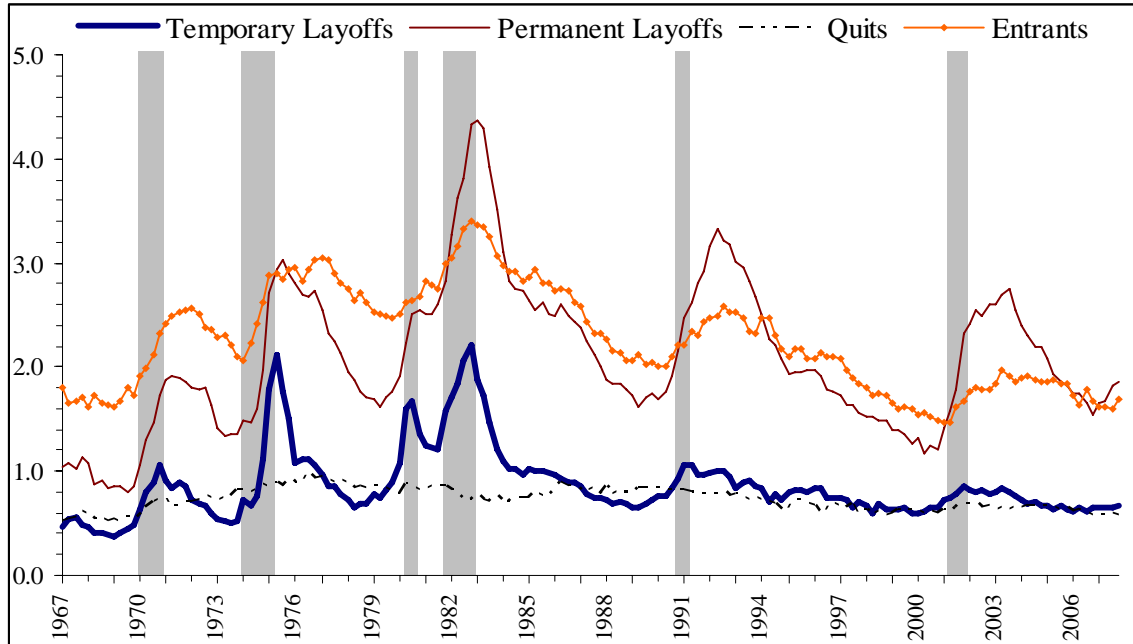
Shimer, Robert, "Reassessing the Ins and Outs of Unemployment," University of Chicago, mimeo (2007).

Spletzer, James R.; Faberman, R. Jason; Sadeghi, Akbar; Talan, David M. and Clayton, Richard L. "Business Employment Dynamics: New Data on Gross Job Gains and Losses," *Monthly Labor Review* (April 2004), 127(4), pp. 29-42.

Stiroh, Kevin J. “Volatility Accounting: A Production perspective on Increased Economic Stability,” Federal Reserve Bank of New York, mimeo (August 2007).

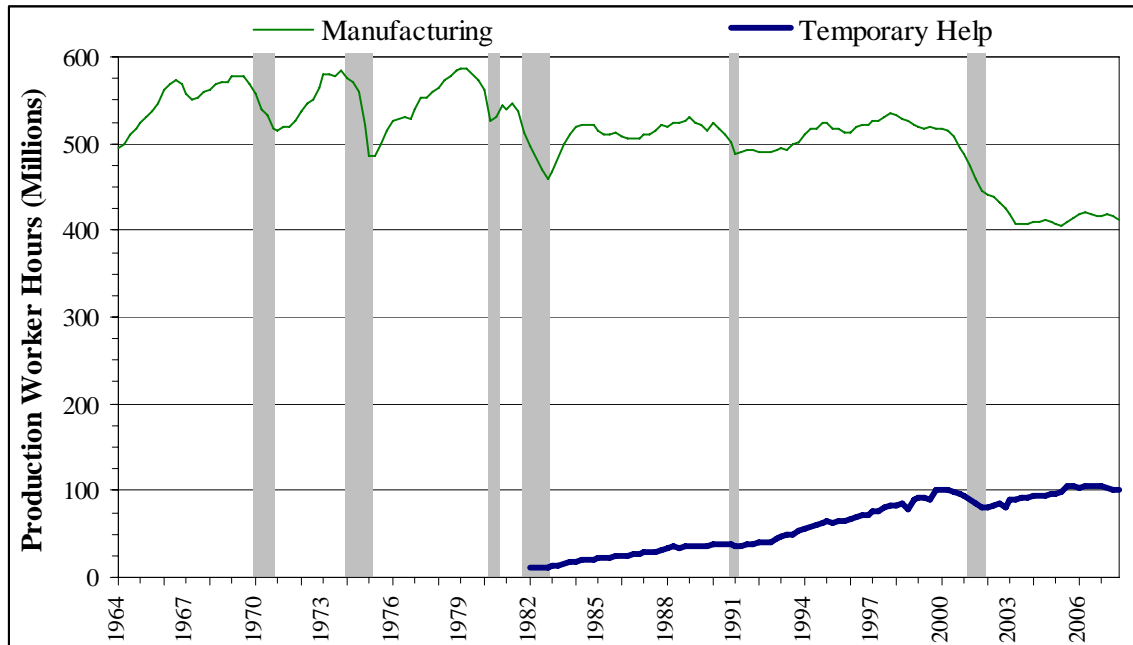
Stock, James H., and Watson, Mark W. “Has the Business Cycle Changed and Why?” in Mark Gertler and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2002*. Cambridge, MA: MIT Press (2002), pp. 159-218.

Figure 1. Unemployment Rates by Reason, CPS Data 1967-2007



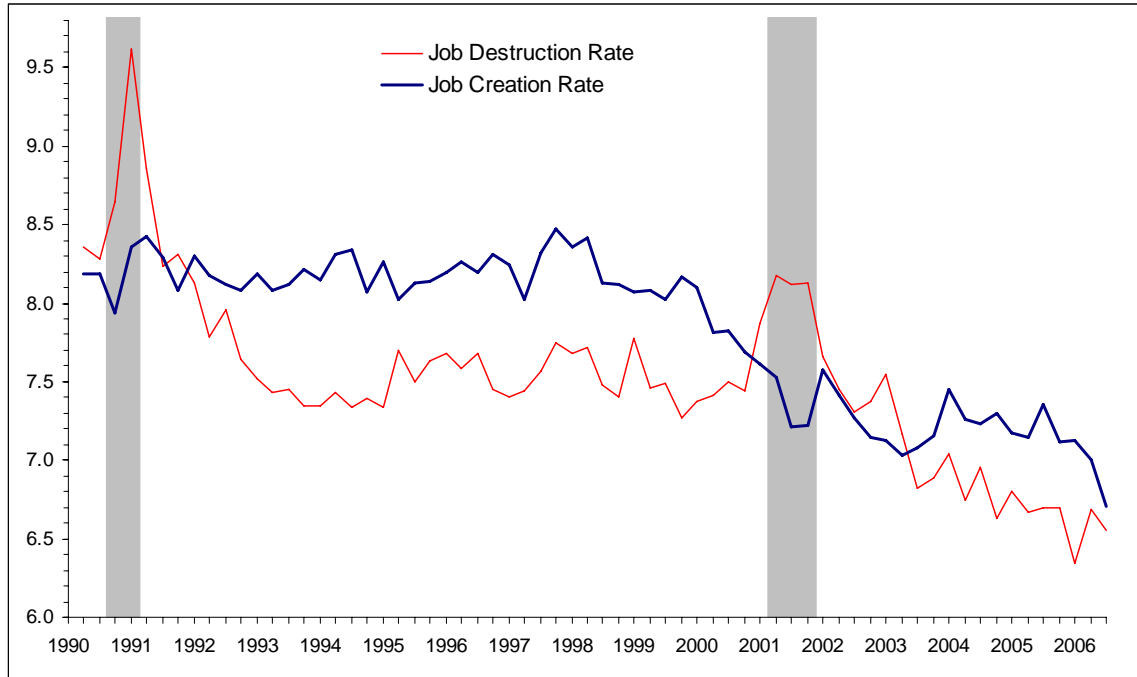
Notes: Figure depicts unemployment rates (as percentages of the labor force) by reason of unemployment. Quits represent voluntary separations into unemployment; entrants represent both new entrants and re-entrants into the labor force. Data come from the Current Population Survey. Shaded areas represent NBER-dated recessions.

Figure 2. Production Worker Hours in Manufacturing and Temporary Help, CES Data 1964-2007



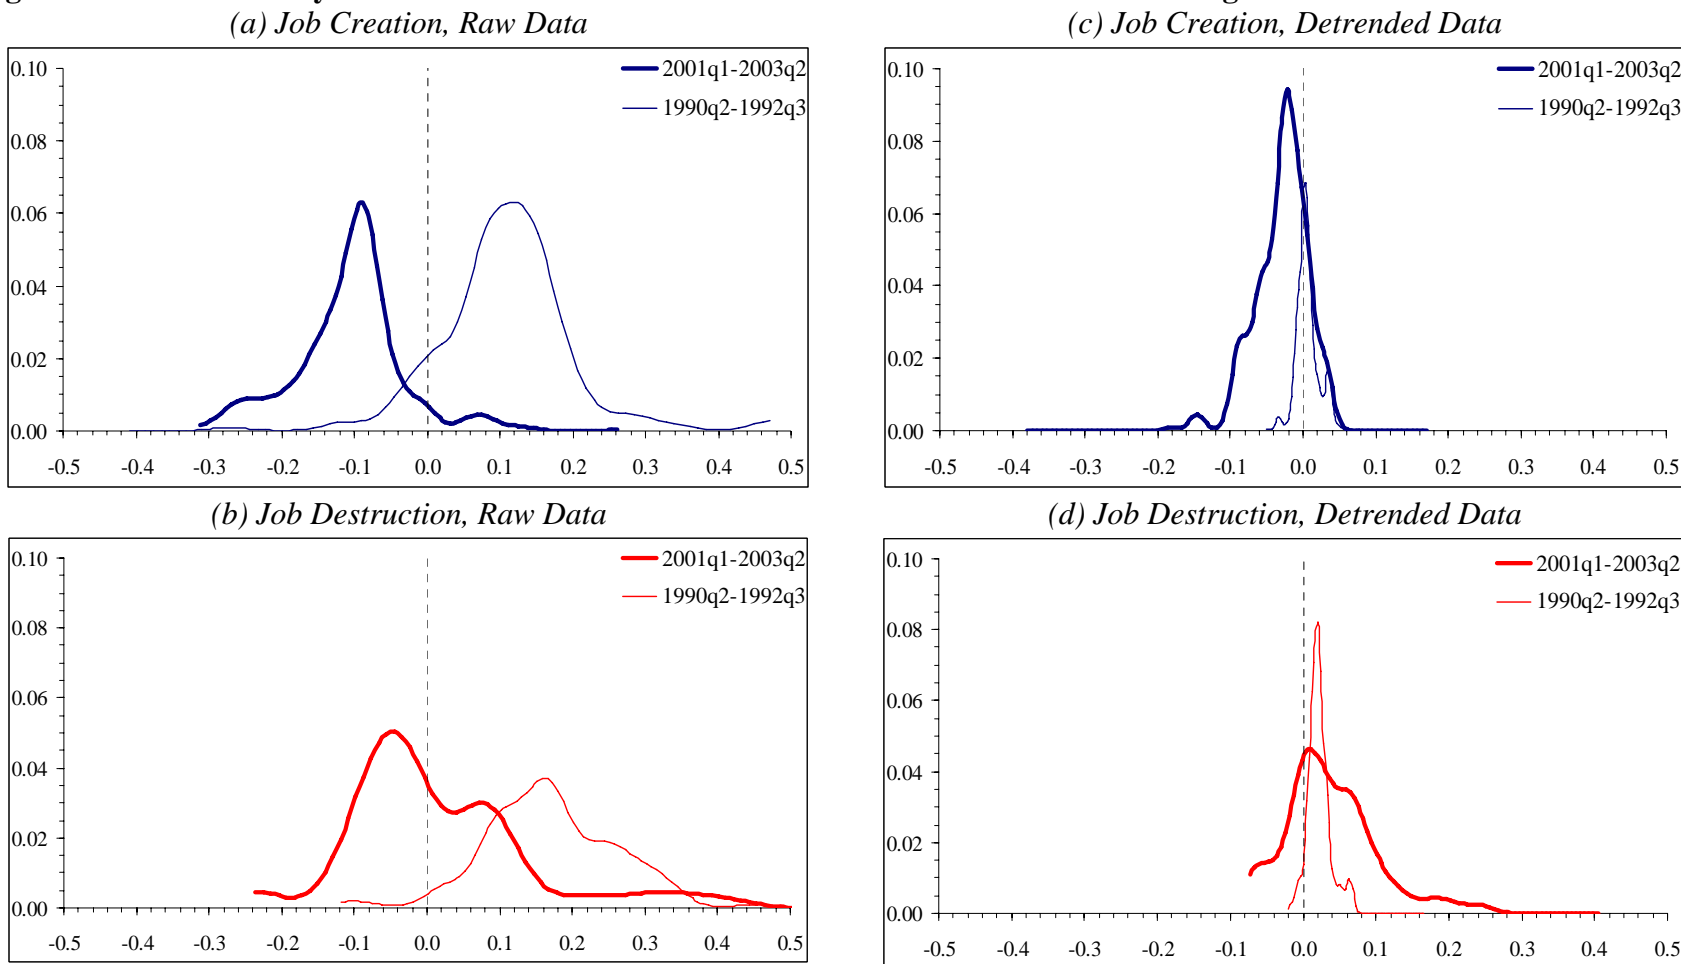
Notes: Figure depicts aggregate production worker hours (in millions) in the manufacturing and temporary help industries. Due to changes in industry classifications, the temporary help data encompass most of the employment services industry, which includes both temporary help and employee leasing agencies. Shaded areas represent NBER-dated recessions.

Figure 3. Job Creation and Job Destruction Rates, Private Sector



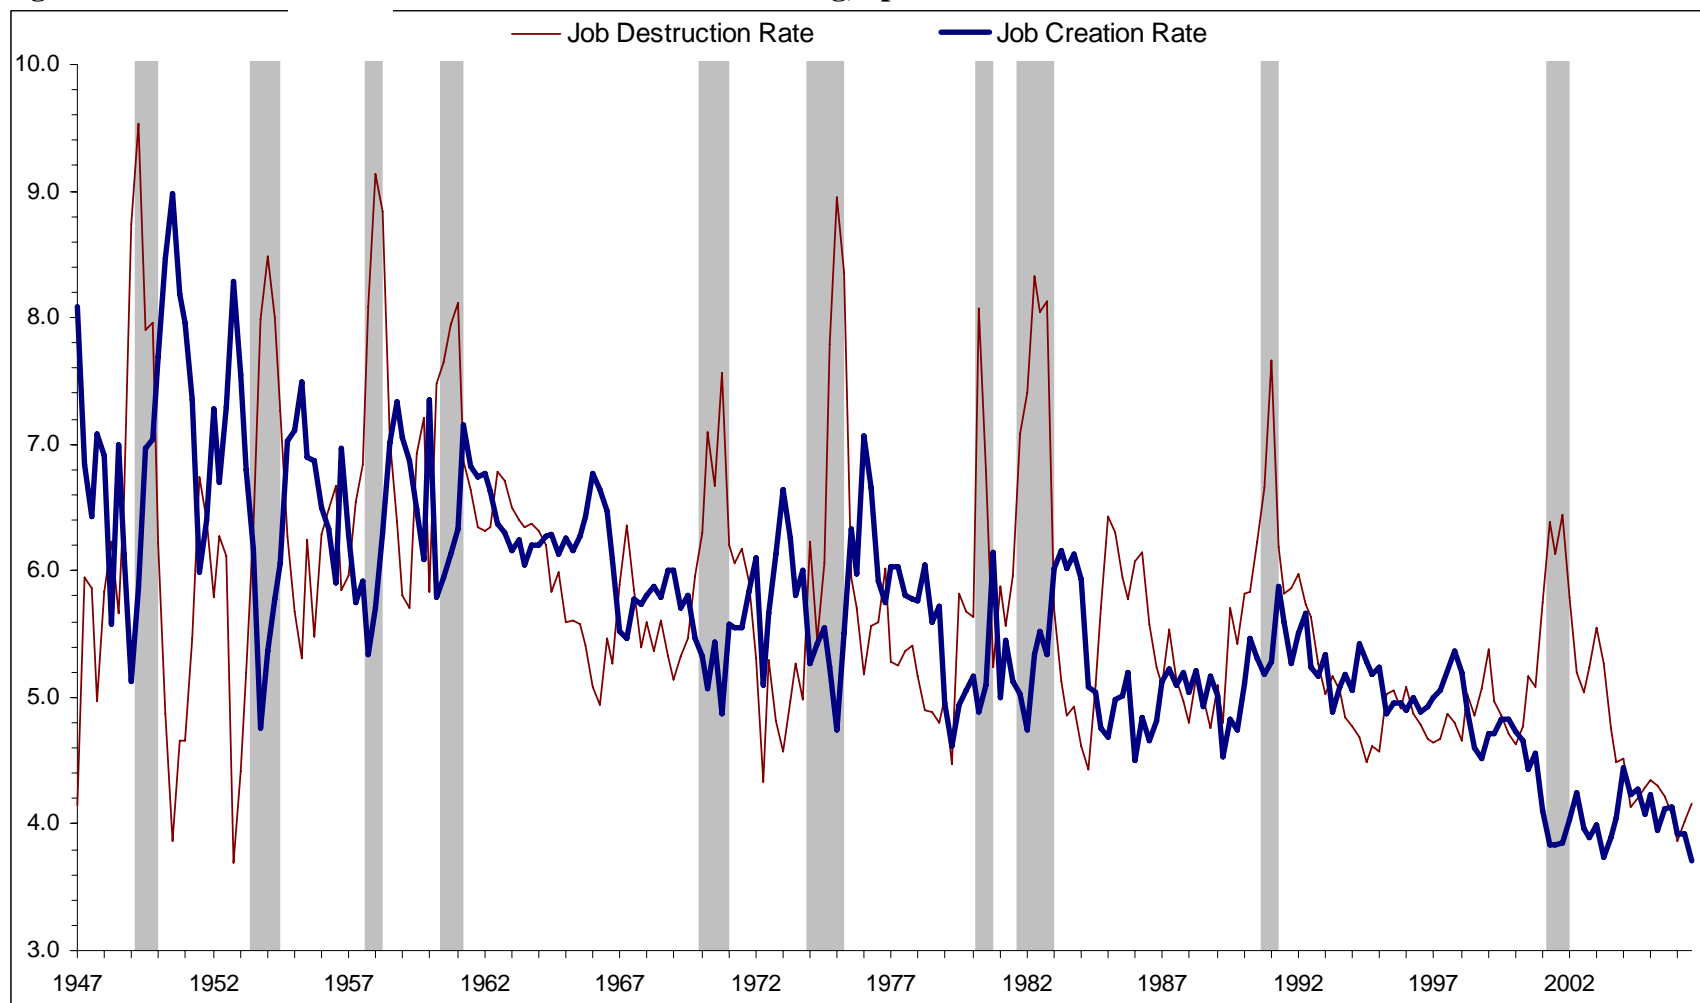
Notes: Job flow rates come from author's calculations using BED data. Shaded areas represent NBER-dated recessions.

Figure 4. Across-Industry Distribution of Shifts in Job Creation and Job Destruction During the Two Downturns



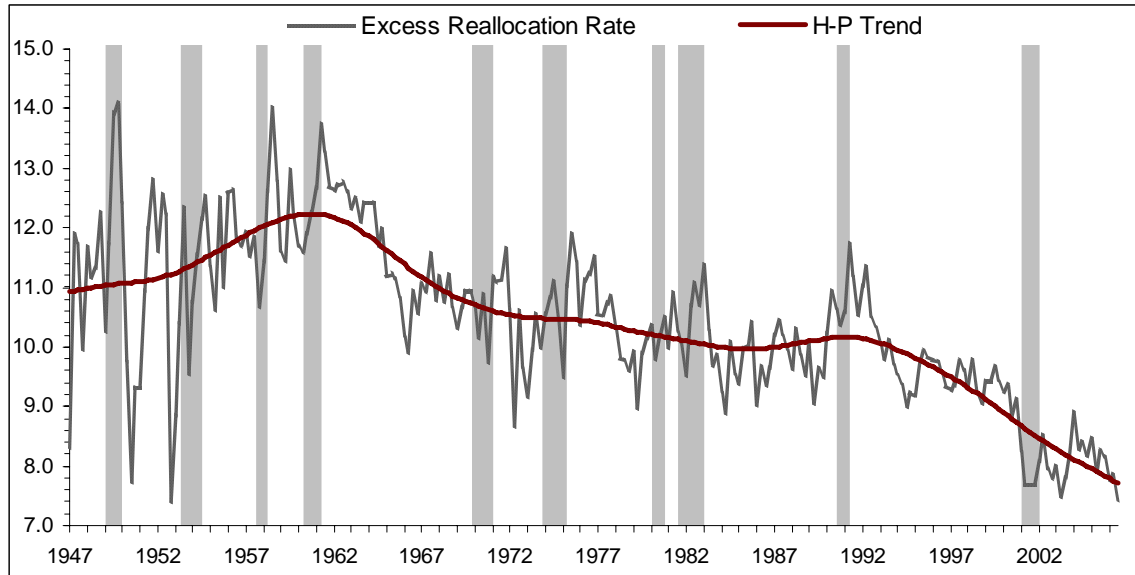
Notes: Panels illustrate kernel density estimates of the employment-weighted distributions of shifts in the average job creation or job destruction rates during the two noted downturns. Shifts are measured using the δ_i coefficients obtained from the regression in equation (1) in the text, using a panel of BED data for 90 3-digit NAICS industries. The shifts for the raw data are relative to the mean job creation or destruction rate for the sample. The shifts for the detrended data are relative to an industry-specific HP-filtered trend, each with a smoothing parameter of $\lambda = 1600$.

Figure 5. Job Creation and Job Destruction in Manufacturing, Spliced Series 1947 - 2006



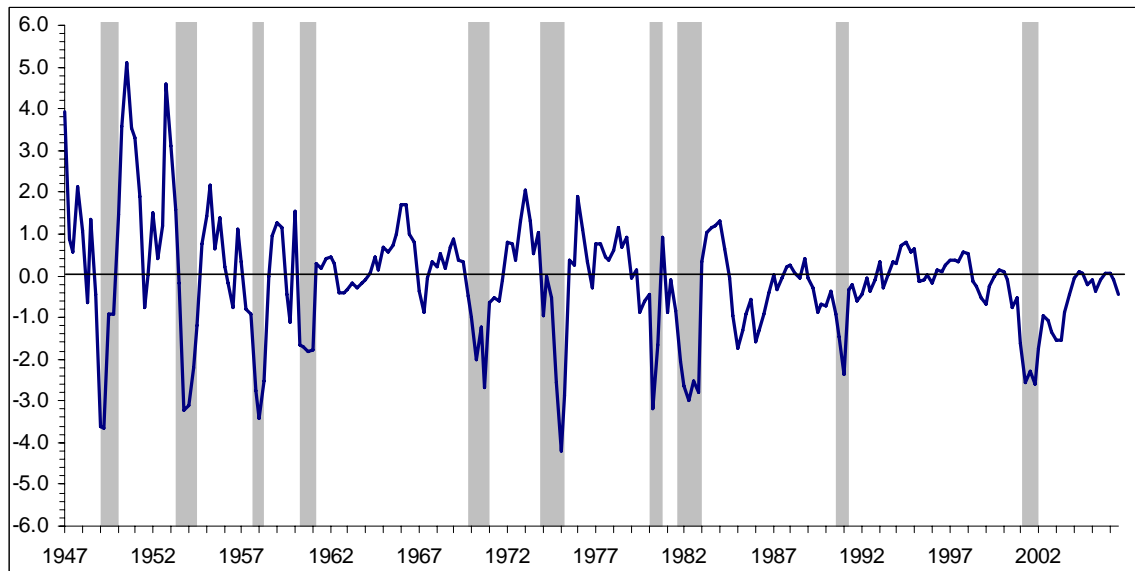
Notes: Estimates are a spliced series of job creation and job destruction rates based on a GMM estimation using BED data along with a merged series of LTS and LRD data from Davis and Haltiwanger (1999). See text for details. Shaded areas represent NBER-dated recessions.

Figure 6. Excess Reallocation in Manufacturing, Spliced Series, 1947 – 2006.



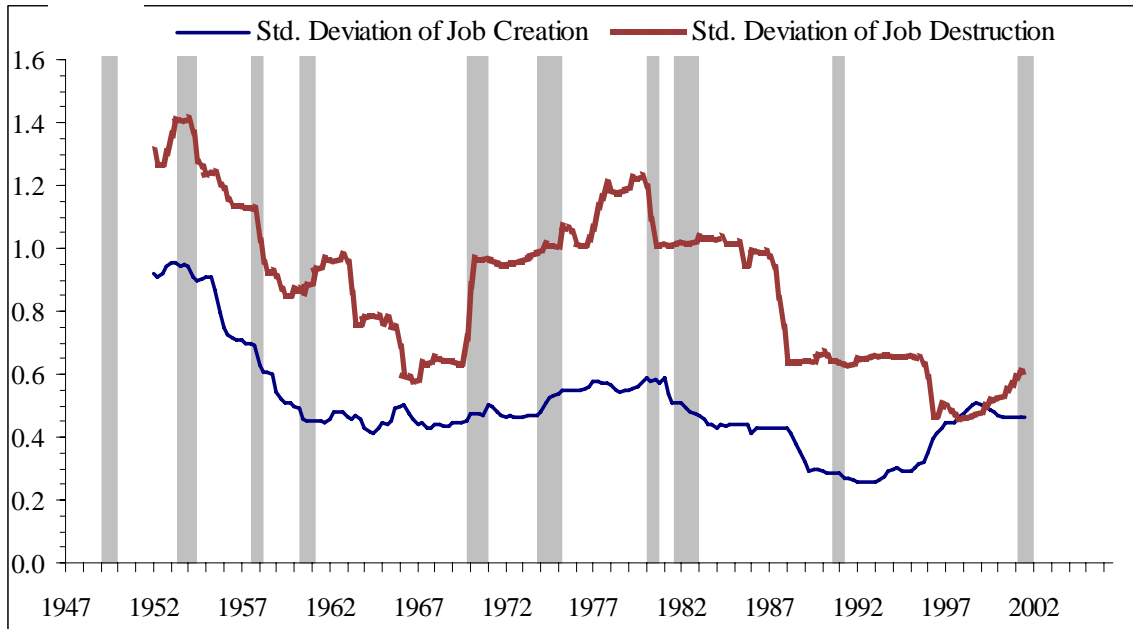
Notes: Excess reallocation estimates come from the spliced manufacturing series described in the text. The trend is HP-filtered with a smoothing parameter of $\lambda = 1,600$. Shaded areas represent NBER-dated recessions.

Figure 7. Net Employment Growth in Manufacturing, Spliced Series 1947 – 2006.



Notes: Net growth rate estimates come from the spliced manufacturing series described in the text. Shaded areas represent NBER-dated recessions.

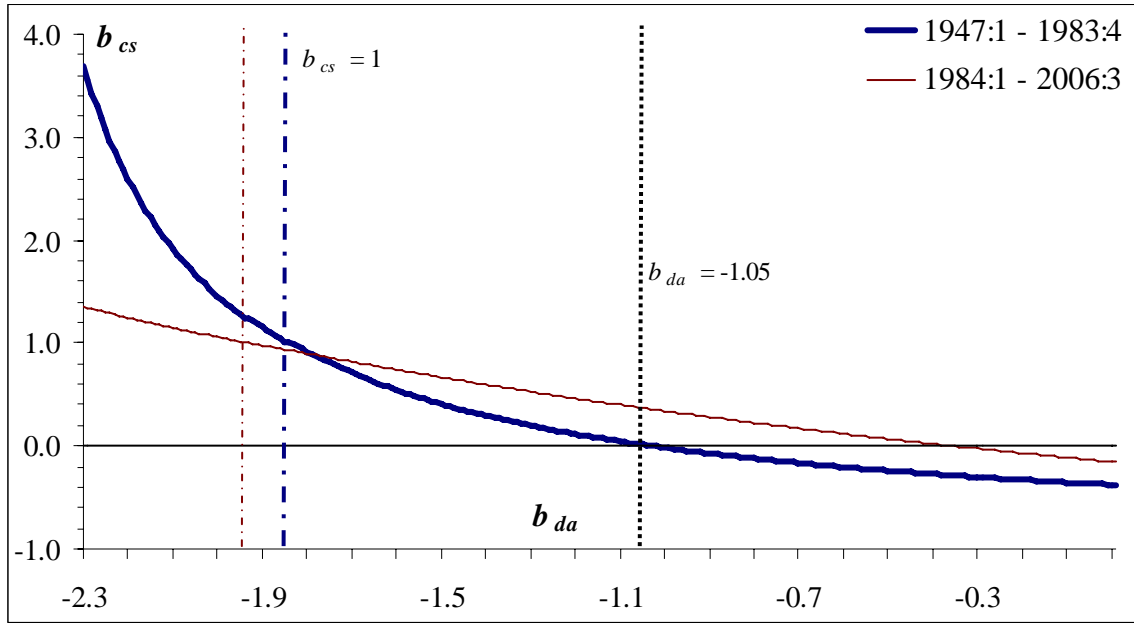
Figure 8. Job Flow Volatility in Manufacturing, Spliced Series, 1947 – 2006



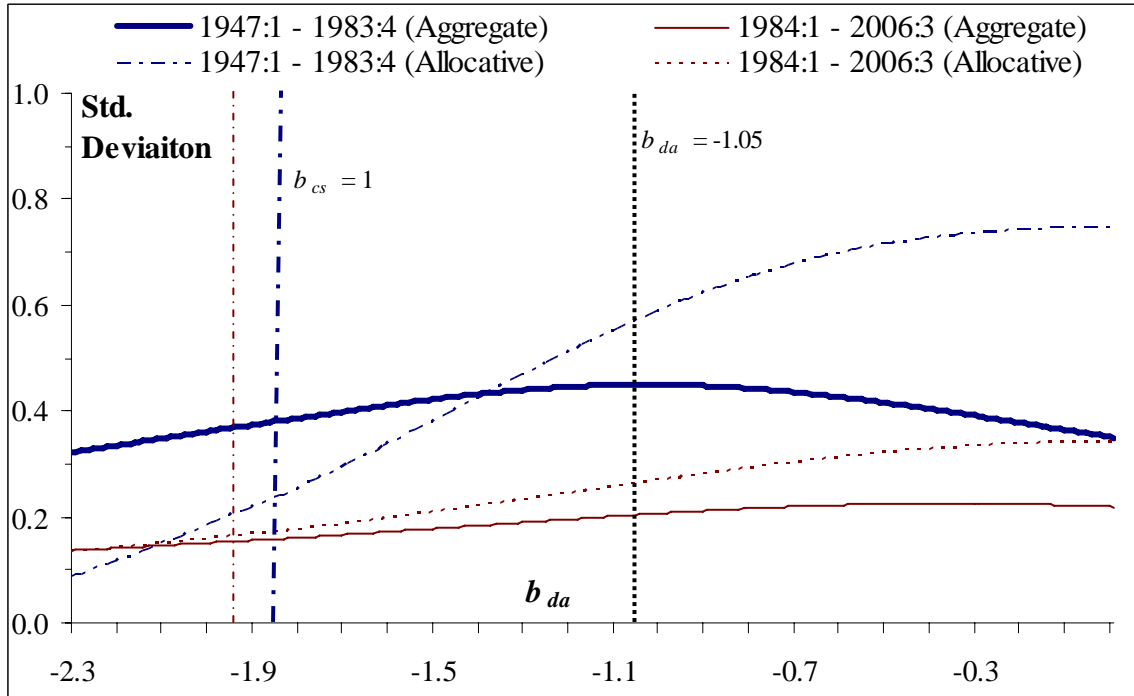
Notes: Estimates are quarterly rates and come from publicly available data from the Job Openings and Labor Turnover Survey (JOLTS). The shaded areas represent the NBER-dated recessions.

Figure 9. Structural VAR Identification Implications for Key Parameters

(a) *The Tradeoff between Identification Assumptions on b_{da} , b_{cs}*

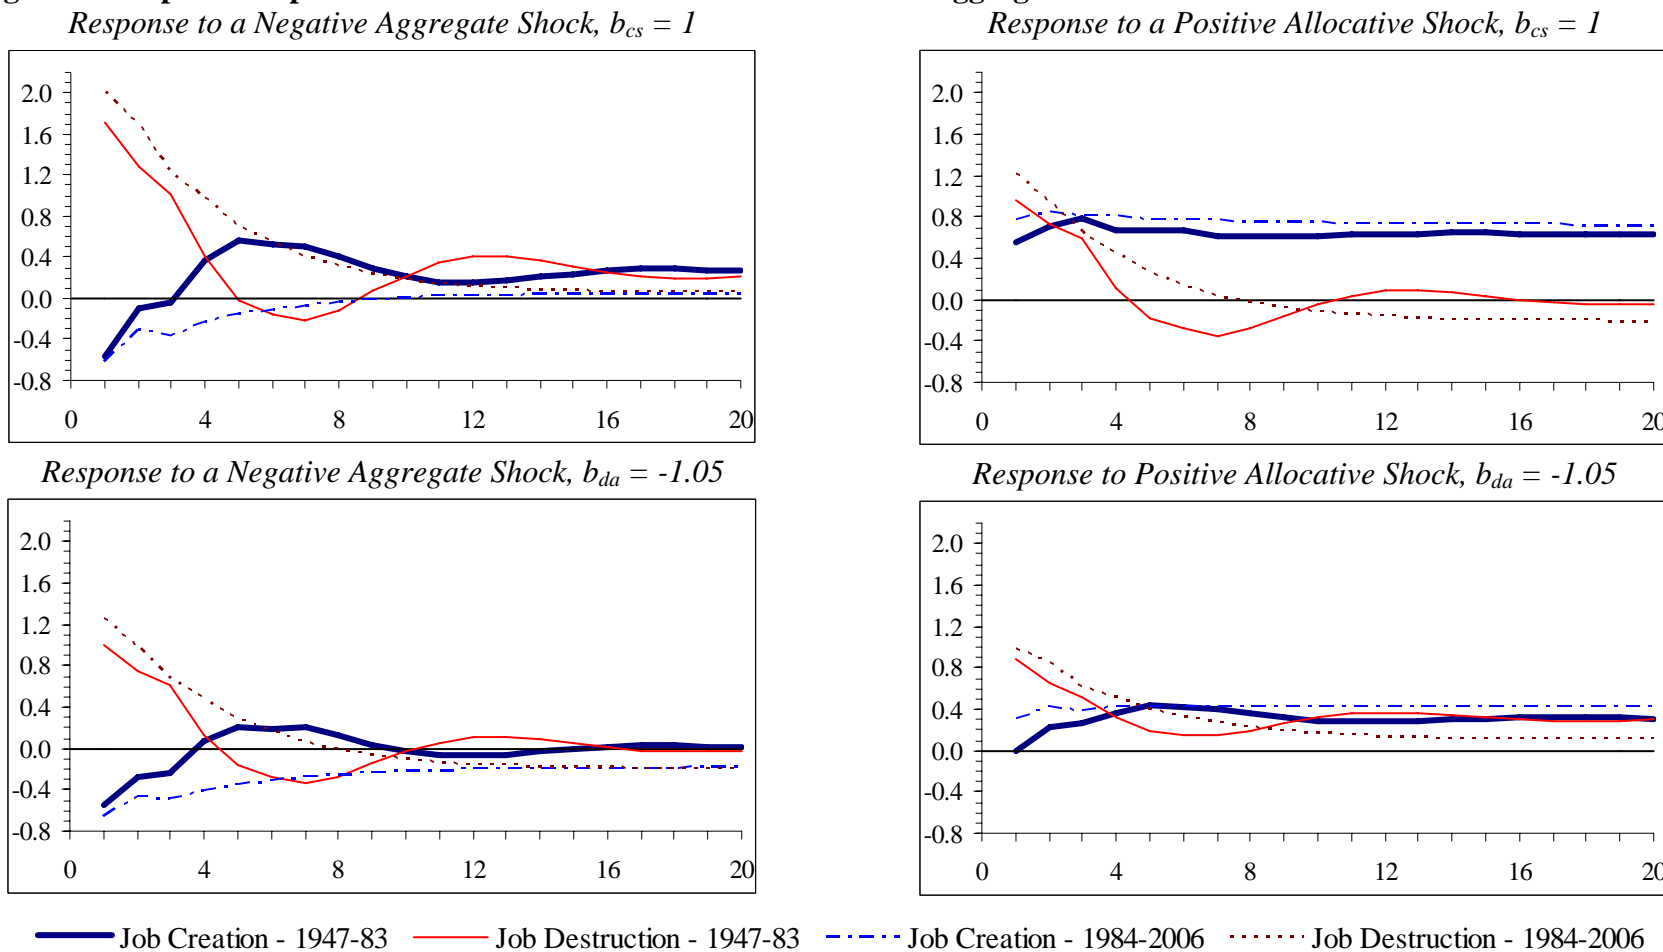


(b) *Identification Implications for Allocative and Aggregate Shock Estimates*



Notes: The figures show the implications of varying the identification assumptions for the structural VAR of job creation and job destruction movements described in the text. The implications are for subsamples of the data split at 1984:1. The upper panel illustrates the implied value of b_{cs} for each value of b_{da} for the early and later sample. The lower panel shows the standard deviations of the allocative and aggregate structural innovations as a function of the identifying assumptions for $\{b_{cs}, b_{da}\}$. See the text for details.

Figure 10. Impulse Responses of Job Creation and Job Destruction to Aggregate and Allocative Shocks



Notes: Panels report the impulse responses to a unity standard deviation change in the noted shock for job creation and job destruction during the early (1947:1-1983:4) and later (1984:1-2006:3) periods, given the listed identifying assumption. Results come from the impulse responses of separate structural VARs for each sample period. See text for details.

Table 1. GMM Splicing Estimation, Moments Matched, and Parameter Estimates*(a) Moments Matched*

Moment	Data Source	Time Period	Value
Mean C_t	BED	1990:2 – 1998:4	5.165
Mean D_t	BED	1990:2 – 1998:4	5.235
Variance of NET_t^C	CES	1947:1 – 1998:4	1.996
Variance of Excess Reallocation, XR_t	LTS-LRD	1947:1 – 1998:4	2.899
Relative Volatility, $\frac{\text{var}(POS_t)}{\text{var}(NEG_t)}$	LTS-LRD	1947:1 – 1998:4	0.658
Correlation of POS_t, NET_t	LTS-LRD	1947:1 – 1998:4	0.662
Correlation of NEG_t, NET_t	LTS-LRD	1947:1 – 1998:4	-0.794

(b) GMM Parameter Estimates

Parameter	Estimate	Parameter	Estimate
α_0	1.424	β_0	2.049
α_1	0.774	β_1	0.664
α_2	-0.072	β_2	-0.097

Notes: Job flow moments are from author's calculations using the listed data source. Parameter estimates are from the GMM estimation described in the text.

Table 2. Job Flows in Manufacturing: Summary Statistics

Period	1947 – 1966	1967 – 1983	1984 – 2006	Full Sample
<u>Mean (Standard Deviation), Unfiltered Data</u>				
Job Creation Rate	6.60 (0.76)	5.63 (0.49)	4.78 (0.53)	5.63 (0.98)
Job Destruction Rate	6.33 (1.17)	5.89 (1.03)	5.16 (0.68)	5.76 (1.09)
Net Growth Rate	0.27 (1.78)	-0.26 (1.35)	-0.38 (0.77)	-0.13 (1.37)
Job Reallocation Rate	12.93 (0.85)	11.52 (0.89)	9.94 (0.95)	11.39 (1.55)
Excess Reallocation Rate	11.59 (1.32)	10.49 (0.67)	9.35 (0.96)	10.51 (1.40)
<u>Relative Volatility, $\sigma_{jd} / \sigma_{jc}$</u>				
Unfiltered Data	1.54	2.11	1.29	1.11
HP-Filtered Data	1.51	2.10	2.08	1.74
<u>Correlations, Unfiltered</u>				
$\rho(JC_t, JD_t)$	-0.69	-0.51	0.22	0.13
$\rho(NET_t, JR_t)$	-0.53	-0.69	-0.26	-0.11
<u>Correlations, HP-Filtered</u>				
$\rho(JC_t, JD_t)$	-0.76	-0.46	-0.25	-0.60
$\rho(NET_t, JR_t)$	-0.55	-0.68	-0.64	-0.59

Note: Statistics are for estimates from the spliced series of job flow rates for manufacturing. See text for details.

Table 3. Volatility of Estimated Structural Shocks in Manufacturing

	Early Period <i>1947:1 – 1983:4</i>	Later Period <i>1984:1 – 2006:3</i>
<i>Assumption of Symmetric Responses to Allocative Shocks</i>		
b_{cs}, b_{da}	1, -1.85	1, -1.94
σ_a (Aggregate std.dev.)	0.381	0.154
σ_s (Allocative std.dev.)	0.239	0.165
Relative volatility, σ_s / σ_a	0.63	1.07
<i>Assumption of Symmetric Responses to Aggregate Shocks</i>		
b_{cs}, b_{da}	0.001, -1.05	0.357, -1.05
σ_a (Aggregate std.dev.)	0.450	0.205
σ_s (Allocative std.dev.)	0.577	0.265
Relative volatility, σ_s / σ_a	1.28	1.29

Note: Statistics are for separate structural VAR specifications for each sample period, given the listed identifying assumptions. See text for details.

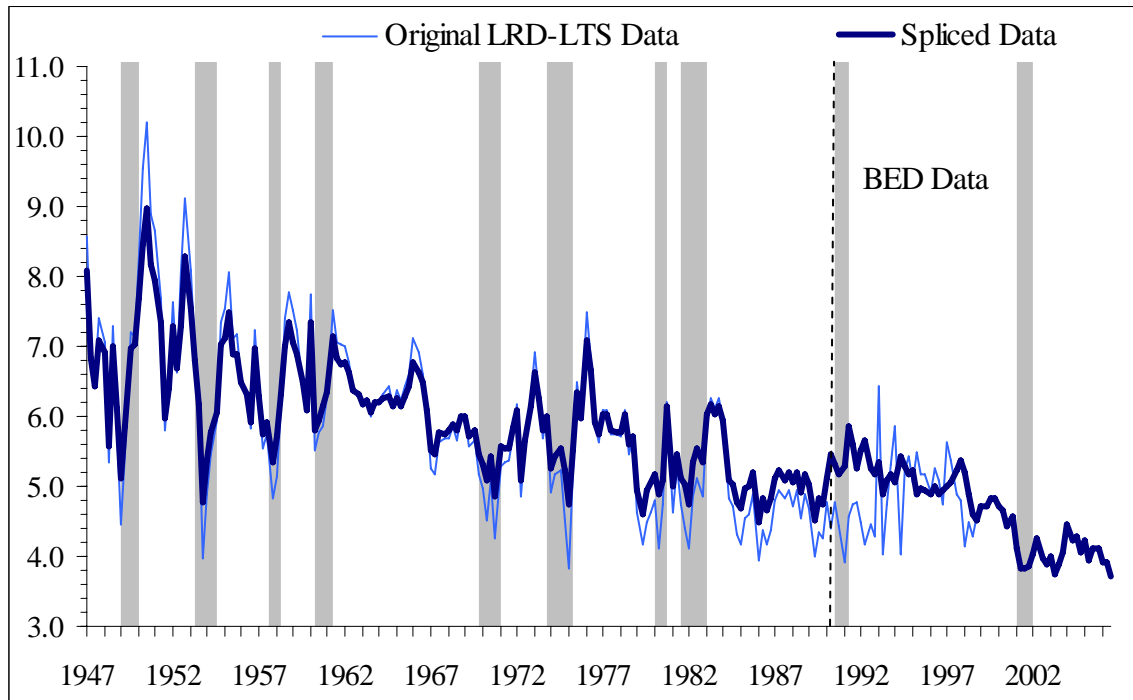
Table A.1. Results of Early BED Match Identification

	Initial Estimate		Corrected Estimate	
	<i>Thousands of Employees</i>	<i>Percent of Employment</i>	<i>Thousands of Employees</i>	<i>Percent of Employment</i>
First Quarter, 1991				
Changes at Openings	5,321	6.0	2,270	2.5
Changes at Closings	5,462	6.1	2,103	2.4
Changes at Expansions	4,402	4.9	4,685	5.3
Changes at Contractions	7,784	8.7	8,376	9.4
Second Quarter, 1992				
Changes at Openings	3,156	3.6	1,887	2.1
Changes at Closings	2,481	2.8	1,226	1.4
Changes at Expansions	6,642	7.5	6,747	7.6
Changes at Contractions	4,310	4.9	4,401	5.0

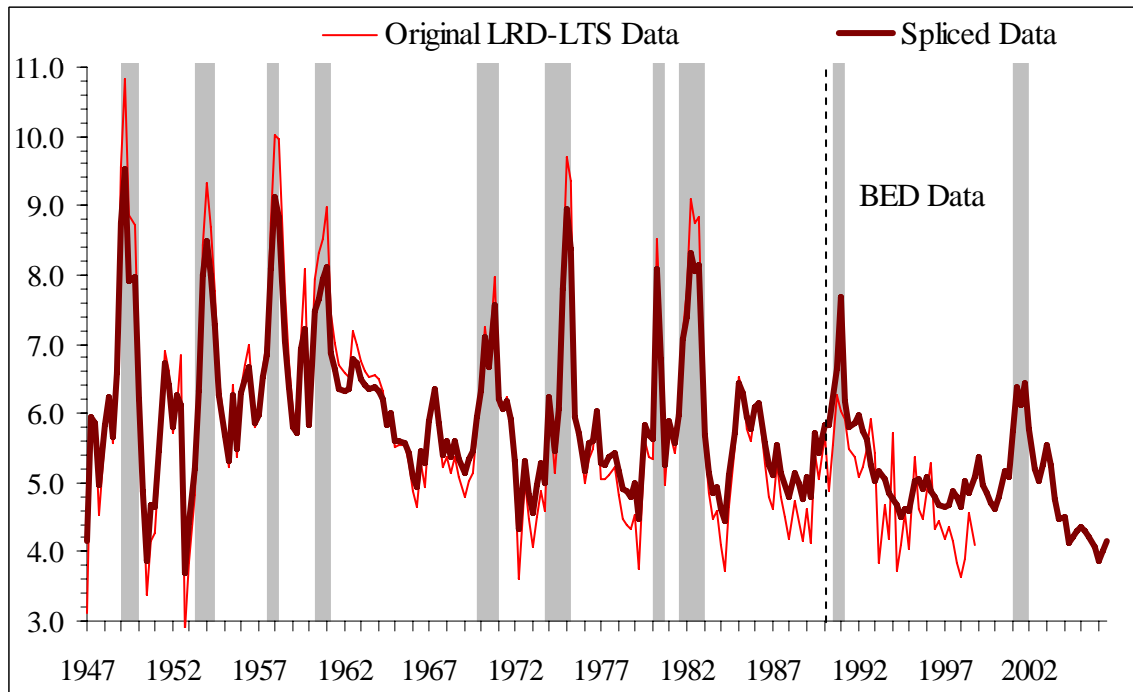
Note: Listed employment changes are prior to seasonal adjustment.

Figure A.1. Original vs. Spliced Manufacturing Job Flow Series

(a) Job Creation



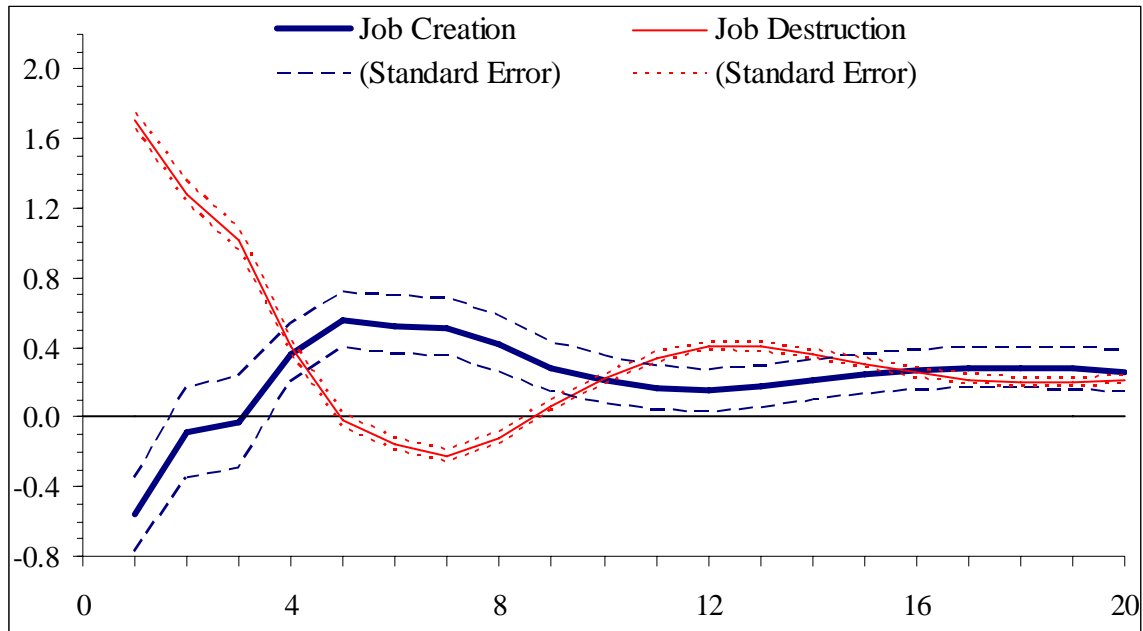
(b) Job Destruction



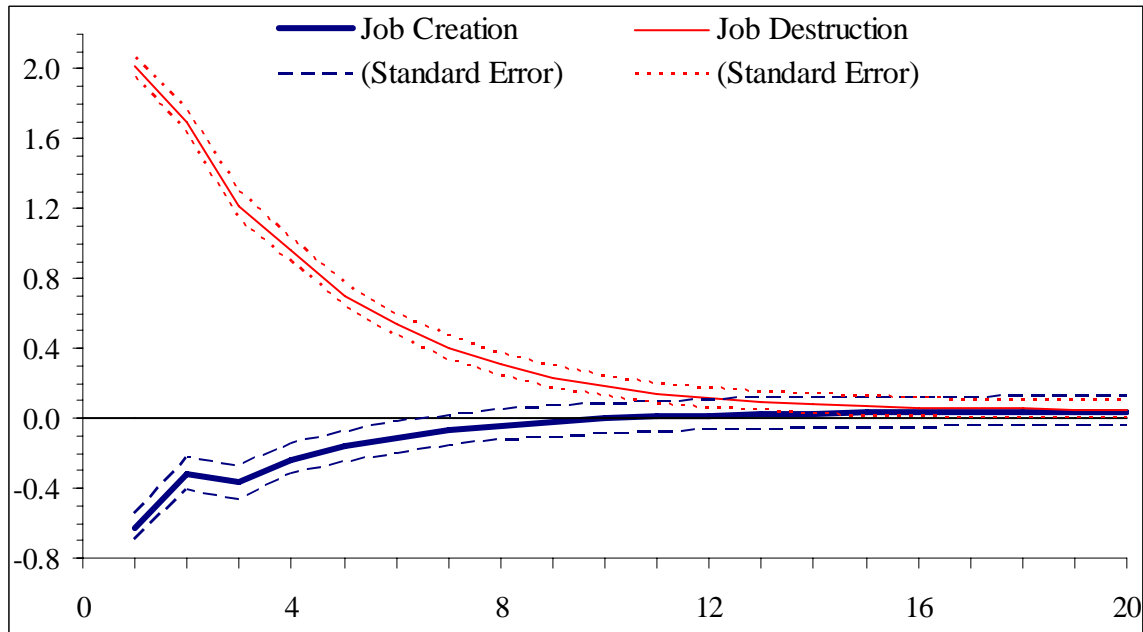
Notes: Panels compare series of spliced BED-LRD-LTS estimates derived from a GMM model to the original merged LRD-LTS series created by Davis and Haltiwanger (1999). The actual BED estimates are used to the right of the dashed vertical line in each panel.

Figure A.2. Impulse Responses (with Standard Errors) to an Aggregate Shock

(a) Early Period (1947:1 – 1983:4, $b_{cs} = 1$)



(b) Later Period (1984:1 – 2006:3, $b_{cs} = 1$)



Notes: Panels report the impulse responses, with standard error bands included, to a unity standard deviation negative aggregate shock for job creation and job destruction during the early (1947:1-1983:4) and later (1984:1-2006:3) periods, under the assumption of symmetric responses to an allocative shock. Results come from the impulse responses of separate structural VARs for each sample period. See text for details.