

NBER WORKING PAPER SERIES

THE CONCENTRATION OF JOB DESTRUCTION

Robert E. Hall

Working Paper 7025
<http://www.nber.org/papers/w7025>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 1999

This research was supported by the National Science Foundation under grants SBR-9410039 and SBR-9730341 and is part of the research program on Economic Fluctuations and Growth of the NBER. John Cochrane made influential comments at an early stage of this work. I am grateful to Éva Nagypál for helpful comments and assistance. The views expressed in this paper are those of the authors and do not reflect those of the National Bureau of Economic Research.

© 1999 by Robert E. Hall. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Concentration of Job Destruction
Robert E. Hall
NBER Working Paper No. 7025
March 1999
JEL No. E24

ABSTRACT

A time series is *concentrated* if the expectation of its current value is a negative function of a moving average of past values up to all but the most recent past. Job destruction has the property of concentration in a model of heterogeneous jobs because an adverse shock destroys jobs in plants close to the margin of shutdown. Until other plants drift close to that margin, there are fewer plants that are vulnerable to another adverse shock. Concentration is easy to spot in the autocorrelations of a time series, which will be negative except for the first few lags. A simple model generates data displaying concentration. Data on job destruction and employment change for U.S. manufacturing show unambiguous evidence of concentration. According to the simple model, job creation is more persistent and thus less concentrated than is destruction, a property reflected in the data as well.

Robert E. Hall
Hoover Institution
Stanford University
Stanford, CA 94305-6010
and NBER
hall@hoover.stanford.edu

I. Introduction

Substantial job destruction occurs during aggregate contractions. Thanks to pioneering work by Davis and Haltiwanger [1990 and 1992], data on job destruction have made possible much more thorough investigation of its role. This paper looks at one critical aspect of job destruction—its concentration during brief periods, typically the most intense parts of cyclical contractions.

I consider the following explanation of the temporal concentration of job destruction: Production occurs in units such as plants. Owners decide each period whether to continue or to shut down each unit by comparing a continuation value to a shutdown value. During periods of calm, the difference—the net continuation value—of some units drifts toward zero as these units depreciate physically or become obsolete. They enter a zone of vulnerability to adverse shocks. When a shock occurs, it clears out the zone of vulnerability. Until the zone begins to fill up again through the normal diffusion of units across the spectrum of productivity, another shock will have less effect. Thus a shock releases job destruction that might otherwise have occurred earlier and also accelerates destruction that would otherwise have occurred later. Job destruction is concentrated in one episode.

Temporal concentration arises when productive units are heterogeneous. I describe the simplest possible model, of the type developed by Caballero and Hammour [1996, 1998]; Caballero, Engle and Haltiwanger [1997]; and others, and show that it results in concentration of episodes of job destruction and inventory runoff. Data on job destruction and employment changes show temporal concentration. On the other hand, the model suggests that job creation is

a more persistent process, driven by the availability of unemployed workers. Data on job creation support this prediction as well.

II. Concentrated Processes

I define a concentrated process as one whose current value is likely to be lower if its values over a span of time up to the recent past have been higher. Let d_t be a covariance-stationary series. I define its *concentration function* with lag t and window N to be

$$c_{t,N} = -\frac{E[d_t | d_{t-t} + \dots + d_{t-t-N+1}]}{\frac{1}{N}(d_{t-t} + \dots + d_{t-t-N+1})}. \quad (2.1)$$

This is negative of the coefficient of the regression of d_t on the lagged moving average, $\frac{1}{N}(d_{t-t} + \dots + d_{t-t-N+1})$. For a window one observation wide, the concentration function is the negative of the autocorrelation function,

$$f_t = -c_{t,1} = \frac{\text{Cov}(d_t, d_{t-t})}{V(d)}. \quad (2.2)$$

A concentrated series has a concentration function with positive values. The concentration function provides a way to determine the degree of concentration of a series and a way to test for concentration. The reason for skipping the first $t - 1$ lags will become apparent from some examples shortly.

Concentration is also manifested in the three standard representations of time series. First, as noted above, a concentrated series will generally have negative autocorrelations except at the shortest lags. Second, in the moving

average representation of a concentrated series, the coefficients on longer lags of the innovation are negative. An innovation causes a unit increase in the series when it occurs, but causes lower future values. For job destruction, an impulse that causes destruction today means less destruction is likely in the future. Third, in the autoregressive representation, the coefficients on longer lags of a concentrated series are negative. Concentration could be detected and measured by estimating an autoregressive-moving average process and then studying either its moving-average or autoregressive representations. However, it appears that the estimation of a full process with many parameters results in weaker tests than the direct estimation of a well-chosen point on the concentration function. I will refer to such a single value of the concentration function as a concentration coefficient.

A concentration coefficient can be estimated as the negative of the coefficient in the regression of d_t on $\frac{1}{N}(d_{t-t} + \dots + d_{t-t-N+1})$. Its sampling dispersion can be measured by the Newey-West robust standard error. The reason for choosing a window of width N greater than one is, as usual, to reduce the effect of sampling error. With very long series, or in situations where the autocorrelations are known with certainty, it is appropriate to use $N = 1$ and thus to study the autocorrelation function or correlogram. The correlogram is also useful to diagnose concentration even when sampling variation is considerable.

A. Concentration with a White-Noise Driving Force

A simple example of a concentrated process is

$$d_t - \bar{d} = -\frac{1}{1-fL}(d_{t-1} - \bar{d}) + e_t . \quad (2.3)$$

Let

$$V_t = \frac{1}{1-fL}(\bar{d} - d_{t-1}), \quad (2.4)$$

interpreted as the stock of vulnerable jobs, which builds up during periods of below-average job destruction. Then the model can be written as

$$d_t = f V_t + e_t. \quad (2.5)$$

Job destruction is more likely when the stock of vulnerable jobs is high.

The correlogram for this process is:

$$f_t = -f \frac{1-f^2+fl}{1-f^2+2fl} (f-l)^{t-1}. \quad (2.6)$$

The autocorrelations are negative, starting near $-f$ at lag 1 and falling in magnitude at rate $f-l$ for longer lags.

B. Concentration with a Persistent Driving Force

A more general model is

$$d_t = -b(L)d_{t-1} + z_t, \quad (2.7)$$

where $b(L)$ embodies concentration and z_t is moderately persistent. The driving force z_t might be a low-order moving average or an autoregressive process with roots not too close to one in modulus. The dynamics of d_t will reflect both those of the driving process and of concentration.

Consider an MA(1) process for the driving force,

$$z_t = (1+qL)e_t, \quad (2.8)$$

with $q \leq 1$ without loss of generality, together with the concentration specification of the previous example,

$$b(L) = \frac{1}{1-fL}, \quad (2.9)$$

The left side of Figure 1 shows the correlogram for $q = 0.7$.

In the case of an AR(1) driving process, say

$$z_t = \frac{1}{1-\gamma L} e_t, \quad (2.10)$$

with $\gamma = 0.5$, the resulting correlogram is shown on the right side of Figure 1, with the same concentration specification.

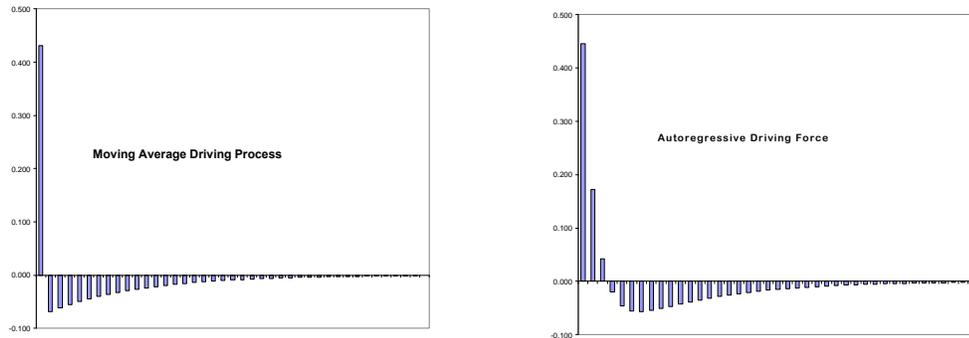


Figure 1. Correlograms of Concentrated Processes with Moderately Persistent Driving Processes

The left panel shows the autocorrelations when the concentration results from a geometric lag process and the driving force is first-order moving average. The right panel refers to the same situation except that the driving force is first-order serially correlated

For both MA and AR driving processes, the resulting correlograms of the series combining the driving process and the concentration effect have the shape

characteristic of a concentrated series except for the low-order correlations. Estimation of a concentration coefficient with an appropriate lag t should reveal the effect of concentration even though it is obscured by the dynamics of the driving process for the low-order coefficients.

As these examples illustrate, the detection of concentration is most effective if the lag, t , is chosen just high enough to avoid contamination by the short-run dynamics of the driving process, and the window width, N , is chosen not so high as to extend the window into ranges where the concentration effect has disappeared. In results reported later in this paper for U.S. job destruction, I use $t = 6$ quarters of lag and $N = 4$ quarters wide.

III. A Model of Shutdowns

The model developed in this section embodies standard ideas about job destruction and the life cycles of units—see my companion paper Hall [1999] for a discussion of a fuller model and the related literature. All job destruction occurs because of shutdowns. The purpose of the model is to make some simple points about the behavior of job destruction, not to create a general equilibrium model of recessions. In the model, there is a discrete set of alternative levels of profit that a unit can earn. The concept of profit is the one suited to the decision-making unit, which may or may not internalize all costs associated with job destruction (see Hall [1999]). A new unit starts at the highest level, of profit, 1. With probability d each period, profit, z_t , falls by a factor $w < 1$:

$$\Pr[z_t = z_{t-1}] = 1 - d \tag{3.1}$$

$$\Pr[z_t = w z_{t-1}] = d .$$

Each period, all units shut down that would otherwise be at or below a cutoff level of profit z_t^* ; the cutoff level evolves as a stationary random variable. Shutdowns define the flow of job destruction in the model. Upon shutdown, workers remain idle for one period and then have a constant probability each period, f , of employment at a new unit, which starts at the highest level of profit.

Figure 2 shows 270 periods of history of job destruction from the model, with the shutdown margin following an AR(2) process. Parameter values are:

<i>Parameter</i>	<i>Interpretation</i>	<i>Value</i>
d	Probability of transition to lower profit level	0.2
	Roots of process for shutdown margin	0.9, 0.5
	Mean of shutdown margin	0.84
	Standard deviation of innovation in shutdown margin	0.1
w	Profit reduction factor	0.98
f	Job-finding rate	0.2

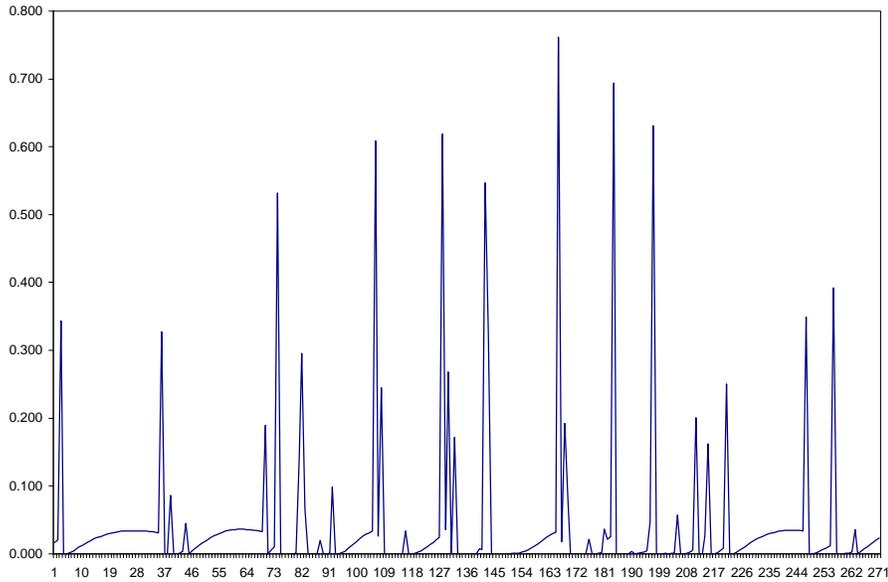


Figure 2. Job Destruction from Model.

The model generates occasional bursts of job destruction, during times when the shutdown margin rises into a region of the current distribution of units across categories where there are substantial numbers of vulnerable units. Figure 3 shows the correlogram of job destruction. For lags past the first, the autocorrelation of job destruction is negative, showing that the series is concentrated. With $t = 6$ and $N=4$, the concentration coefficient is 0.198 with a standard error of 0.039 (based on 1300 observations).

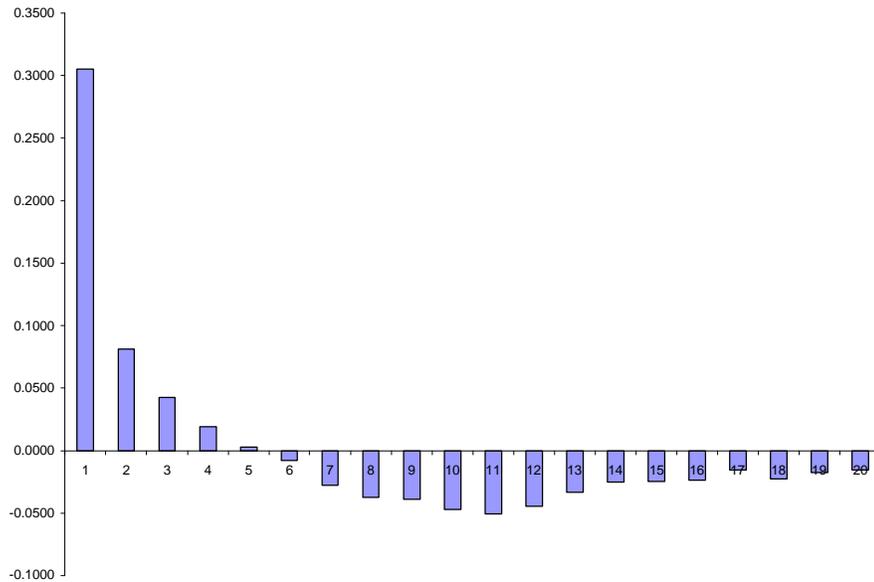


Figure 3 Correlogram of Job Destruction from Model.

Employment dynamics depend on job creation as well as job destruction. A simple view of job creation is that a job is always created when a qualified worker encounters an employer. The model of this section embodies that view. The flow of job creation is proportional to the number of unemployed workers searching for new jobs. Job creation shares the highly persistent dynamics of unemployment. Aggregate shocks have their immediate effect on job destruction in the short run and only affect job creation through the cumulating stock of searching workers.

In the simple model, an impulse that results in a sharp spike in job destruction causes a prolonged period of unemployment and job creation. Job creation is a geometric moving average of job destruction, so the correlogram of creation has much more positive serial correlation at short and medium lags. Eventually, the concentration of job destruction shows through in negative serial

correlations of creation, but only after a considerable lag. Figure 4 shows the correlogram of job creation for the simple model.

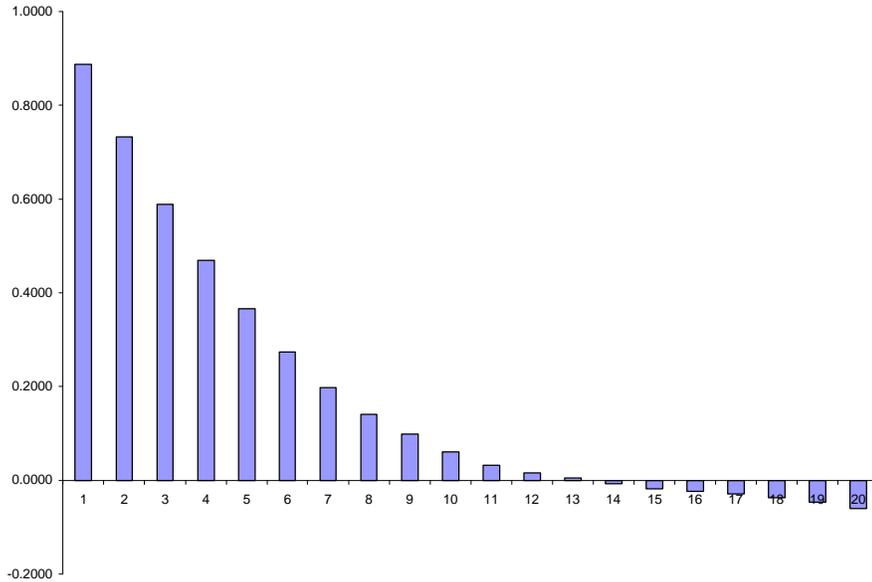


Figure 4. Correlogram of Job Creation from Model

IV. Job Destruction in U.S. Manufacturing

The Davis-Haltiwanger data break job destruction into components from plant shutdowns and from employment reductions in plants that continue to operate. The second type of job destruction includes the elimination of shifts and the closing of productive units within plants, as well as reductions in staffing that do not involve shutting down discrete units. Over the entire sample period, 1972-88, 12 percent of total job destruction in manufacturing occurred because of plant shutdowns. Plant-shutdown job destruction is more sensitive to aggregate shocks than is the remainder of job destruction—shutdown destruction rose to 18.5 percent of total destruction in 1974, 21.3 percent in 1981, and 23.9 percent in

1986. Davis, Haltiwanger, and Schuh [1996, p. 29] report that another 44 percent of destruction takes the form of employment reductions of 25 percent to 99 percent.

Figure 5 shows job destruction from plant closings, seasonally adjusted, as measured by Davis and Haltiwanger from the Longitudinal Research Database. For details on the construction of the data, see Davis, Haltiwanger, and Schuh [1996]. In brief, the database reports employment each quarter in a large random sample of manufacturing plants. A closing is defined as a change to zero employment. The series is calculated as the ratio of employment in all closed plants in the quarter before closing to total employment in the sample.

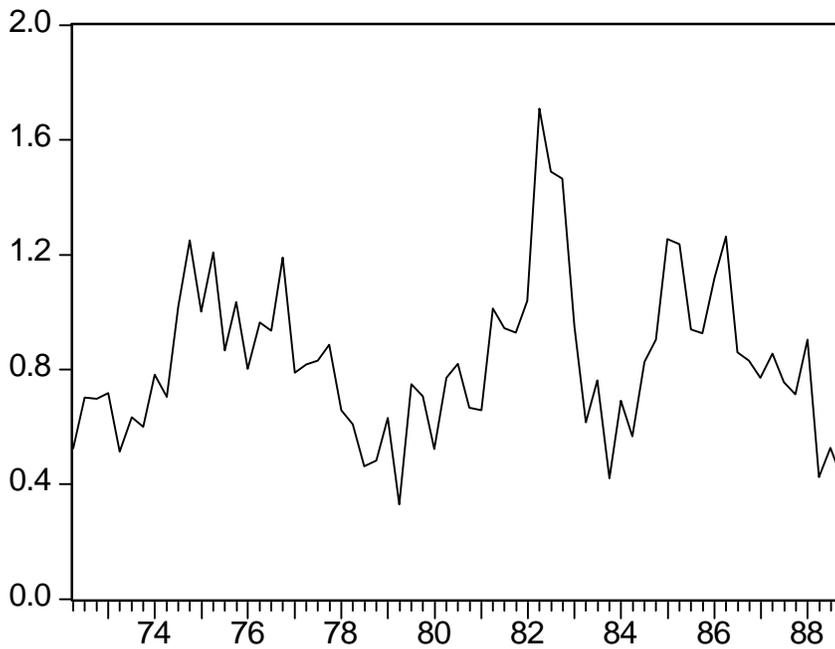


Figure 5. Shutdown Job Destruction Rate in U.S. Manufacturing, 1972-1988

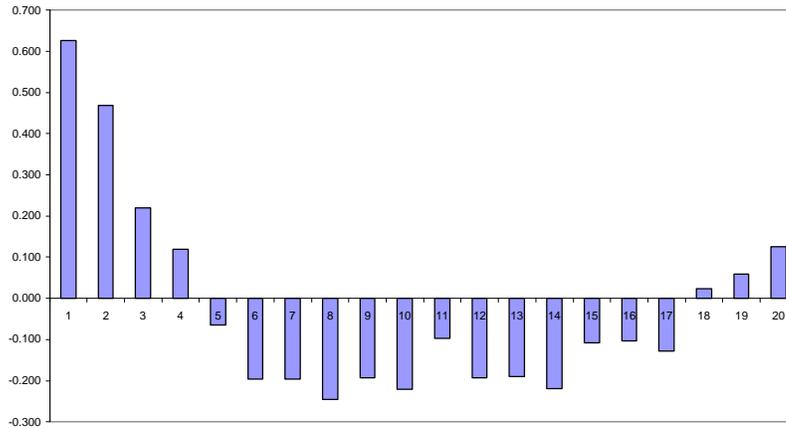


Figure 6 Correlogram of Shutdown Job Destruction

Figure 6 shows the correlogram of the shutdown job destruction rate. The positive values for the first three lags appear to relate to the dynamics of the driving force. The autocorrelations for lags 5 through 17 are negative, suggesting considerable concentration. The concentration coefficient with lag 6 and width 4 is 0.395 with a standard error of 0.183, statistically unambiguous evidence of concentration ($p = 0.035$).

The data in Figures 5 and 6 relate to plant closings. The bulk of job destruction arises not from plant closings but from events that cause employment in a plant to shrink to levels still above zero. Shifts may be eliminated, units within plants shut down, or there may be reductions in staffing without explicit shutdowns. Although these decisions involve stronger interactions with other influences than the decision to shut a plant, the character of the decision is the same. In particular, all forms of job destruction should be concentrated for the reasons discussed earlier in this paper.

Figure 7 shows job destruction at continuing plants, seasonally adjusted, as compiled by Davis and Haltiwanger. This component of job destruction also tends to concentrate in bursts, such as in 1982 and 1975.

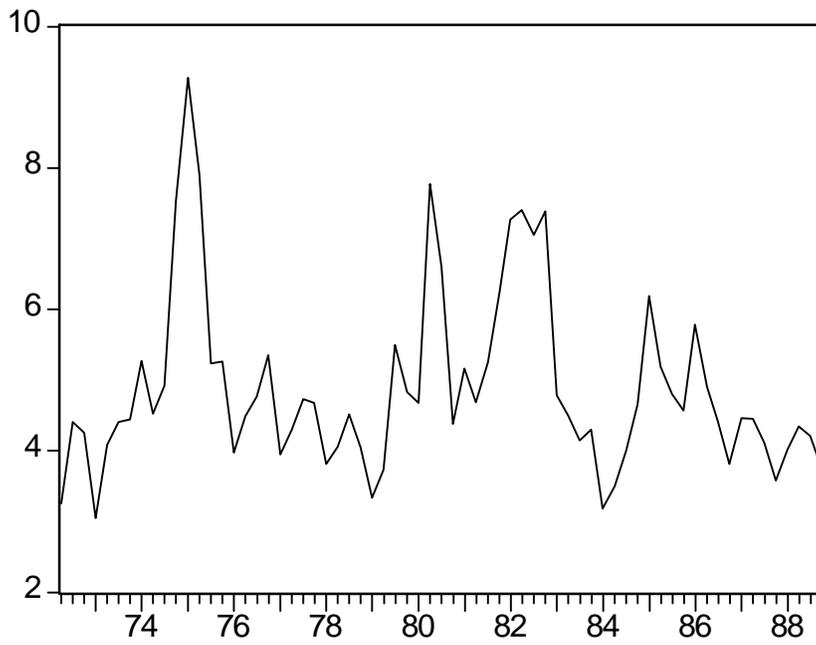


Figure 7. Job Destruction Rate for Continuing Plants in U.S. Manufacturing, 1972-1988

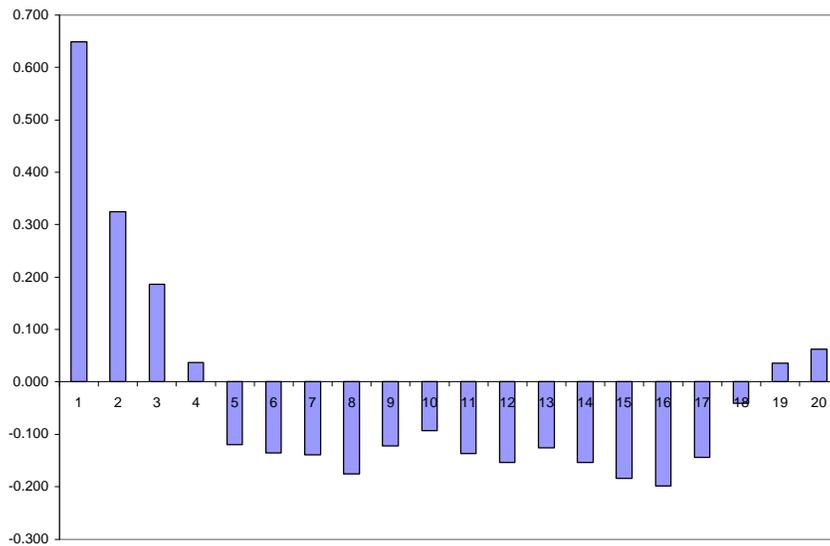


Figure 8. Correlogram of Job Destruction at Continuing Plants

Figure 8 shows the correlogram of job destruction at continuing plants. Again, the first few autocorrelations are positive, presumably reflecting the short-run dynamics of the driving force. At lag 5, they turn negative and remain consistently negative through lag 18. The concentration coefficient with lag 6 and width 4 is 0.285 with a standard error of 0.228, reasonable but not statistically unambiguous evidence of concentration ($p = 0.217$). A longer window would probably result in a nominally stronger finding, but the test is hard to interpret if it is based on a preliminary inspection of the correlogram.

V. Job Creation

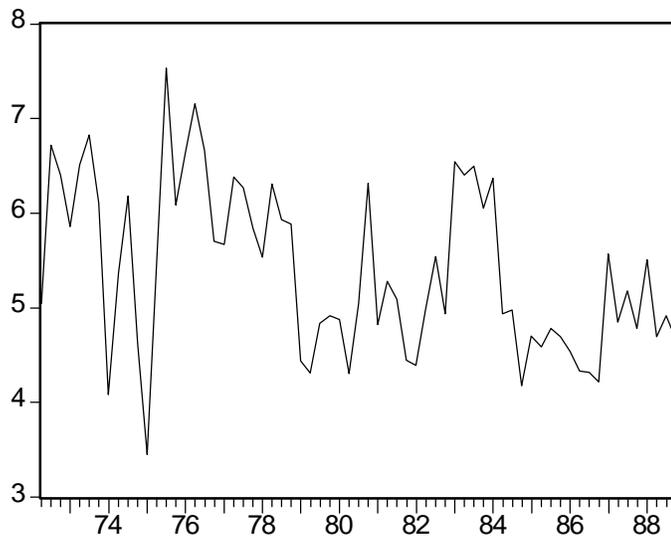


Figure 9. Job Creation Rate in U.S. Manufacturing

Figure 9 shows job creation in U.S. manufacturing as measured by Davis and Haltiwanger. There is no doubt that job creation reaches its highest levels around the time that unemployment reaches its maximum during a recession,

usually a few quarters after the initial shock. On the other hand, in some recessions, notably in 1975, job creation falls sharply, contrary to the simple model. But the correlogram in Figure 10 shows that the model's predictions about concentration of job creation are consistent with the data. There is no systematic tendency for the autocorrelation of job creation to turn negative after a few quarters.

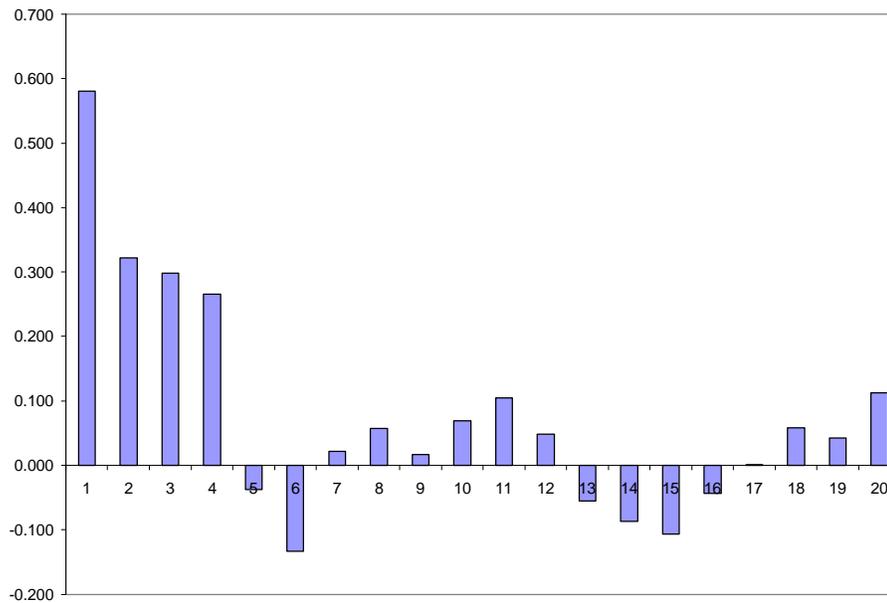


Figure 10. Correlogram of Job Creation.

VI. Concentration of Employment Changes

In the simple model, most of the short-run dynamics of employment come from job destruction. Data on job destruction and creation are generally consistent with this view, though creation has more short-run movements than is strictly consistent with proportionality of creation and unemployment. Nonetheless, it seems reasonable to expect to find concentration in employment changes, which

are the negative of the difference between job destruction and creation. Figures 11 and 12 show the correlograms of the log-changes in employment in durables manufacturing and non-durables.

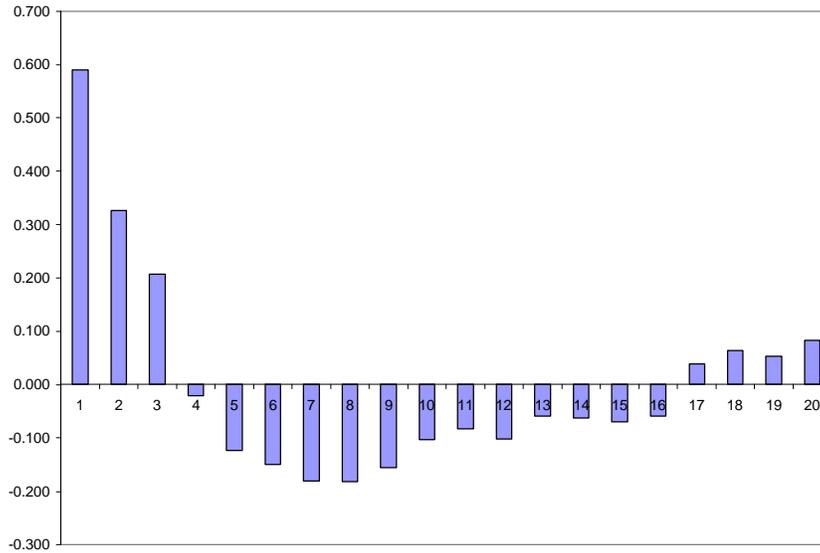


Figure 11. Correlogram of Employment Change, Manufacturing Durables

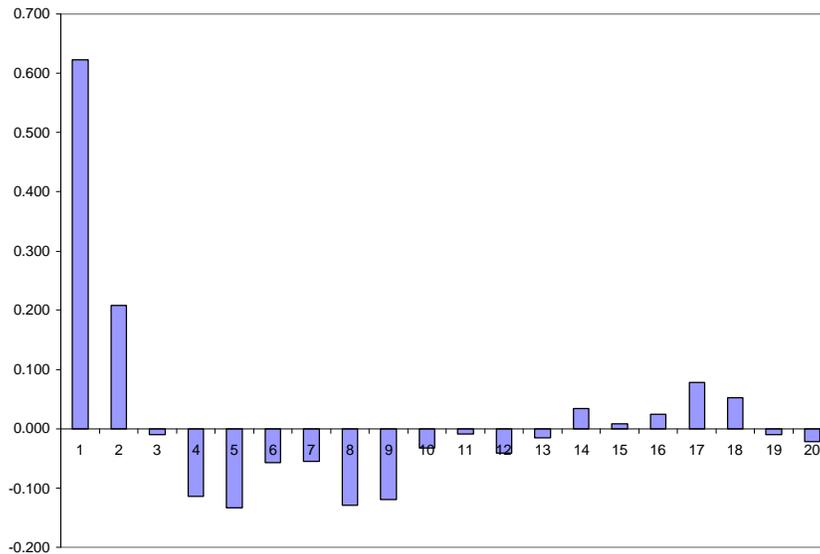


Figure 12. Correlogram of Employment Change, Manufacturing Non-Durables

Concentration is pronounced in durables and is statistically unambiguous. The concentration coefficient for lag 6 and window width 4 is 0.277, with a standard error of 0.10 (p value of 0.006). Though the correlogram for non-durables shows signs of concentration, the evidence is ambiguous: the coefficient is 0.172 with a standard error of 0.147 (p value of 0.243).

VII. Concluding Remarks

A simple model of heterogeneous jobs has a sharp testable implication: The flow of job destruction should have the econometric property of concentration. The likelihood of additional job destruction falls below normal a few quarters after a spike of job destruction. On the other hand, job creation tracks the stock of unemployed workers, and creation is more persistent than is destruction. Data from the U.S. manufacturing sector supports these predictions.

References

- Caballero, Ricardo J. and Mohamad L. Hammour. 1996. "On the Timing and Efficiency of Creative Destruction," *Quarterly Journal of Economics*. 111:805-852. August.
- Caballero, Ricardo J. and Mohamad L. Hammour. 1998. "Jobless Growth: Appropriability, Factor Substitution, and Unemployment," *Carnegie-Rochester Conference Series on Public Policy*. 48: 51-94, June.
- Caballero, Ricardo J., Eduardo M.R.A. Engel, and John Haltiwanger. 1997. "Aggregate Employment Dynamics: Building from Microeconomic Evidence," *American Economic Review* 87:115-137, March.
- Davis, Steven J., and John Haltiwanger. 1990. "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," *NBER Macroeconomics Annual* 5: 123-168.
- Davis, Steven J., and John Haltiwanger. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics* 107:819-863.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*. Cambridge: MIT Press.
- Hall, Robert E. 1999. "Aggregate Job Destruction and Inventory Liquidation." NBER Working Paper 6912, January.