Job Reallocation and the Business Cycle: New Facts for an Old Debate

Scott Schuh and Robert K. Triest*

Can the reallocation of factors of production among firms and sectors, or the restructuring of production technology, cause business cycle fluctuations? Theoretically, the answer is unequivocally "yes." But *do* reallocation and restructuring actually cause fluctuations? The answer to this question—like most questions about economic causality—is much less clear. Indeed, determining whether or not reallocation and restructuring cause fluctuations is as fundamentally difficult as identifying supply and demand. The purpose of this paper is to examine the relationship between the business cycle and reallocation and restructuring, with a particular emphasis on trying to learn whether reallocation causes recessions. We conduct our investigation by briefly assessing existing evidence and theories of reallocation and restructuring, and then providing new evidence designed to improve our understanding of the relationship.

Held a generation ago, this conference probably would not have included a session devoted to understanding the role of reallocation and restructuring (henceforth, simply "reallocation") in business cycle fluctu-

*Economists, Federal Reserve Bank of Boston. The authors have benefited from comments by Hoyt Bleakley, Lynn Browne, Ricardo Caballero, Steven Davis, Jeff Fuhrer, and Michael Klein. Kevin Daly and Marie Willard provided excellent research assistance, Georgeanne DaCosta excellent typesetting assistance. We thank Joyce Cooper, Steven Davis, Tim Dunne, Lucia Foster, John Haltiwanger, and C.J. Krizan for providing data, documentation, and general data assistance. The research in this paper was conducted while the authors were Census Bureau research associates at the Boston Regional Data Center. Research results and conclusions expressed are those of the authors and do not necessarily indicate concurrence by the U.S. Bureau of the Census, the Federal Reserve Bank of Boston, or the Board of Governors of the Federal Reserve System. This paper has been screened to ensure that no confidential data are revealed.

ations. But during the 1970s and 1980s, the U.S. economy endured an intense period of reallocation associated with large, persistent energy price increases, increasing international trade, widespread deregulation, demographic changes from the Baby Boom, regional migration, sweeping financial market innovations, and shifts in the level and composition of government spending. It was also an economically painful period, with stubbornly high rates of inflation and unemployment, permanent job loss by high-wage experienced workers, plant closings and permanent "downsizing," slower real wage and trend productivity growth, and an increasing wage gap between skilled and unskilled workers. Traditional macroeconomic models at this time did not incorporate reallocation and thus were unable to explain many of these phenomena and the economic turmoil they generated.

Reflecting on this period, economists began investigating the role that reallocation among firms and sectors may have played in producing or amplifying the macroeconomic problems of the time. Initially, attention focused on two features: (1) heterogeneity and changes in the demographic characteristics of unemployed workers, and (2) dispersion in employment growth across highly aggregated sectors of the economy. The former was subsumed in calculations of the natural rate of unemployment and had little effect on the fundamental behavioral characteristics of macroeconomic models. The latter was dismissed initially as merely a by-product of heterogeneity in the cyclical sensitivity of sectors. Thus, the prevailing macroeconomic view continued to be that recessions were periods of temporarily low aggregate demand and that firms responded to this reduced demand by temporarily laying off workers, recalling them when government demand-management policies kicked in to raise aggregate demand. What prevented reallocation from being taken seriously as a factor in business cycle fluctuations was the lack of convincing empirical evidence.

Evidence of an important business cycle role for reallocation began unfolding in the late 1980s. Building on a limited base of earlier work, Steven Davis and John Haltiwanger (later joined by Scott Schuh) embarked on an extensive project of measuring the gross flows of jobs across U.S. manufacturing establishments using a unique, and particularly well-suited, new data base at the U.S. Bureau of the Census called the Longitudinal Research Database (LRD).[1] As summarized in their recent

[1] The research includes Davis and Haltiwanger (1990, 1992, 1995, 1998) and Davis, Haltiwanger, and Schuh (1990, 1996). Precursors include: the U.S. Bureau of Labor Statistics Manufacturing Turnover Data Base, Leonard (1987), Dunne, Roberts, and Samuelson (1989), and Blanchard and Diamond (1990). Other studies with U.S. data include: Troske (1993); Lane, Isaac, and Stevens (1994); Anderson and Meyer (1994); and Foote (1998). See Table 2.2 in Davis, Haltiwanger, and Schuh (1996) for references to international evidence on job flows.

book, Davis, Haltiwanger, and Schuh (1996), the extent of gross job flows is remarkable: One in five manufacturing jobs is either created or destroyed each year, on average. But these job flows are not the result of transitory, heterogeneous responses of plants to business cycle fluctuations. Rather, they reflect primarily *permanent* job creation and destruction that cause *permanent* relocation of workers to new jobs throughout the economy.

More important for the purpose of this paper is the cyclicality of the gross job flows. Not surprisingly, job destruction increases regularly and dramatically during recessions. The surprising result is that job creation does not decrease in recessions nearly as much as destruction increases—in fact, sometimes it even *increases* during recessions. Furthermore, sharp increases in job destruction are often followed by surges in job creation, as dislocated workers find new jobs elsewhere over time. Together, these job flow dynamics produce a strongly countercyclical rate of total job reallocation (the sum of creation and destruction). In other words, gross job reallocation increases during recessionary periods and decreases during expansionary periods. This reallocation of jobs often entails permanent displacement of workers across plants, destruction of human capital, and reduction in permanent income.

Many economists find these results intriguing on a variety of dimensions, but the most pertinent feature for macroeconomists is the surprising countercyclicality of permanent job reallocation. Why does permanent job reallocation rise during recessions? Traditional macroeconomic models are essentially silent about the heterogeneity of individual agents and firms and thus have difficulty explaining this phenomenon. Underlying economic churning—the rise and fall of particular firms, the employment and unemployment of particular individuals—is acknowledged to exist but viewed as benign for understanding macroeconomic fluctuations. To a first approximation, inflation, unemployment, and interest rates depend only on the level of aggregate demand in traditional macroeconomic models. Thus, traditional models must be modified and enhanced to incorporate and explain countercyclical reallocation.

In recent years, economists have begun proposing numerous theories to explain countercyclical job reallocation. Proposed theories must confront the following question: Does reallocation cause business cycles, or do business cycles cause reallocation? If the factors that determine the allocation of economic activity across plants and sectors change and induce costly, time-consuming reallocation, then aggregate economic activity will decline and reallocation will have caused the business cycle. Alternatively, if the economy experiences a business cycle slump and the slump leads plants and sectors to permanently destroy and create jobs, then the business cycle will have caused reallocation. More specifically, the countercyclical movements in job reallocation rates are initiated by sharp increases in job destruction prior to, and during, recessions. Thus, the theories must articulate what causal force(s) drive the increased job destruction, though they may also explain more generally the dynamic patterns of gross job creation and destruction that follow.

Broadly speaking, theories of countercyclical reallocation can be classified by how they answer these questions. One type stresses the role of *allocative* forces that induce reallocation across firms and sectors. Because reallocation is costly and time-consuming, aggregate demand declines. A second type stresses the role of *aggregate* forces that reduce aggregate activity, producing recessions. Either the reduction in aggregate activity directly increases (decreases) job destruction (creation), thereby setting off reallocation activity, or it indirectly induces reallocation activity by creating incentives to engage in such activity when the opportunity cost is low. Although most theories provide a role for both allocative and aggregate forces during recessions, the causal ordering of allocative and aggregate forces is usually a defining feature of the theories.

Despite the vast wealth of new empirical evidence on reallocation, and despite the technical and intellectual impressiveness of the new theories of countercyclical reallocation, no consensus has been reached on whether reallocation causes business cycle fluctuations. One key reason the issue is unresolved is its sheer complexity. What we know about reallocation is just the tip of the iceberg, and modern theories of reallocation model only a small fraction of what we know. Much like the fabled blind men trying to describe an elephant, most of our attempts to understand the complex process of reallocation have been just one piece of the puzzle at a time.

To gain a better understanding of countercyclical reallocation, prior empirical and theoretical analyses must be expanded, deepened, and corroborated along several dimensions. First, most analysis has relied on one-dimensional characteristics: industry, region, size, age, and the like. Much is yet to be learned from simultaneous disaggregation along multiple dimensions. Second, most of the analysis has focused solely on *job* reallocation. But firms choose labor simultaneously with other factors of production, and factor prices, productivity, and inventories also affect these choices. Jointly examining all aspects of production should improve our understanding of reallocation. Third, virtually all analysis has focused on the supply side. But demand factors, such as the level of demand, product innovation, product mix, market structure, and regional economic conditions, are also important determinants of reallocation. Finally, consideration of how expectations, uncertainty, and learning affect reallocation has been limited.

Our goal in this paper is to begin addressing a modest number of these issues. In the next two sections, we characterize the relationship between job reallocation and the business cycle with an updated review of the existing empirical evidence and a nontechnical summary of theories of countercyclical reallocation. In the remainder of the paper, we provide some new evidence, similar in style to that in Davis, Haltiwanger, and Schuh (1996), that provides a deeper understanding of this relationship and an overview of the research we intend to pursue. Although we do not formally test theories of countercyclical reallocation, we use them to formulate appropriate and interesting empirical exercises. Based on the evidence, we draw inferences about reallocation and business cycles, note potential implications for theories, and suggest areas for further investigation. Finally, we close with some new and up-to-date information about current job reallocation developments, and we hazard some guesses about the likelihood of future reallocation.

Although we do not conclusively resolve the question of causality, two general findings emerge that advance our understanding of job reallocation and business cycles. First, much of the cyclical fluctuation in gross job creation and destruction occurs in larger plants with relatively moderate employment growth that tends to be transitory, especially at medium-term horizons (up to five years). Unusually large employment growth rates, especially plant start-ups and shutdowns, are primarily small-plant phenomena and tend to be permanent, less cyclical, and to occur later in recessions than moderate, transitory job flows. Further, high rates of job flow rates occur primarily in plants that recently have been experiencing sharp employment contractions or expansions. Second, we discover that some of the key variables that should determine the allocation of factors of production across plants and sectors do in fact appear to be related to gross job flows, particularly to job destruction. Relative prices, productivity, and investment all exhibit suggestive time series correlations with the process of job reallocation that lead us to suspect that allocative driving forces may contribute significantly to business cycle fluctuations.

Facts about Job Reallocation and the Business Cycle

We begin with a brief review of job reallocation in U.S. manufacturing by summarizing the evidence from Davis, Haltiwanger, and Schuh (1996) (henceforth, "DHS") for the period 1972 to 1993, which includes new evidence on job reallocation during the latest recession. This section defines terminology, presents the data in graphical and tabular form, and restates the salient features of the data. Readers interested in more details may consult DHS.

Measurement of job reallocation begins at the plant level. The Longitudinal Research Database (LRD) contains employment data for about 50,000 to 70,000 plants each year. These plants are linked over time, and their employment change is measured each period. Plants whose employment increases are said to have created jobs; plants whose employment decreases are said to have destroyed jobs.[2] Gross job creation is the sum of employment gains at all plants with increasing employment, and gross job destruction is the sum of employment losses at all plants with decreasing employment. These gross job flows are expressed as rates relative to total employment. Gross job reallocation is the sum of gross job creation and destruction, and net employment (job) growth is the difference between gross job creation and destruction. The persistence of job flows is the fraction of jobs created (destroyed) that still (do not) exist in some future period (for example, one year later).

Four basic facts emerge from the gross job flow data, which are plotted and summarized in Figure 1 and Table 1. First, the gross flow rates are *large* at all times. Roughly one in 10 jobs is created, and one in 10 destroyed, each year on average in the annual data. The annualized flow rates are much larger for the quarterly data because of seasonal fluctuations, temporary layoffs, measurement error, and other transitory factors. These data reveal that a substantial fraction of employment—sometimes more than one-fourth—is involved in a continuous process of gross job reallocation, which can be costly and time-consuming. Even in periods of "full employment" at the ends of expansions, extensive job reallocation occurs.[3]

A second, more important fact for this paper, is that gross job reallocation is *countercyclical*. Countercyclical reallocation results from asymmetric cyclicality between job creation and destruction: Job destruction increases sharply in recessions but job creation decreases relatively little, or even increases, during recessions. One measure of this asymmetry—the ratio of variances of job destruction and creation, calculated from Table 1—shows that job destruction is 3½ times more volatile than job creation.[4] Some of the cyclical fluctuation in gross job reallocation is associated with big (and transitory) swings in net employment growth, and some is attributable to a lag between the initial occurrence of job destruction and subsequent job creation. But during a typical recession-

[2] Implicit in this definition is that a job is equated with employment, which neglects the fact that there may be unfilled jobs. Also, plant-level employment change is net of internal job flows. Plants create and destroy many jobs that are not measured by the LRD data, so employment change is the number of jobs created net of jobs destroyed. Further, a plant 's net employment change does not fully reflect the actual amount of worker flows through jobs at the plant, because plants may hire and fire workers without changing the stock of jobs.

[3] Evidence on the large magnitude of gross job flows has been known for a long time. Although simultaneous and large rates of job creation and destruction motivate a reconsideration of some aspects of macroeconomic thinking, the magnitude may be irrelevant for business cycle analysis. Thus, in this paper, we are not concerned with explaining why job flows are large; rather, we focus on why they are correlated with the business cycle.

[4] The asymmetry between the variances of job destruction and creation appears to be unique to manufacturing. Ritter (1994) and Foote (1998) provide evidence that job creation varies about the same amount as job destruction in nonmanufacturing industries.

Figure 1
Gross Job Flows in U.S. Manufacturing

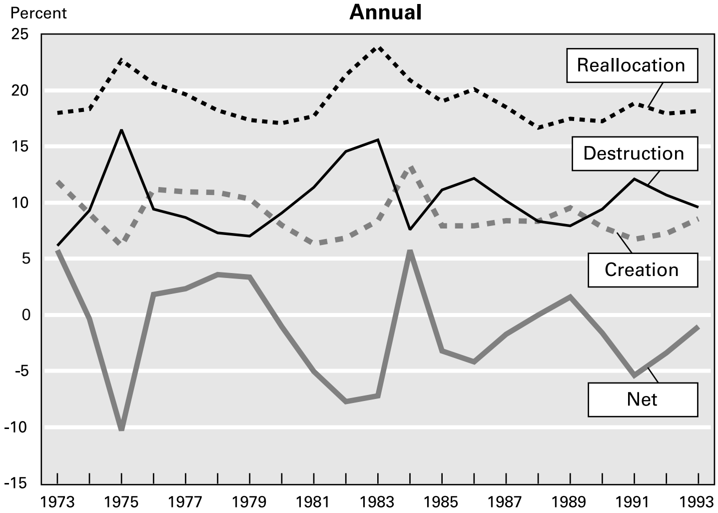

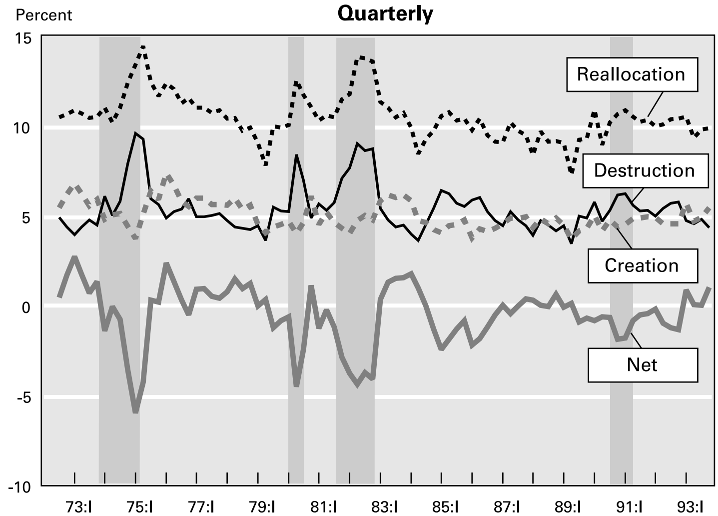

Source: Census/DHS job flows data.

Table 1
Gross Job Flows in U.S. Manufacturing: Summary Statistics

	Rates (%)							
	1972–93				1972–88			
	Annual		Quarterly		Annual		Quarterly	
	Avg.	S.D.[a]	Avg.	S.D.	Avg.	S.D.	Avg.	S.D.
Job Creation (C)	8.8	1.9	5.1	.8	9.1	2.1	5.2	.9
Job Destruction (D)	10.2	2.8	5.5	1.3	10.3	3.1	5.5	1.7
Job Reallocation (R)	19.0	1.9	10.6	1.3	19.4	2.1	10.7	1.6
Net Employment Growth (N)	−1.3	4.4	−.4	1.7	−1.1	4.8	−.3	2.1

	Quarterly Correlation Matrices					
	1972–93			1972–88		
	C	D	R	C	D	R
Destruction	−.29			−.36		
Reallocation	.68	.90		.71	.91	
Net Growth	.31	−.82	−.49	.22	−.83	−.53

[a]Standard deviation.
Source: Authors' calculations using the Census/DHS job flows data, and DHS (1996, Table 2.1).

ary period, which can last several years, the underlying trend rate at which jobs are permanently reallocated across plants, above and beyond net employment growth, increases. This surge in reallocation during recessions and the relative stability of job creation are surprising and intriguing features of the macroeconomy.

Two other basic facts help fill out the picture. One possible explanation for countercyclical reallocation is that the sharp increase in job destruction is merely the result of temporary layoffs, with workers being recalled to their jobs when aggregate demand rebounds. A third basic fact largely rules out this possibility—most job reallocation is *persistent*, or permanent. On average, about three-fourths of job destruction and one-half of job creation persist for two years, or twice the average length of a recession. A fourth basic fact addresses the unevenness of the job flows across plants. A common assumption in traditional macroeconomic models is that plants and industries rise and fall together. Already it can be seen from the high rates of job creation during recessions that this assumption is in trouble. But in addition, the plant-level data show that job flows are *concentrated* primarily in a relatively small number of plants that experience unusually large employment changes. Indeed, two-thirds of job creation and destruction occurs in plants that expand or contract employment by 25 percent or more per year.

In closing this section, let's look at the new gross job flow data containing the 1990–91 recession. Once again, job destruction and reallocation both increased. Actually, the increases occurred prior to the NBER-dated recession because manufacturing employment began contracting in 1989. Compared to earlier recessions, however, job destruction and job reallocation rose much less in the most recent recession, apparently because this recession was less severe, at least in manufacturing. Reallocation in this recession also remained high for an unusually long time. By the end of 1993, job destruction was still above its level five years earlier at the peak of manufacturing employment growth; during this period, job creation was rising steadily. Thus, while the rise in job reallocation was less dramatic, the cumulative amount of reallocation was substantial. More generally, Table 1 demonstrates that the data for 1989–93 do not change the basic time series properties of the earlier data for 1972–88.

Business Cycle Theories with Countercyclical Reallocation

The surprising countercyclicality of job reallocation has sparked an interest in developing macroeconomic theories to explain this phenomenon. Prior to the LRD evidence, most multisectoral theoretical models were not designed to produce countercyclical job reallocation or asynchronous job creation and destruction.[5] Instead, sector-specific shocks generally were assumed to cancel out as the number of sectors increased, so that only an unusual—and, in the view of many, implausible—confluence of shocks could generate a large aggregate effect. The "sectoral shifts" literature provided empirical evidence of countercyclical dispersion in employment growth and argued that reallocation was a driving force behind aggregate economic fluctuations.[6] But formal models of countercyclical reallocation were lacking.

In recent years, new theoretical models have been developed in which job reallocation is an endogenous behavioral response to changes in the surrounding economic environment.[7] Broadly speaking, two types of theories have arisen that explain the cyclical properties of job creation

[5] Rigorous multisector general equilibrium models date back at least to Johnson (1962). More recently, part of the literature focuses on how allocative shocks generate unemployment due to costly and time-consuming search and worker reallocation, including Lucas and Prescott (1974), Hall (1979), Diamond (1981), Mortensen (1982), and Hamilton (1988). Other multisectoral general equilibrium models include Kydland and Prescott (1982), Long and Plosser (1983), Rogerson (1987), Hopenhayn and Rogerson (1993), and Greenwood, MacDonald, and Zhang (1996).

[6] This literature includes Lilien (1982), Black (1987), Loungani (1986), and Davis (1987), and many subsequent articles.

[7] A detailed and comprehensive technical review of these theories is beyond the scope of this paper. See Chapter 5 of DHS (1996), Davis and Haltiwanger (1998), and Hall (1998) for further reading and additional references.

and destruction. One type stresses the role of allocative driving forces in generating business cycle fluctuations and determines that job reallocation causes business cycles. The other type stresses the role of aggregate driving forces and determines that business cycles cause reallocation. In Hall's (1997a) business cycle classification terminology, job reallocation is an "impulse" in allocative theories and job reallocation is a "propagation" and/or "amplification" mechanism in aggregate theories.

Not wanting to lose this simple organizing thought, we must note that both types of theories acknowledge a contemporaneous role for the other driving force as well as intertemporal feedback between the driving forces. But the core issue for this paper is: Can allocative driving forces, by themselves, induce aggregate fluctuations large enough to generate business cycles, or do aggregate driving forces cause essentially all business cycles?

Defining Allocative and Aggregate Driving Forces

Much of the debate over whether allocative or aggregate driving forces cause business cycle fluctuations can be attributed to a fundamental ambiguity about what these forces are. Thus, before discussing alternative theories it is necessary to describe how we define terms and interpret the ambiguities in assessing the debate.

Simply put, aggregate driving forces are economic factors that initially affect firms or consumers in a similar direction and magnitude, whereas allocative driving forces are economic factors that initially affect firms or consumers in a dissimilar direction or magnitude. The most common aggregate driving forces are aggregate demand and aggregate productivity; similarly, the most common allocative driving forces are sectoral demand, sectoral productivity, and relative prices. Demand, measured by income, output, or employment, and relative prices are observable. Productivity, on the other hand, is unobservable and usually estimated as a Solow residual—a concept fraught with measurement difficulties, as illustrated by the Basu article in this volume, and ultimately unsatisfying conceptually.

Ambiguity about allocative versus aggregate driving forces arises in at least three ways. First, changes in aggregate driving forces often affect sectors differentially. A good example is the difference in the cyclical sensitivities of industries producing durable goods and nondurable goods. Another good example is monetary policy, where some firms and consumers are more adversely affected than others by rising interest rates and restricted credit. In these cases, the initial impact of the driving force seems to be both aggregate and allocative in nature.

Second, changes in aggregate demand typically are not spread evenly throughout the economy. A good example is government spending, a classic measure of aggregate demand in traditional macroeconomic models. Government programs, such as military purchases, often are targeted at narrow groups of industries, plants, and regions, as California and Massachusetts found out earlier this decade. Thus, this aggregate driving force seems to be initially allocative in nature, although the concomitant reduction in GDP is likely to affect all agents later as well.

Third, allocative forces such as relative prices appear to have aggregate implications as well. Oil price changes are a quintessential example of an allocative driving force that may not be purely allocative: Only oil producers experience the output price changes (allocative), but virtually every consumer and every other producer buys products or services based on oil (aggregate). Oil price changes also affect the aggregate price, at least in the short run. Exchange rates are another case. Not all industries and plants engage in international trade, but most consumers and plants purchase some foreign goods. Even monetary policy falls into this category, because it is still an open question whether all firms or consumers are adversely affected by interest rate increases.

This discussion suggests most driving forces have both an aggregate and an allocative nature. But most observers tend to view driving forces as being either aggregate or allocative, with the other serving as a means of amplifying and propagating the initial driving force. Empirically, it seems reasonable to conclude that both types of driving forces operate within the common frequencies of macroeconomic data (monthly, quarterly, and annual). Evaluating the nature of the driving force as frequency increases (daily, hourly), however, seems to point more often toward allocative factors as the initial driving force behind fluctuations.

Theories Based on Allocative Driving Forces

Theoretical models with allocative driving forces introduce heterogeneity in workers, plants, capital, products, and the like. Often plants are grouped by common characteristics, such as industry, which form a sector. The allocation of factors of production across plants and sectors is determined primarily by relative prices of goods and factors, relative productivity, and consumers' tastes and preferences for goods. Allocative driving forces cause a change in the desired allocation of factors across plants.

Multisectoral models in which allocative forces drive recessions usually focus on one particular driving force that disrupts the optimality of existing factor allocation. In Davis and Haltiwanger (1990), sectoral productivity shocks govern the exogenous evolutions of high- and low-productivity plants. Often sectoral productivity shocks are amplified and propagated by things such as nonconvexities and complementarities (Cooper and Haltiwanger 1993), sparse input-output structures (Horvath 1998b, 1998c), and uncertainty and learning (Horvath 1998a). In Hamilton (1988), oil price increases alter the efficient allocation of labor (as well as reducing aggregate real income).[8] Other observable allocative forces include real exchange rate fluctuations (Gourinchas 1998) and geographical movements (Blanchard and Katz 1992). Davis, Loungani, and Mahidhara (1997) examine reallocation across geographic regions driven by changes in oil prices and military spending. In all of these models, the driving forces induce desired reallocation across plants and sectors. Actual reallocation ultimately depends on the magnitude, timing, permanence, and uncertainty associated with the driving forces.

In a world without frictions, factor reallocation would occur instantly. But the real world is full of costly and time-consuming frictions that prevent factors from being instantly reassigned to the plant where they are most highly valued. Plants that become unprofitable due to allocative shocks may destroy jobs quickly, but the job creation process often takes more time. Construction of new structures, and delivery and installation of new equipment, may involve significant lead times. And matching displaced workers to the newly created job openings often requires workers and firms to acquire new information, retrain, or shift geographic location. All of these types of frictions typically involve forgone output and a reduction in aggregate activity.

Theories Based on Aggregate Driving Forces

Theoretical models based on aggregate driving forces generally take the aggregate force as a shock given from outside the model, and then focus on explaining how factor allocation changes in response to the shock. Although there are many rich explanations, three basic classes of models have emerged.

One class of models develops direct links between aggregate shocks and factor reallocation, specifically job creation and destruction. In Mortensen (1994) and Mortensen and Pissarides (1994), negative aggregate shocks destroy the profitability of worker-job matches relative to alternative productive opportunities, which must be found through search. In Caballero and Hammour (1994), aggregate shocks reduce the profitability of low-productivity jobs with old capital and cause them to shut down. Caballero and Hammour (1996, 1998) extend their framework to incorporate inefficiencies due to incomplete contracting, financial market imperfections, and suboptimal government policies that amplify and propagate the aggregate shock. Hall's (1997b) model also produces plant shutdowns based on reductions in the expected present value of profits tied to discount rate changes (monetary policy). And Garibaldi (1997) links monetary policy with job flows in a Mortensen-Pissarides

[8] Other articles focusing on the role of oil price shocks include Loungani (1986), Davis and Haltiwanger (1996, 1997), and Bresnahan and Ramey (1993).

framework. Through a variety of complex mechanisms, these models try to match the asymmetry between creation and destruction found in the job flow data.

A second class of models develops an indirect link between aggregate activity and factor reallocation. These models, exemplified by Davis and Haltiwanger (1990), Cooper and Haltiwanger (1993), and Hall (1991), embody a "reallocation timing hypothesis" (RTH). The RTH says that when the level of aggregate demand is low, as in a recession, the opportunity cost of reallocation—forgone output—is also low. Thus, while there is some steady underlying rate of reallocation in the economy, it is optimal to bunch reallocation into periods of low opportunity cost. This intertemporal substitution generates countercyclical reallocation.

The third class of models assumes the presence of microeconomic nonconvexities that produce discrete and infrequent employment adjustment from (S,s) or adjustment hazard-type policy rules. Examples include Bertola and Caballero (1990), Caballero and Engel (1993), Caballero, Engel, and Haltiwanger (1997), and Campbell and Fisher (1998). In these models, the likelihood of plant employment adjustment depends on the gap between actual and desired employment. Aggregate shocks represent small shifts in the average employment gaps but generate large employment changes by pushing some plants over the adjustment threshold. Although not always designed explicitly to explain countercyclical job reallocation, these models are able to produce a sharp spike in job destruction. Furthermore, they provide a nice accounting framework for evaluating the relative importance of aggregate versus allocative shocks. Caballero, Engel, and Haltiwanger (1997) find that aggregate shocks dominate.

Implications for Causality

So, do any of these theories tell us anything about what causes business cycles? Unfortunately, the answer is no. Generally speaking, neither allocative nor aggregate theories of countercyclical reallocation provide much guidance on what *causes* business cycles. Although the reasons differ, both types of theories fail to provide guidance largely because they do not really try to explain *why* business cycles occur. More often than not, fluctuations are assumed to be exogenous shocks—an increasingly tenuous strategy as more and more dimensions of the economy are being explained endogenously in macroeconomic models. Instead, these models address the relatively easier task of explaining *how* business cycles occur.

Theories based on aggregate driving forces do not explain what causes business cycles because they assume the preexistence of business cycles. These theories take as given negative aggregate demand or productivity shocks—that is, recessions—and then try to explain how reallocation occurs in response to the shock. Except for the distinction between aggregate demand and aggregate productivity, these theories really do not depend on the causes of the recession, and thus can treat them as ambiguous, exogenous forces. Nevertheless, the theories do provide some guidance about which broad empirical features might be observed in the data simultaneously with the job reallocation.

Theories based on allocative driving forces have the potential to be more specific about what causes business cycles, but thus far they have not been. One key reason is that they are based on the same kind of unobservable and unsatisfying exogenous productivity shocks on which aggregate theories are based. Another reason is that many of the more promising observable sources of allocative driving forces, such as relative prices, supply linkages, international factors, geographic factors, and product market factors, remain relatively unexplored. Many rich alternative supply-side explanations would arise from a more detailed look at production—a point that resonates with the Basu paper in this volume. Likewise, many demand-side explanations would arise from a more detailed look at the product market environment—we are unaware of any efforts in this regard. Lastly, advocates of allocative driving forces are still a minority, and multisectoral general equilibrium models are technically challenging.

New Evidence on the Nature of Gross Job Flows

This section extends the work of DHS (1996) on the nature of plant-level job flows and presents new tabulations from the LRD. In the first subsection, we examine the cyclical properties of the concentration and persistence of gross job flows, and the relation of these two properties to plant size, in search of clues about what causes job flows, especially job destruction. In the second subsection, we examine the dependence of gross job flows on plants' recent employment growth, in an effort to better understand the nature of plant-level employment adjustment and the possibility that it is governed by nonconvexities.

The Cyclicality of Job Flows by Concentration, Persistence, and Size

In describing the nature of plant-level job flows, DHS (1996, pp. 146–49) paint a picture in which old, large plants play a central role in recessions. "Job flow dynamics in good times are dominated by the creation and destruction of jobs among relatively young and small plants. . . . During recessions, older and larger plants experience sharply higher job destruction rates, so that their contribution to the job and worker reallocation process rises. This time of intense job destruction by older and larger plants coincides with the rise in layoff unemployment,

Figure 2
Percentage of Total Manufacturing Jobs Destroyed, by Degree of Plant's Contraction

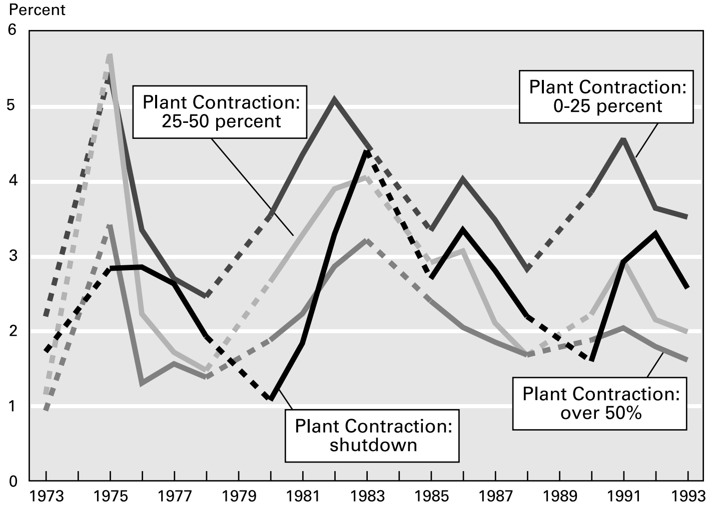

Note: Dotted lines represent interpolated numbers.
Source: Authors' computations from Longitudinal Research Database.

especially among prime-age workers." This subsection refines and sharpens these ideas.

Although an increase in the number of manufacturing plants undergoing large decreases in employment occurs in recessions, this is not primarily a phenomenon of plants closing. As Figure 2 shows, the percentage of jobs destroyed in plant shutdowns increases during recessions, but so does job destruction due to much more modest reductions in plants' employment.[9] The figure shows a strong countercyclical pattern to the job destruction at plants undergoing all four degrees of contraction. Job destruction in plant closings tends to appear somewhat later in recessions than does job destruction due to more modest degrees of contraction, however.

An interesting relationship exists between the size of a plant and the degree of its contraction: Small plants tend to destroy jobs in much more concentrated contractions than do large plants. This result is documented

[9] Because new panels were introduced in the Annual Survey of Manufactures (ASM) in 1974, 1979, 1984, and 1989, job destruction rates were not calculated for those years. The dashed lines in the figure interpolate between the adjacent years for which data are available.

in Figure 3, which shows the concentration of job destruction by plants' average level of employment (measured over all years in which the plant appears in the LRD) separately for recession and non-recession years.[10] About one-third of the job destruction in plants with less than 50 employees is due to plant shutdowns, while less than 15 percent of the job destruction in plants with more than 1,000 employees is due to shutdowns. Interestingly, for all four employment size classes shown in the figure, shutdowns account for a somewhat smaller proportion of job destruction in recession years than they do in non-recession years. Job destruction due to plant closings is countercyclical, but less so than job destruction due to employment reductions in continuing plants.

In general, larger plants exhibit a higher percentage of job destruction occurring in relatively modest employment contractions. Half of the jobs destroyed by plants in the largest employment size class were lost in employment reductions of 25 percent or less, while only 20 percent of the jobs destroyed by plants in the smallest employment size class were in this range.

These differences between the ways large and small plants destroy jobs are potentially quite important. Plant shutdowns may be determined by processes that are substantively different from those that produce more moderate reductions in a plant's staffing. One indication of this, shown in Table 2, is that the persistence of newly destroyed jobs increases with the degree of the plant's contraction.[11] Job destruction due to plant shutdowns is nearly always permanent, while about one-half of the jobs destroyed in contractions of less than 25 percent will be restored within five years.[12] Table 2 reveals that although the one-year persistence rates associated with employment reductions of any magnitude short of a full shutdown vary relatively little with the degree of the plant's contraction, the differences across the concentration classes increase as one examines persistence over longer time horizons.

The average persistence rates in Table 2 hide considerable and informative heterogeneity in plant-level persistence rates. We investigate the heterogeneity of persistence by plotting the distributions of plant-

[10] We define 1974–75, 1981–83, and 1990–91 as recession years, and other years in the 1973–93 period as non-recession years, when using annual data. However, our calculations exclude 1974, 1979, 1984, and 1989 because new ASM panels were introduced in those years.

[11] The persistence rates shown in Tables 2 and 3, and Figures 4 and 8, are based on tabulations from the LRD. Persistence rates were calculated in each year for plants that were included in the LRD sample for five years beyond that year (except in the case of plant closings, where closed plants are assumed to remain closed if they do not reappear in the LRD in a future year). The rates were then aggregated over plants and years, weighting each plant-year observation by the number of jobs destroyed or created times the sample weight.

[12] We extend the DHS persistence measure horizon from two to five years to see whether some of the "permanent" plant-level employment adjustment in the short run might be "transitory" in the medium term. Changing capital stocks, production technologies, and product lines could easily take more than two years.

Figure 3
Concentration of Job Destruction, by Size of Plant

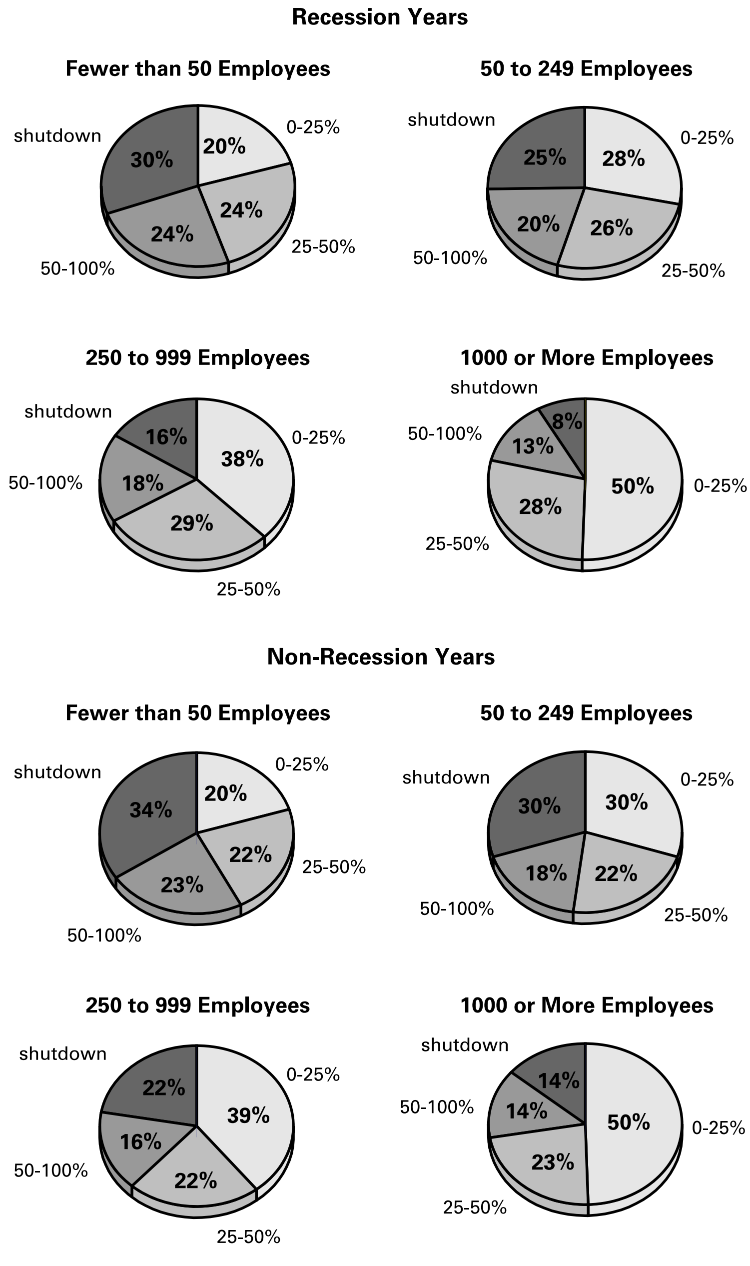

Source: Authors' computations from Longitudinal Research Database.

Table 2
Job Destruction Persistence Rates, by Degree of Plant's Contraction, 1973 to 1988
Percent

Plant Contraction	One-Year Horizon	Two-Year Horizon	Five-Year Horizon
0–25 Percent	75	64	51
25–50 Percent	79	69	58
Over 50 Percent	80	74	68
Shutdown	99	98	97

Source: Authors' calculations using Longitudinal Research Database.

level job destruction persistence rates in Figure 4. Because we are primarily interested in the extent to which job flows are permanent versus transitory, we focus on a very simple discrete distribution with three cells: permanent (85 to 100 percent), intermediate (15 to 85 percent), and transitory (0 to 15 percent).[13] Figure 4 shows that plant-level destruction persistence tends to be bimodal, with most mass clustered in the tails near 0 (transitory) and 100 (permanent), rather than being bell-shaped, with most mass concentrated near the mean persistence rate.[14] As the time horizon increases, the proportion of plants in the permanent range falls and the proportion in the transitory range increases. For example, among plants contracting less than 25 percent, 64 percent of job destruction is permanent and only 16 percent is transitory after one year. But after five years, only about 40 percent is permanent and another 40 percent is transitory. In future work, we plan to develop a better understanding of why some plants permanently destroy jobs and others temporarily destroy jobs.

The persistence of job destruction associated with moderate degrees of plant employment reductions (0 to 25 percent) follows an interesting pattern over the business cycle. As Figure 5 shows, one- and two-year persistence rates tend to increase just before a recession begins and then drop in the recession's final year. This suggests that the persistence of jobs lost in moderate plant contractions may be very sensitive to aggregate demand. A drop in aggregate demand will increase the persistence of job

[13] Because plants' persistence rates are weighted by the number of jobs destroyed in computing the distribution, the frequencies show the proportion of destroyed jobs located in plants with persistence rates in each of the three categories.

[14] The actual probability density function of persistence rates is determined by the driving process for plant-level employment growth rates, which we do not specify but know from the evidence in DHS is bell-shaped (though definitely not normal). A more detailed and careful treatment of this issue is clearly warranted but beyond the scope of this paper. However, we note that the observed bimodal distribution is markedly different from even a uniform distribution, in which we would expect to see 70 percent of persistence in the intermediate range (15 to 85 percent). In fact, however, a much smaller proportion of destruction falls in this range.

Figure 4
Distribution of Job Destruction Persistence Rates, by Degree of Plant's Contraction

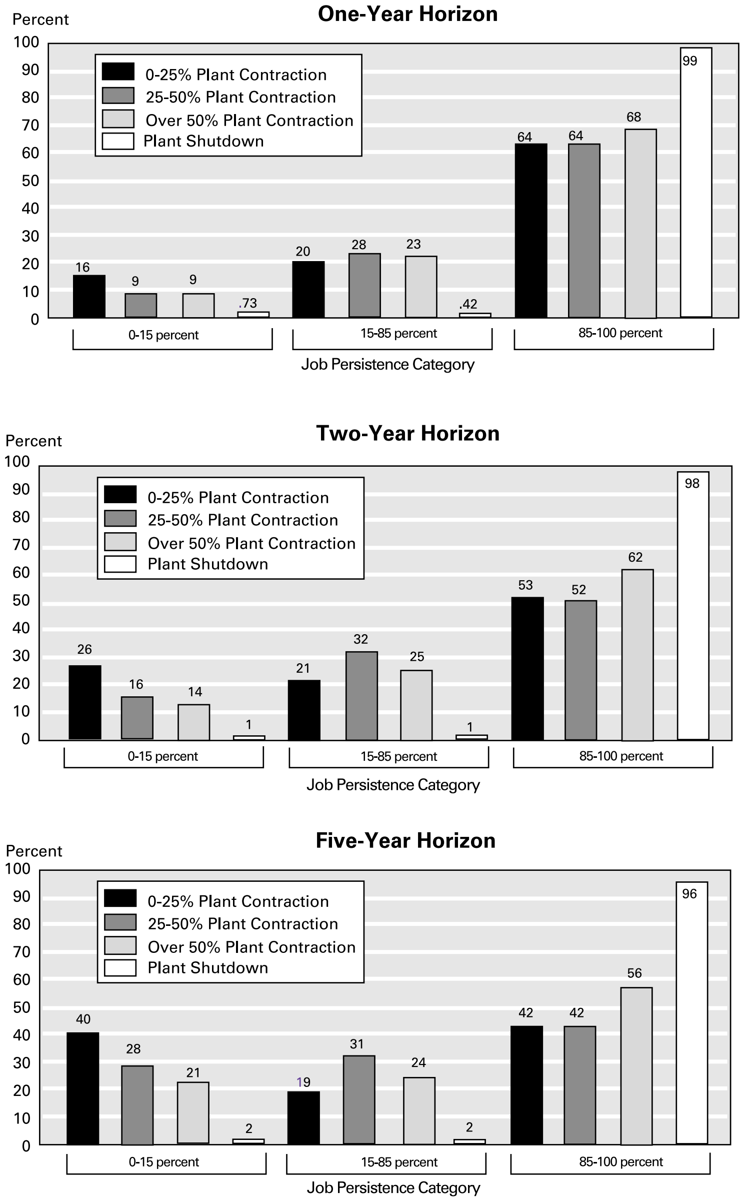

Source: Authors' computations from Longitudinal Research Database.

Figure 5
Persistence of Job Destruction, by Degree of Plant's Contraction

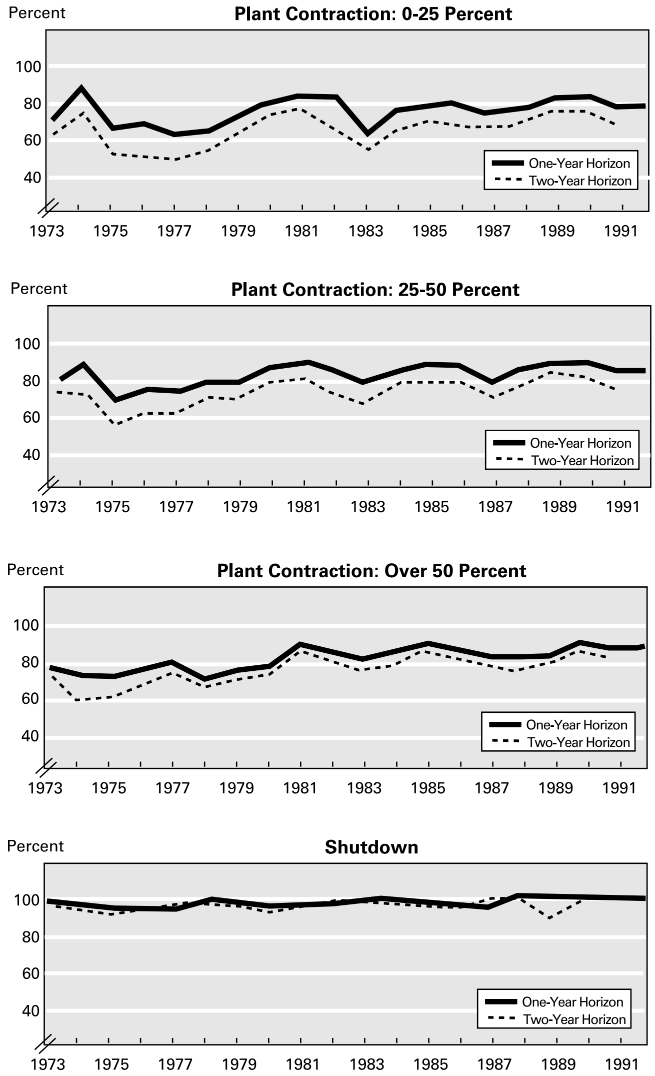

Source: Authors' computations from Longitudinal Research Database.

Figure 6
Percentage of Total Manufacturing Jobs Created, by Degree of Plant's Expansion

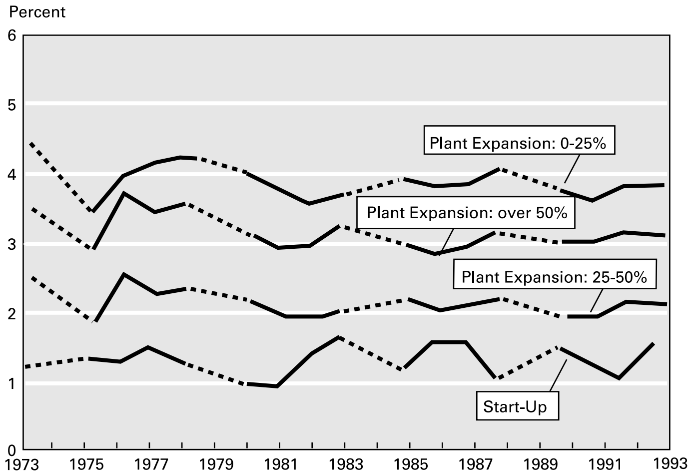

Note: Dotted lines represent interpolated numbers.
Source: Authors' computations from Longitudinal Research Database.

destruction, while an increase in demand will make the job destruction more transitory. This pattern also shows up, although to a lesser degree, for jobs destroyed in plants contracting between 25 and 50 percent. For plants contracting over 50 percent, however, the main pattern seems to be an upward drift of persistence rates over time, and plant shutdowns are nearly always permanent in all years.

Turning to the job creation side of employment reallocation, Figure 6 shows the pattern over time of the percentage of total manufacturing jobs created each year, according to the degree of plant expansion. The first three concentration classes, existing plants, display job creation patterns that are very procyclical during the 1970s, but less so in the 1980s and 1990s. However, new plants (start-ups) follow a pattern that is acyclic or possibly even countercyclic. One possibility is that much of the job creation in the moderate expansion classes following a recession is due to the restoration of jobs temporarily destroyed during the recession. Recall that the persistence rates of job destruction are relatively low for moderate degrees of job destruction, and drop toward the end of recessions. This would lead one to expect an increase in job creation rates in the moderate expansion classes immediately following recessions.

As with job destruction, there are marked differences by plant size in the concentration of job creation. As Figure 7 shows, a much larger

Figure 7
Concentration of Job Creation, by Size of Plant

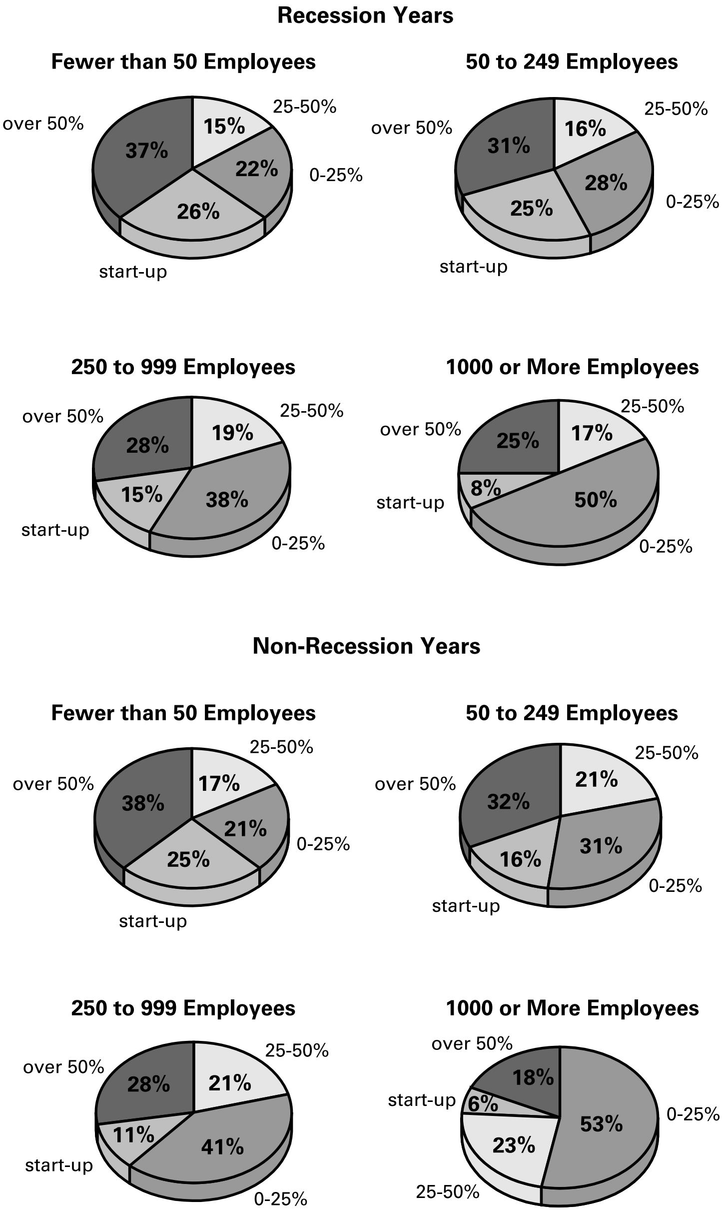

Source: Authors' computations from Longitudinal Research Database.

Table 3
Job Creation Persistence Rates, by Degree of Plant's Expansion, 1973 to 1988
Percent

Plant Expansion	One-Year Horizon	Two-Year Horizon	Five-Year Horizon
0–25 Percent	66	50	28
25–50 Percent	71	55	33
Over 50 Percent	72	59	39
Start-Up	74	62	44

Source: Authors' calculations using Longitudinal Research Database.

proportion of the jobs created by small plants is in start-ups or created through very high rates of expansion than is the case for large plants. These patterns are similar to those we documented earlier for the concentration of job destruction. Small plants tend to adjust employment through start-ups/shutdowns and truly massive percentage changes in their employment levels, while large plants tend to make smaller percentage adjustments both when expanding and when contracting. Also note that for large plants, the higher proportion of job creation through start-ups or very large (more than 50 percent) increases in employment is higher in recession years than in non-recession years. Recall that we found job destruction in large plants is more concentrated in shutdowns during non-recession years than during recessions. The job creation and destruction patterns together suggest that start-ups, shutdowns, and other massive employment changes in large plants may tend to be the result of long-run planning, while smaller percentage employment changes may be caused by fluctuations in product demand.

As with job destruction, the persistence of job creation increases with the degree of the plant's employment change (as shown in Table 3): Jobs created in start-ups are most persistent, while jobs created in expansions of less than 25 percent are least persistent. Like the job destruction persistence rates, job creation persistence rates tend to be concentrated in the tails of the distribution (as shown in Figure 8). Many fewer plants have persistence rates in the intermediate (15 to 85 percent) range than one would expect if the persistence rates were uniformly distributed. As the horizon over which persistence is measured increases from one to five years, the proportion of plants in the high persistence (85 percent or greater) category shrinks and the proportion in the low persistence category (less than 15 percent) increases; the proportion in the intermediate class remains relatively stable. Further research is needed to understand why some plants create long-lived jobs, while others increase their employment only temporarily.

Figure 9 shows that the persistence (at the one- and two-year horizons) of jobs created by plant start-ups is largely acyclic, but the

Figure 8
Distribution of Job Creation Persistence Rates, by Degree of Plant's Expansion

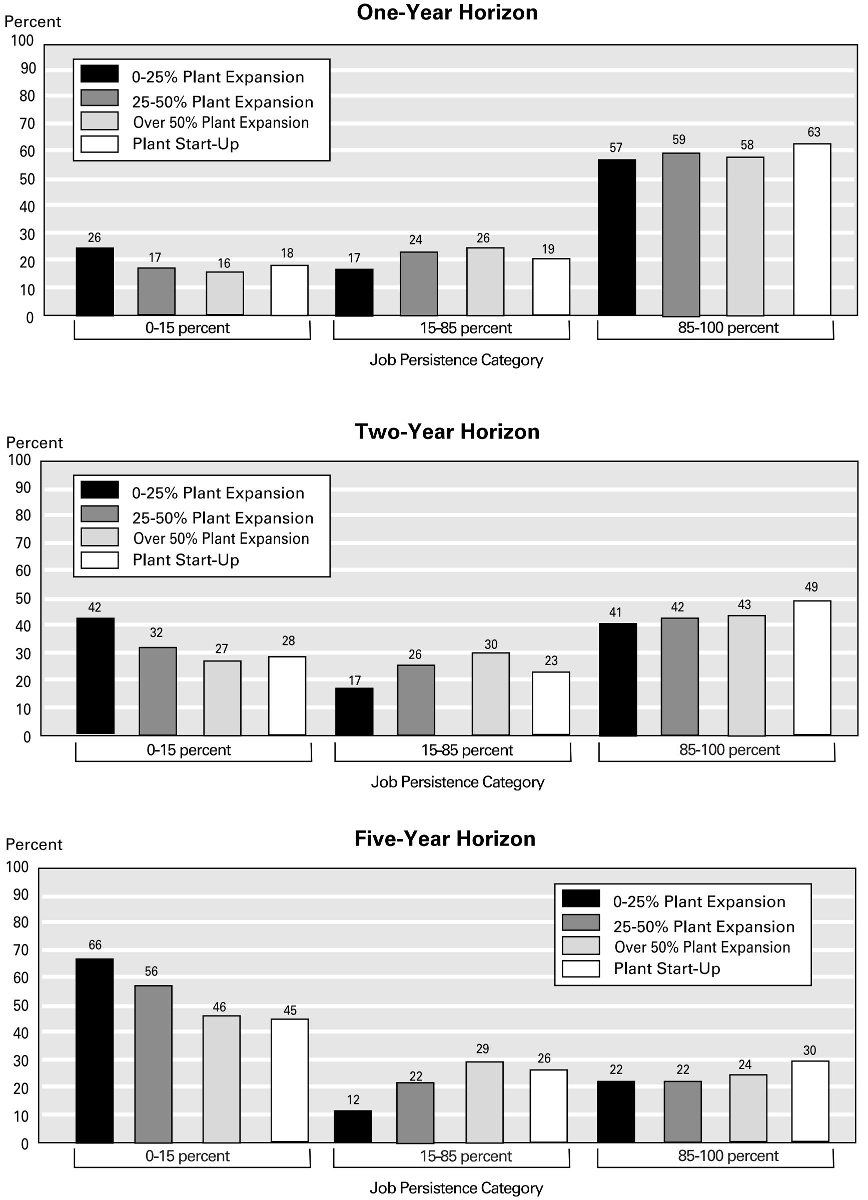

Source: Authors' computations from Longitudinal Research Database.

Figure 9
Persistence of Job Creation, by Degree of Plant's Expansion

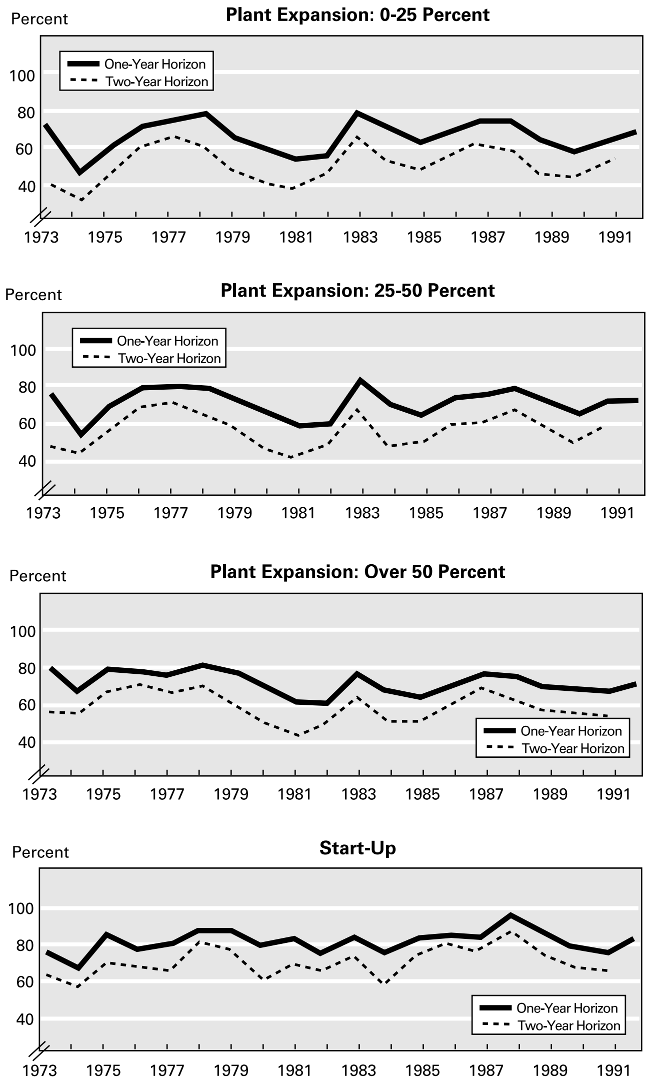

Source: Authors' computations from Longitudinal Research Database.

persistence rates of jobs created in moderate percentage expansions of plants' employment follow a pattern of decreasing early in a recession and then increasing before the recession ends. Jobs created relatively early in a recession seem to be at risk of being eliminated because the plant is contracting before the recession ends, while jobs created at the end of a recession are more likely to survive because of the more favorable demand conditions coming in future years.

Implications and Extensions. Putting together the facts on the concentration and persistence of job destruction yields a number of interesting enhancements to the basic DHS picture of plant job flows. Although larger plants exhibit a greater cyclical asymmetry between job creation and destruction, large plant job flows are likely to be much more moderate than small plant job flows. Small plant job flows are much more likely to be concentrated in shutdowns or massive contractions and expansions. Because higher job flow rates tend to be much more persistent, small plant job flows tend to be more permanent and large plant job flows more transitory. High job flow rates tend to be less cyclical than moderate job flow rates, and the predominance of large plant job flows (especially destruction) during recessions occurs in relatively moderate (less than 50 percent) contraction and expansion classes. Further, these moderate expansions and contractions are more likely to be transitory, especially in the medium term. To summarize, while large plants dominate movements in job flows during recessions, their flows are smaller and less persistent than inferred from DHS.

To the extent that small plants tend to be part of smaller firms than are large plants, the greater concentration of small plants' job flows in sharp expansions and contractions could be due to credit market imperfections. Large firms may reduce employment and output because of a transitory decrease in product demand or increase in costs, but in many cases they eventually resume operating at their previous level. Small firms may instead close down entirely under the same circumstances because of lack of credit and depletion of internal funds. The fact that job destruction due to plant closings tends to occur relatively late in recessions adds further credence to this hypothesis.

Small plants might also employ workers with less organization-specific human capital than workers in large plants. In this case, small plants might be more likely to close or permanently downsize owing to the smaller capital loss in dissolving existing worker-firm matches. Another possible explanation is that large plants are more diversified (or are parts of more diversified firms) than are small plants. A reallocative shock may prompt a large plant to temporarily decrease employment and output while it retools to produce new products, while a small plant might find that its optimal response to the shock is to shut down. Yet another possibility is that the differences in the ways in which small and large plants destroy jobs are partly due to differences in the distribution of plant sizes across industries. For example, large plants may be more likely to be part of cyclically sensitive durable goods industries than are small plants. Clearly, more research is needed in order to better understand the differences in the job destruction patterns of small and large plants.

Plant-Level Time Series Characteristics of Employment Growth

Aside from the facts that plant-level employment adjustments are often quite large and that flows of newly created and destroyed jobs tend to persist for at least two years, little is known about the plant-specific time series properties of employment growth. Much of the literature, such as Hammermesh (1989), Caballero and Engel (1993), and Caballero, Engel, and Haltiwanger (1997), views the large, persistent plant-level employment adjustment as arising from microeconomic nonconvexities that also cause adjustments to be infrequent. This view contrasts sharply with the standard macroeconomic treatment of employment being subject to adjustment costs that induce serial correlation and partial adjustment. The contrast of these two views motivates a deeper look at the plant-specific time series properties of employment growth.[15]

We conduct a simple exercise to obtain an approximate estimate of the historical time dependence in plant-level employment growth. Unfortunately, the short time series dimension of the LRD (at most 20 years, much less for most plants) inhibits direct estimation of plant-specific autocorrelation properties. Instead, for each plant we calculate the average employment growth rate during the prior two years and tabulate the gross job flows by eight prior employment growth categories, plus a missing category.[16] This latter group contains primarily plant start-ups and, in certain years, new panel entrants.

Table 4 reports gross job flow rates and shares by prior employment

[15] Without a formal model of the process of plant-level employment adjustment, it is impossible to know what the exact nature of employment time dependence is. If plant-level employment is stationary, microeconomic nonconvexities suggest that time dependence should be minimal and employment growth prior to large employment adjustments should be quite small (averaging close to zero). However, if plant-level employment contains a deterministic or stochastic trend, nonconvex employment adjustment could exhibit time dependence. But trending plant employment—especially upward trending—is hard to conceive without joint adjustment of capital (buildings and equipment). Abel and Eberly (1997) argue that nonconvexities in capital adjustment can cause large, infrequent employment adjustments even in the presence of convex labor adjustment. Cooper and Haltiwanger (1998) find that a model with a mix of convex and nonconvex employment adjustment fits the plant-level data best.

[16] Results are similar using the previous one, two, or three years of prior employment growth, but longer periods entail more missing values. Two years smooths out the "regression to the mean" behavior associated with transitory employment changes while providing a sense of the trend in employment growth and has fewer missing observations than three years.

Table 4
Gross Job Flows in U.S. Manufacturing, by Prior Employment Growth

	Rates (%)				Shares (%)[a]		
Prior Growth[b]	C	D	R	N	C	D	E
Recession years:							
−50% or less	66.3	18.8	85.1	47.4	2.0	.3	.2
−50% to −25%	14.7	19.9	34.6	−5.2	6.5	4.8	3.2
−25% to −10%	5.5	15.1	20.6	−9.6	6.6	9.4	8.3
−10% to 0%	3.5	11.8	15.4	−8.3	12.0	20.9	23.5
0% to 10%	3.7	10.1	13.8	−6.4	13.4	19.6	26.0
10% to 25%	5.3	12.9	18.2	−7.5	8.2	10.5	10.9
25% to 100%	7.6	19.1	26.7	−11.6	5.0	6.7	4.7
More than 100%	12.2	18.2	30.4	−5.9	7.1	7.5	6.2
Missing	25.9	14.3	40.3	11.6	38.8	20.1	17.0
Expansion years:							
−50% or less	62.2	20.6	82.7	41.6	1.8	.6	.3
−50% to −25%	17.8	16.7	34.5	1.1	7.3	5.9	3.7
−25% to −10%	8.8	10.4	19.2	−1.6	9.7	10.2	9.8
−10% to 0%	5.4	8.0	13.4	−2.6	15.9	22.1	26.0
0% to 10%	5.1	7.1	12.2	−2.0	14.9	18.9	25.5
10% to 25%	7.0	9.7	16.7	−2.6	9.7	11.9	12.0
25% to 100%	8.9	14.8	26.8	−5.9	5.6	8.4	5.5
More than 100%	14.3	16.6	30.9	−2.2	3.1	3.3	1.9
Missing	28.1	11.9	40.1	16.2	32.0	18.6	15.2

[a] Totals may not sum to 100 because of rounding.

[b] Prior employment growth rate classes represent the average plant-level total employment growth during the preceding two years.

Note: C = Job Creation, D = Job Destruction, R = Job Reallocation, N = Net Employment Growth, E = Employment.

Source: Authors' calculations using Longitudinal Research Database.

growth for recession and expansion years. The table reveals a clear U-shaped pattern between gross job flows and prior employment growth: Plants with higher absolute rates of prior employment growth have higher job flow rates. The pattern is slightly different between creation and destruction. Job creation rates are flatter in the moderate prior growth classes, then rise sharply in the higher absolute growth classes. Particularly notable is the enormous rate of job creation for plants that had been radically contracting (–50 percent or less). Job destruction rates are more V-shaped, and destruction is not notably higher in radically contracting plants.

The table also reveals a strong negative correlation between net employment growth and prior employment growth. Excluding the missing category, the net employment growth rate tends to decline as prior employment growth rises. Bearing in mind that total manufacturing employment and average plant employment size were declining on average in the LRD sample during this period, it is still quite interesting that only two extreme categories experienced above-average net employment growth. The missing category has strong net employment growth because it is dominated by plant start-ups (including first-time plant births), which tend to grow relatively rapidly in their early years—see especially the job creation rates. But the only other category with positive net employment growth is the most rapidly contracting plants. In fact, net employment growth for these plants is so large that it goes a long way toward eliminating the earlier decline.[17]

The last three columns of the table provide perspective on the importance of the variation in job flow rates across growth rate categories, by reporting the shares of job creation and destruction and employment. About half of all plants had very moderate prior employment growth (–10 percent to 10 percent). The shares of employment in the large (in absolute value) classes are considerably smaller, but these categories account for disproportionately high shares of job flows.

A final result from the table pertains to the cyclical pattern of time dependence. The U-shaped pattern of job flows is roughly the same in recession and expansion years, indicating that the result is not driven by business cycle effects. The U-shaped pattern in job destruction is much flatter across categories in recessions, but the pattern in job creation is somewhat deeper (except in the categories of very high absolute growth rates). Perhaps most striking is the cyclical change in the shares of job creation, destruction, and employment in the highest prior growth class (more than 100 percent), which all increase dramatically in recessions. Two kinds of plants may be in this category. Some plants might overexpand employment leading up to a recession, perhaps as a result of forecast errors, and destroy jobs when their expectations are updated. Other plants might be growing rapidly and, when their demand remains high during the recession, are encouraged to expand further. The relative acyclicality of job creation in startups shown earlier may also help explain this result, particularly if start-ups in larger plants are spread gradually over several years.

Implications and Extensions. The simple tabulations in this section imply that theories of employment adjustment must exhibit historical time dependence in gross job flow behavior to fit the data. Job flow rates are systematically related to prior employment growth, with higher job flow rates being associated with plants whose employment previously

[17] One possible explanation for the U-shaped pattern in rates is that it is driven by differences in plant size, age, average wage, and the like. Larger, older, higher-wage plants have markedly lower rates of gross job flows, on average, so if they tended to have smaller absolute prior employment growth rates as well, the plant characteristics would explain the pattern. However, calculations not reported here indicate that the same general U-shaped pattern appears across all size classes.

has been changing dramatically. Also, the more a plant's employment grew recently, the more likely it is to decline a lot currently. Theories based on convex adjustment exhibit time dependence inherently; theories based on nonconvex adjustment must introduce time dependence in a sensible way.

Combining the results on prior employment growth with the results in the previous subsection paints an interesting but complex picture of employment adjustment. Recall that most plant-level job flows, especially extreme expansions and contractions, are persistent. Consequently, high rates of creation and destruction will show up as high rates of employment expansion or contraction for the next two years. Thus, large employment adjustments appear to be a common, rather than unusual, feature for plants that exhibit high job flow rates. In other words, high job flow rates seem to come in batches or are an inherent feature of certain plants.

Consecutive large rates of annual job flows do not necessarily indicate plant employment is trending, though. The data show roughly equal evidence of four basic patterns of plant employment adjustment: (1) rapid expansion for two years, then rapid expansion for two more; (2) rapid expansion for two years, then rapid contraction for two years; (3) rapid contraction for two years, then rapid contraction for two more; and (4) rapid contraction for two years, then rapid expansion for two years.

What does this evidence suggest about convexity of employment adjustment? If plant employment is stationary, the time dependence favors convex adjustment. If plant employment is trending, cases (1) and (3) could reflect nonconvex adjustment because trending employment will hit upper and lower adjustment thresholds that induce large adjustments. This argument is made by Foote (1998), based on the fact that total manufacturing employment was trending downward during the LRD period. However, even the prior employment growth results in this section do not provide concrete evidence on the relationship between the employment trends of the specific plants that exhibited large employment adjustments. The real empirical challenge for theories of nonconvex employment adjustment are cases (2) and (4), where large positive and negative employment trends are suddenly reversed by high destruction and creation rates, respectively.

Large job flow rates that reverse large employment trends, such as high job destruction in rapidly expanding plants, suggest either a deliberate, large, and transitory employment adjustment, such as temporary layoffs or retooling, or employment adjustment associated with forecast errors, information surprises, and the like. Consecutive large employment adjustments in the same direction, such as high job destruction in rapidly contracting plants, indicates that large, permanent employment adjustments are spread out over several years. If so, this phenomenon may imply that marginal employment adjustment costs are rising, or perhaps that uncertainty and learning may slow the adjustment of employment to its desired level. Clearly more detailed investigation is required on these issues.

New Evidence on the Connection between Job Reallocation and the Business Cycle

This section presents new empirical evidence on the connection between job reallocation and business cycles and the causal ordering between the two. The empirical exercises were designed to address some of the goals for expanding our understanding of job reallocation as described in the introduction. Although not formal tests of modern theories of reallocation, the exercises yield new evidence that is relevant for many of the theories. Analyses in this section are conducted at the industry and sectoral levels using Census/DHS gross job flows and related data.

Time Series Causality

A logical and simple first step is to look for evidence of standard time series causality between reallocation and business cycles. With the usual strong caveats about causality in mind, we conduct two types of exercises designed to detect allocative effects on reallocation. First, we look for differences in gross job flow behavior among sectors defined by plant characteristics. To keep the results manageable, plant characteristics are simplified into two-way classifications, such as large and small. Despite the high degree of parsimony, important differences emerge. The goal is to discover differences among plants in gross job flows that may yield clues about the causes of recessions. Second, we look for lead-lag relationships between sectoral job flows and aggregate activity that may reflect causality.

At this point, we want to temporarily narrow our focus to job destruction. Figure 10 plots the gross job flows for the entire postwar period (see Davis and Haltiwanger (1995) for data construction). A sharp increase in job destruction is a regular feature of every postwar recession, not just the LRD period. Although job creation is more variable prior to 1972, the asymmetry between destruction and creation variances, and the countercyclicality of total job reallocation, are evident in this period as well. Furthermore, the relatively modest decline in job creation during recessions is almost always followed by a surge in job creation. Thus, the tendency for job destruction to change more and earlier during recessionary periods than job creation suggests that whatever causes job destruction may be what causes recessions.

Table 5 reports cyclical characteristics of job destruction by plant-

Figure 10
Quarterly Gross Job Flows in Manufacturing

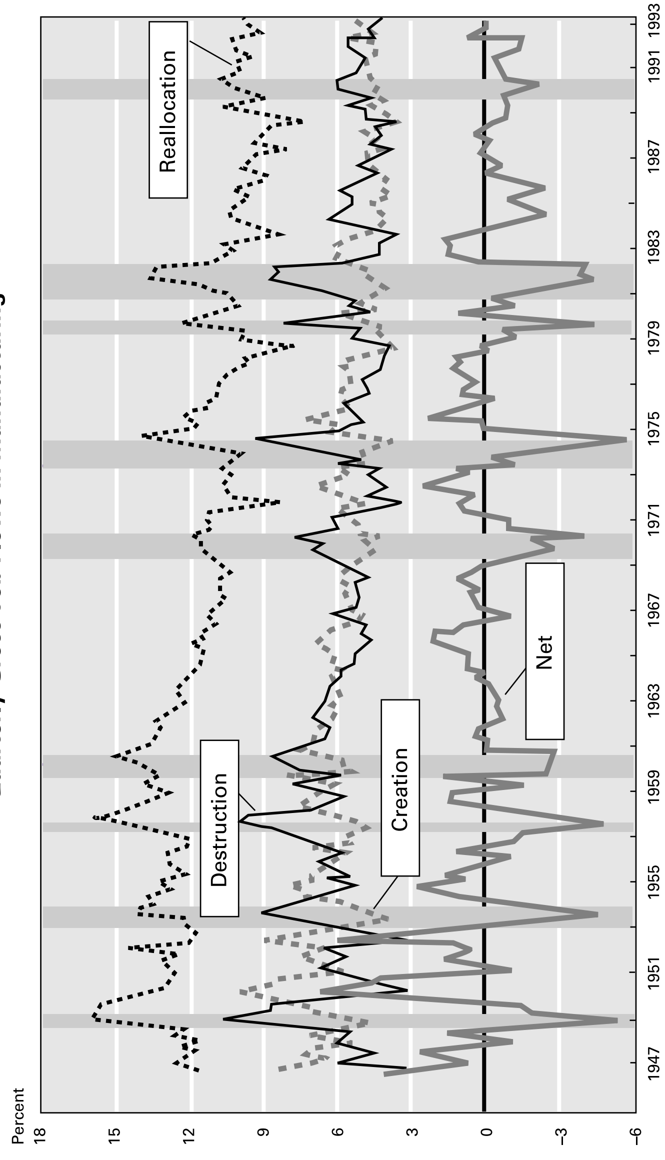

Source: Census/ DHS job flows data.

Table 5
Quarterly Characteristics of Job Destruction over the Business Cycle[a]

Plant classification	Variance Ratio[b]	Average Job Destruction in:		Business Cycle Change[c]		
		Expansions	Recessions	1975	1980	1982
Total Manufacturing	3.4	5.1	7.2	5.7	3.6	3.2
Nondurable Goods Industries	2.0	5.4	6.7	3.9	1.6	2.3
Durable Goods Industries	3.4	5.1	7.8	6.9	5.2	4.0
East Region	3.7	4.8	6.4	4.9	2.6	2.8
North Region	3.7	4.9	7.8	7.7	6.6	3.9
South Region	2.8	4.7	6.6	5.5	2.7	3.6
West Region	1.4	6.7	8.8	5.4	3.4	3.3
Small Plants	2.2	6.2	8.4	4.5	3.2	3.6
Large Plants	3.1	4.2	6.4	6.9	3.9	3.0
Young Plants	1.3	6.1	8.1	5.5	3.4	2.9
Old Plants	4.1	4.8	6.9	5.8	3.6	3.3
Specialized Plants	3.4	5.4	7.9	6.2	3.7	3.0
Diversified Plants	2.5	4.5	6.5	5.2	3.5	3.3
Low-Wage Plants	2.7	5.6	7.6	5.3	3.0	3.1
High-Wage Plants	3.8	4.2	6.8	6.2	4.7	3.9
Low Energy-Intensity Plants	3.3	5.0	7.0	4.9	3.5	3.0
High Energy-Intensity Plants	3.2	5.1	7.4	6.4	3.6	3.4
Low Capital-Intensity Plants	3.5	5.8	7.6	5.2	3.1	3.4
High Capital-Intensity Plants	3.3	4.4	6.8	6.1	4.2	3.4

[a] Based on quarterly data from 1972 to 1988.

[b] Ratio of the variance of job destruction to the variance of job creation.

[c] The increase in job destruction from the average value in the four quarters before the recession to the maximum value during the recession.

Source: Authors' calculations based on data from Davis and Haltiwanger (1996).

characteristic sectors. The table includes three types of statistics: the ratio of the variance of job destruction to the variance of job creation; the average levels of job destruction in expansions versus recessions; and the business cycle change in job destruction during recessions.[18] The variance ratios reveal considerable heterogeneity in the extent of countercyclical reallocation among plants—the larger the ratio, the greater the asymmetry between destruction and creation, and hence the more reallocation. Certain types of plants exhibit relatively high variance ratios and reallocation: durable goods, eastern and northern, large, old, high-wage, and

[18] The business cycle change in job destruction is the increase in destruction from the business cycle peak (average rate during the peak and previous three quarters) to the maximum value during the recession.

specialized. Ironically, plants with high variance ratios tend to have lower average rates of job destruction in both expansions and recessions. But during recessions, these types of plants experience larger increases in job destruction than their counterparts. To summarize, whatever causes recessions must account for the fact that these types of plants experience disproportionate increases in job destruction.

Table 5 shows that despite the asymmetry between creation and destruction rates, *all* types of plants experience substantial increases in job destruction rates during recessions. Business cycle changes in all sectors are broadly similar to total manufacturing, except for the industry sectors. On average, job destruction in durable goods industries rises twice as much as in nondurable goods industries, presumably because the demand for individual durable goods is inherently lumpy and more sensitive to credit conditions and the business cycle. Note that plants in durable goods industries also tend to be larger, older, higher-wage, and disproportionately located in the east and north. Together these facts beg the question: Do the plant characteristics arise because of industry differences (demand) or do industry differences arise because of plant production characteristics (supply)? Most theories of countercyclical reallocation are predicated on the latter, but the data hint at the former. However, the industrial differences may apply only to the transitory, rather than the permanent, component of destruction. More detailed disaggregation is needed to sort out these issues.

One can see further the similarity of job flows among sectors during recessions in Figure 11. Despite differences in relative variances, destruction (creation) rates clearly rise (fall) in all sectors during recessions. The magnitudes of change differ across sectors, but the time series patterns are remarkably similar. No sector averts the recession on either the creation or destruction margin. In contrast, Figure 12 plots gross job flow shares and highlights some important differences across sectors. The share of job destruction (creation) tends to rise (fall) during recessions for large and durable goods plants. One obvious potential explanation for this result is that these types of plants engage more in temporary layoffs. Another interesting difference, which is secular but timed around the 1981–82 recession, is a permanent increase in the share of job destruction and creation in old and, to a lesser extent, high-wage plants.

Two results stand out thus far. First, countercyclical reallocation prevails in all of the sectors. But, second, no clear sectoral differences arise to indicate that allocative driving forces hit some sectors but not others during recessions. Instead, the sectoral job flow characteristics resemble the early empirical results on dispersion in employment growth rates, where quantitative rather than qualitative differences could easily be explained by differing cyclical sensitivity of sectors to aggregate driving forces. However, these sectors are highly aggregated and may be masking clearer differences at more disaggregated levels. Finer sectoral classifications might identify sectoral differences that would more clearly suggest evidence of allocative driving forces. Based on the results so far, disaggregating by industry, interacted with plant characteristics, seems to be particularly promising for finding substantial sectoral differences.

Our second exercise looks for lead-lag relationships in the spirit of traditional econometric causality tests, as pioneered by Granger (1969) and Sims (1972). Consider first graphic evidence from the so-called butterfly plots in Figures 13a, 13b, and 13c for each of the three NBER official recessions during the 1972–88 period. The butterfly plots show sectoral job flows for one year before and after each business cycle peak (solid vertical line).[19] These plots are designed to detect cases where job destruction increases before recessions and hence might be said to have caused the recession. Similarly, these plots can detect lead-lag relationships between the sectoral rates of job destruction.

These plots show that destruction begins rising in virtually every sector prior to recessions. Although the spike in job destruction occurs after the recession begins, an unmistakable and substantial (often several percentage points) increase in job destruction occurs before all three recessions. Interestingly, job creation does not tend to change much before recessions except in 1981–82, when it declines a bit for some sectors. Note, however, that the plots also show that the leading nature of job destruction is generally similar across sectors. Except for a mild blip in job destruction in Figure 13c, there is no visual evidence of a lead-lag relationship between job flows within sectors.

To quantify this finding more rigorously, we estimated small-scale VAR models with job creation and destruction similar to those found in the literature.[20] These models impose structural restrictions designed to identify fundamental driving forces called allocative innovations and aggregate innovations. The relative importance of allocative versus aggregate driving forces is inferred from variance decompositions. A very disappointing common feature of these models is that estimates of this relative importance are highly sensitive to the identifying assumptions, and thus the models provide inconclusive estimates of the relative importance of allocative and aggregate driving forces.

Although our econometric models are similar, our focus is slightly different. Following Davis and Haltiwanger (1997), we model sectoral job flow rates together with an observable measure of aggregate activity—in

[19] Recall that NBER recession dating pertains to the entire U.S. economy, whereas the job flow data pertain only to manufacturing. Thus, if manufacturing activity tends to decline before nonmanufacturing in recessions, then job flows would spuriously lead the business cycle.

[20] This literature includes Blanchard and Diamond (1989, 1990), Davis and Haltiwanger (1996, 1997), Haltiwanger and Schuh (1998), and Campbell and Kuttner (1996).

Figure 11

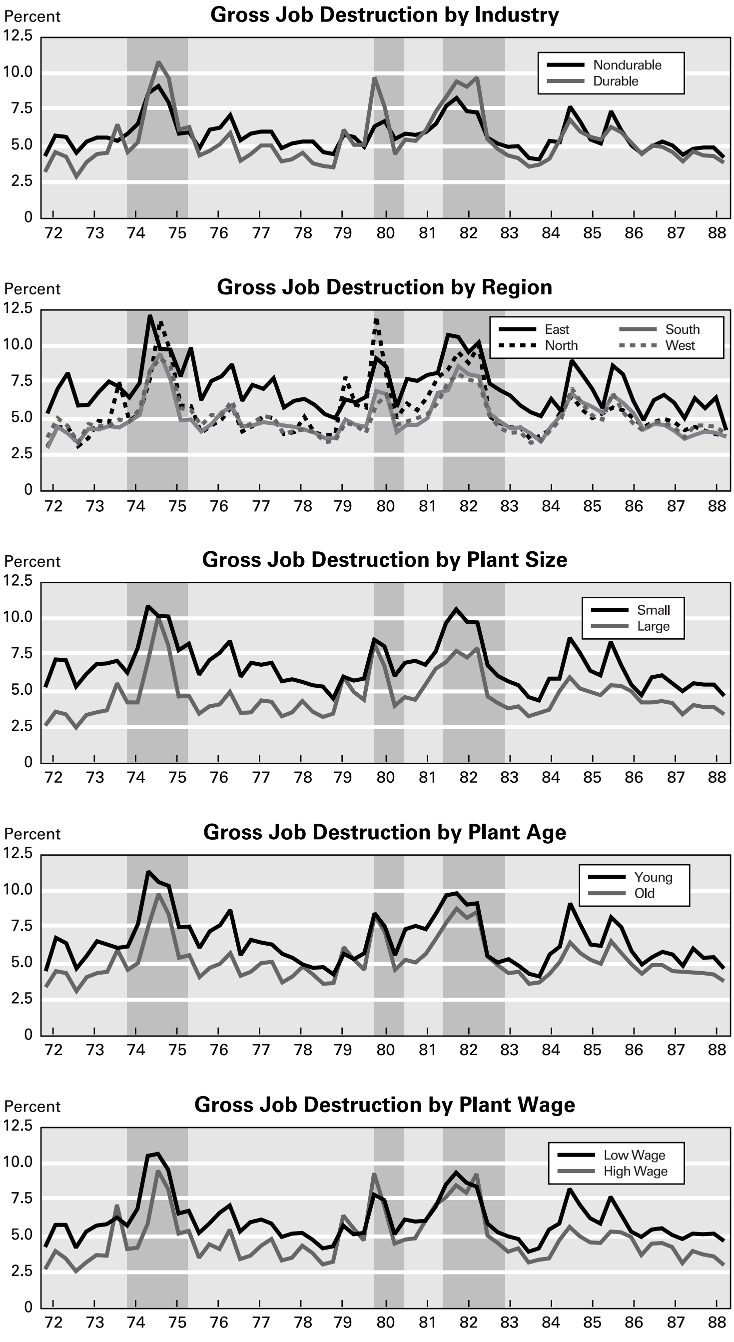

Source: Census / DHS job flows data.

Figure 11, continued

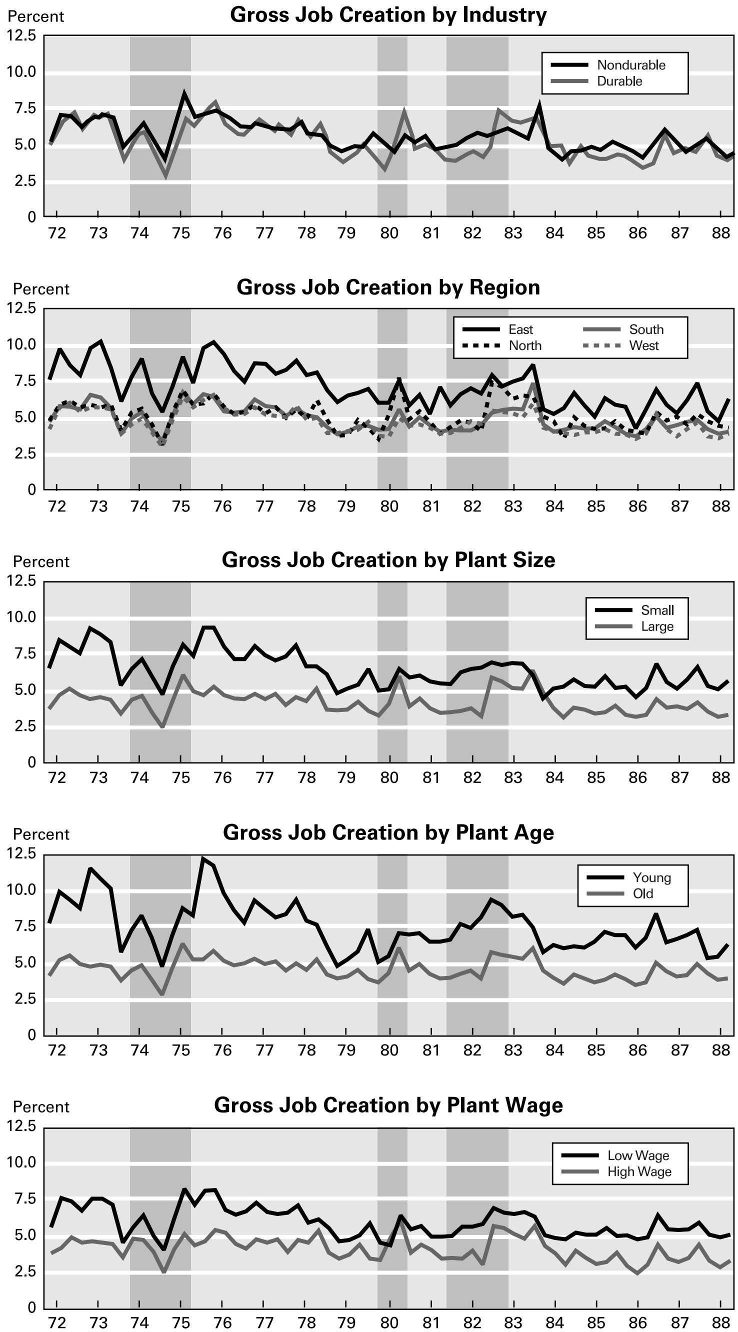

Source: Census / DHS job flows data.

Figure 12

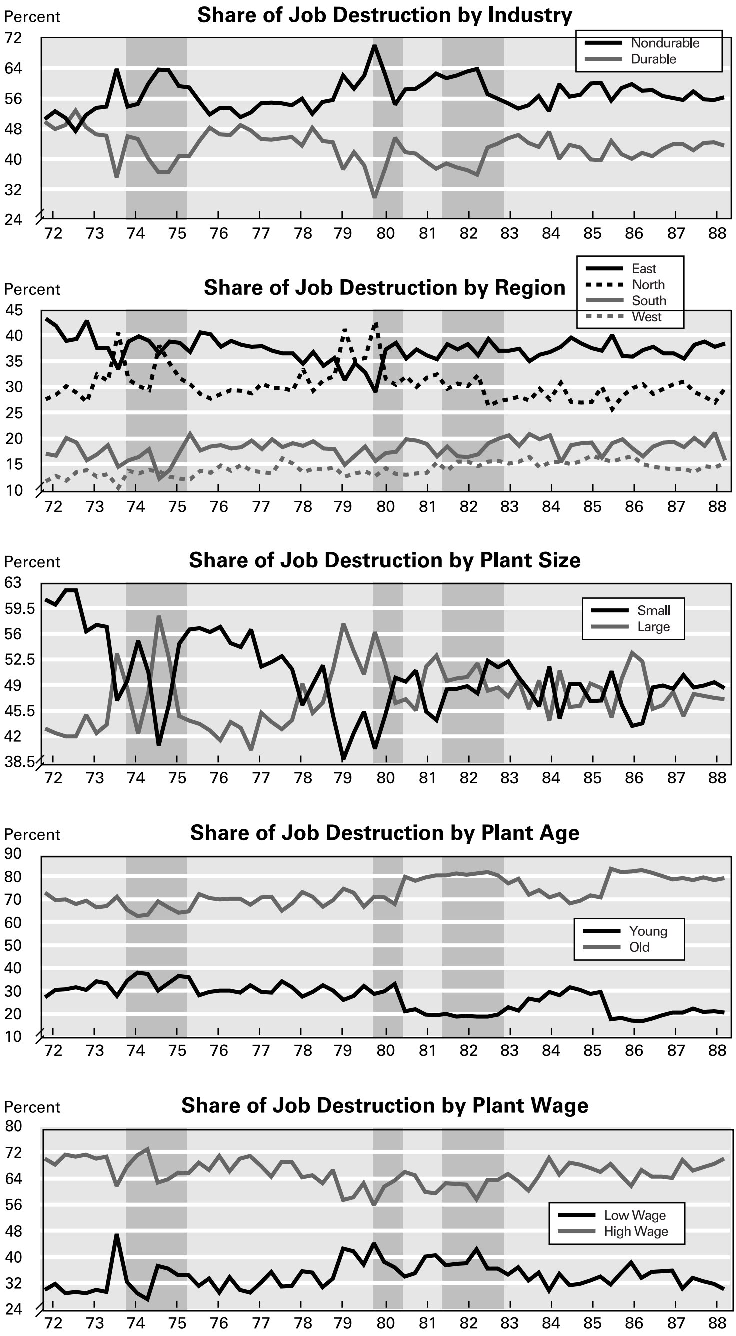

Source: Census / DHS job flows data.

Figure 12, continued

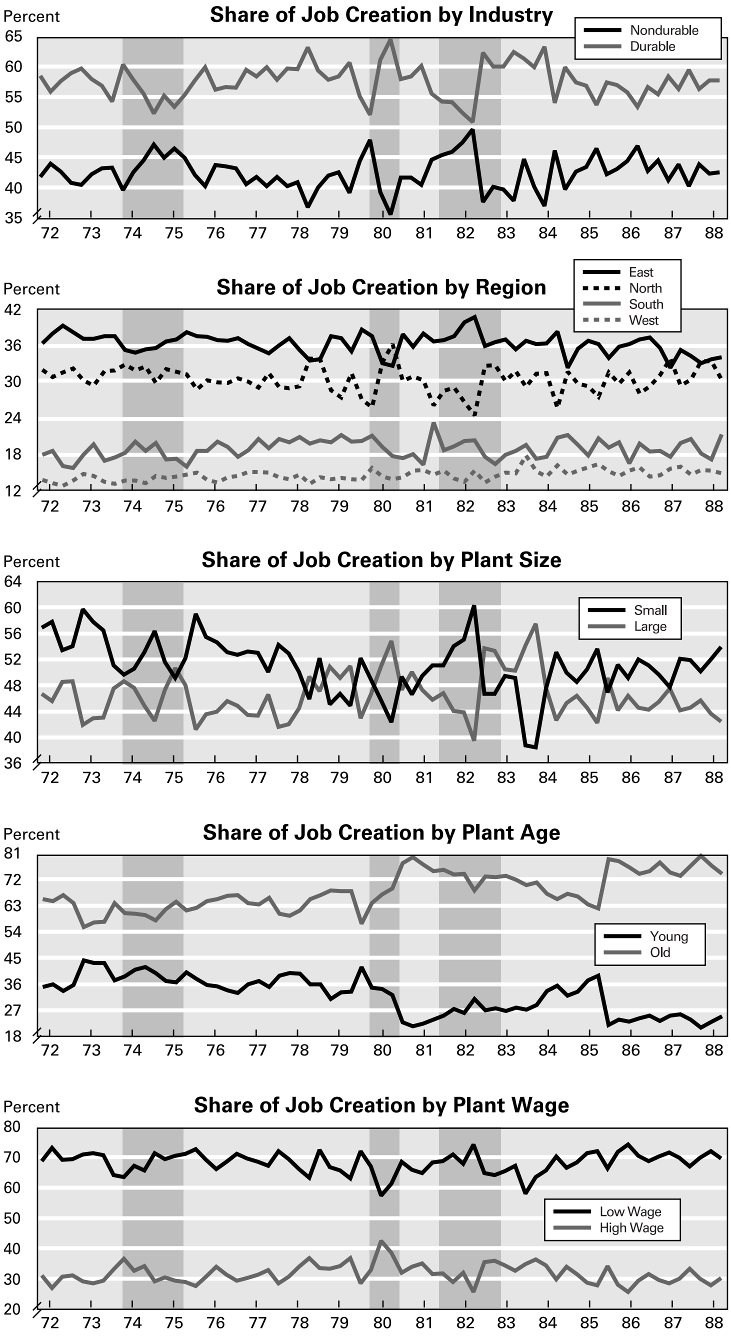

Source: Census / DHS job flows data.

Figure 13a
Sectoral Job Flows around 74:I Business Cycle Peak

Source: Census / DHS job flows data.

Figure 13a, continued

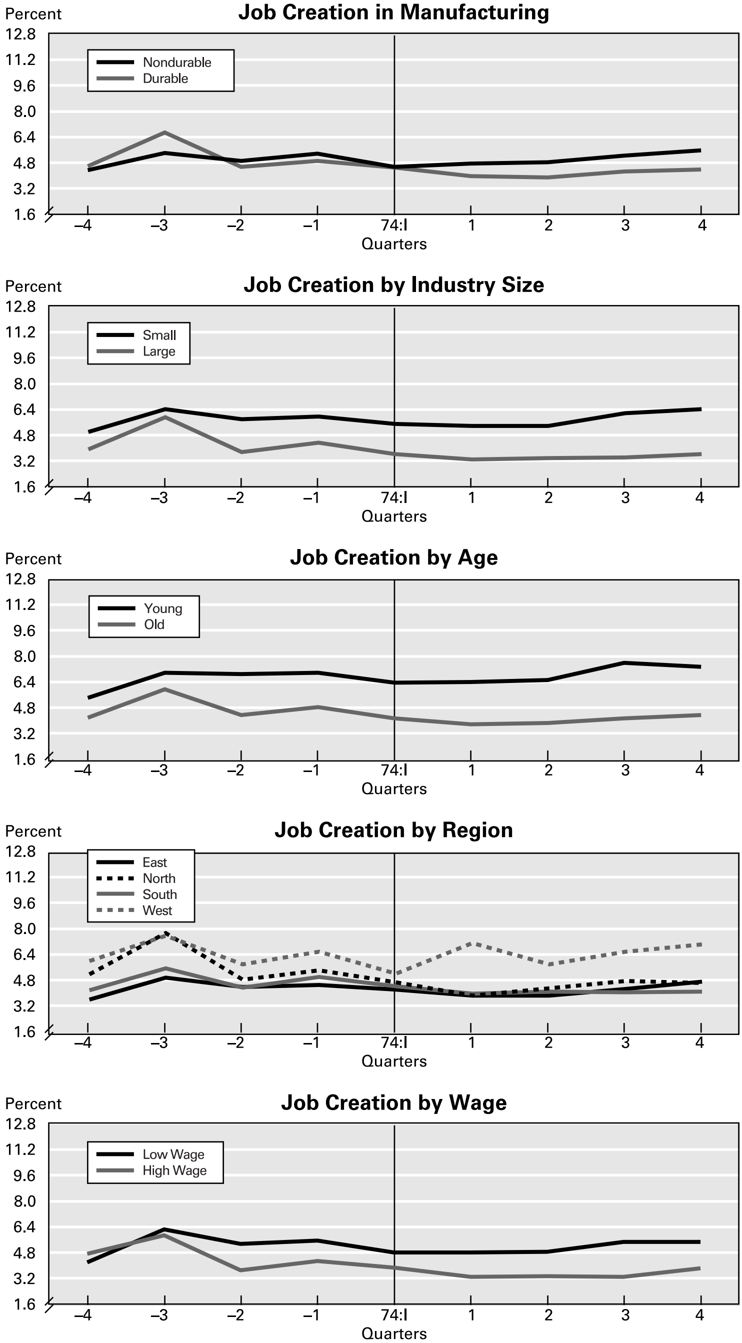

Source: Census / DHS job flows data.

Figure 13b
Sectoral Job Flows around 80:I Business Cycle Peak

Source: Census / DHS job flows data.

Figure 13b, continued

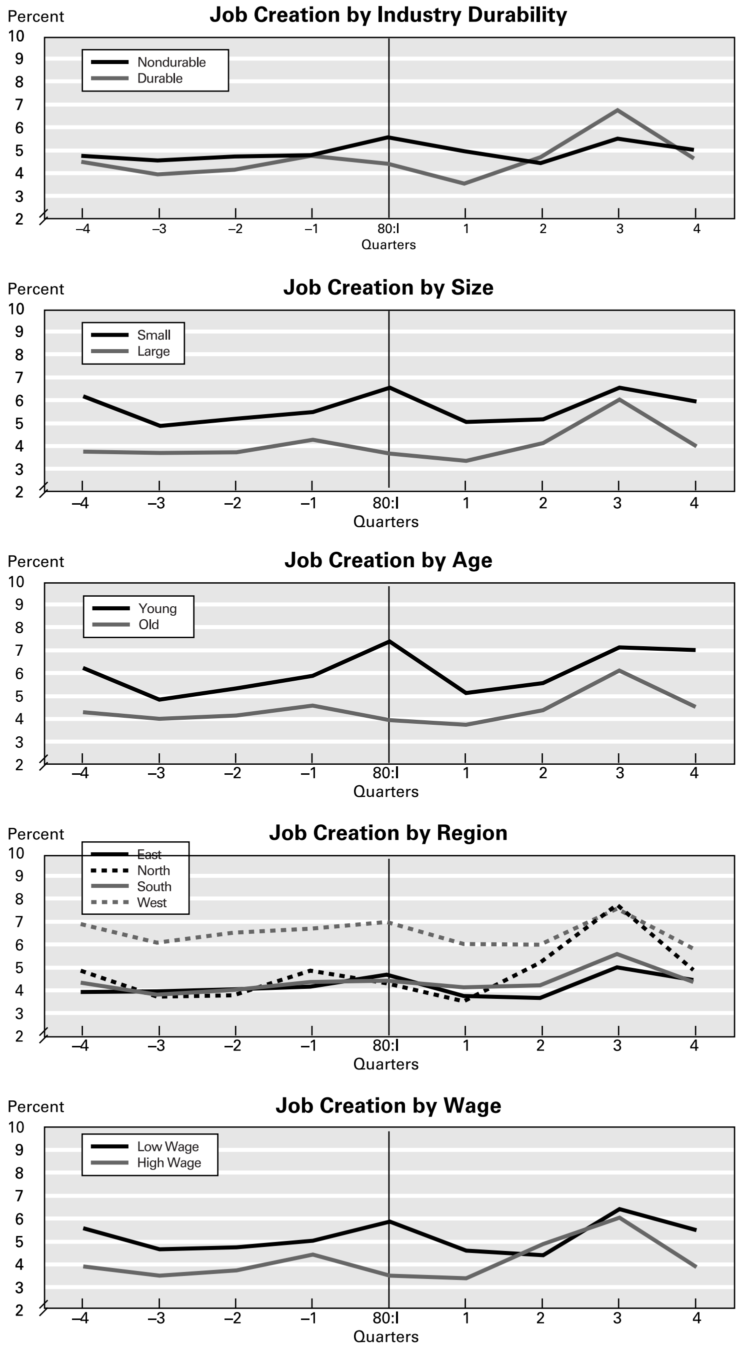

Source: Census / DHS job flows data.

Figure 13c
Sectoral Job Flows around 81:III Business Cycle Peak

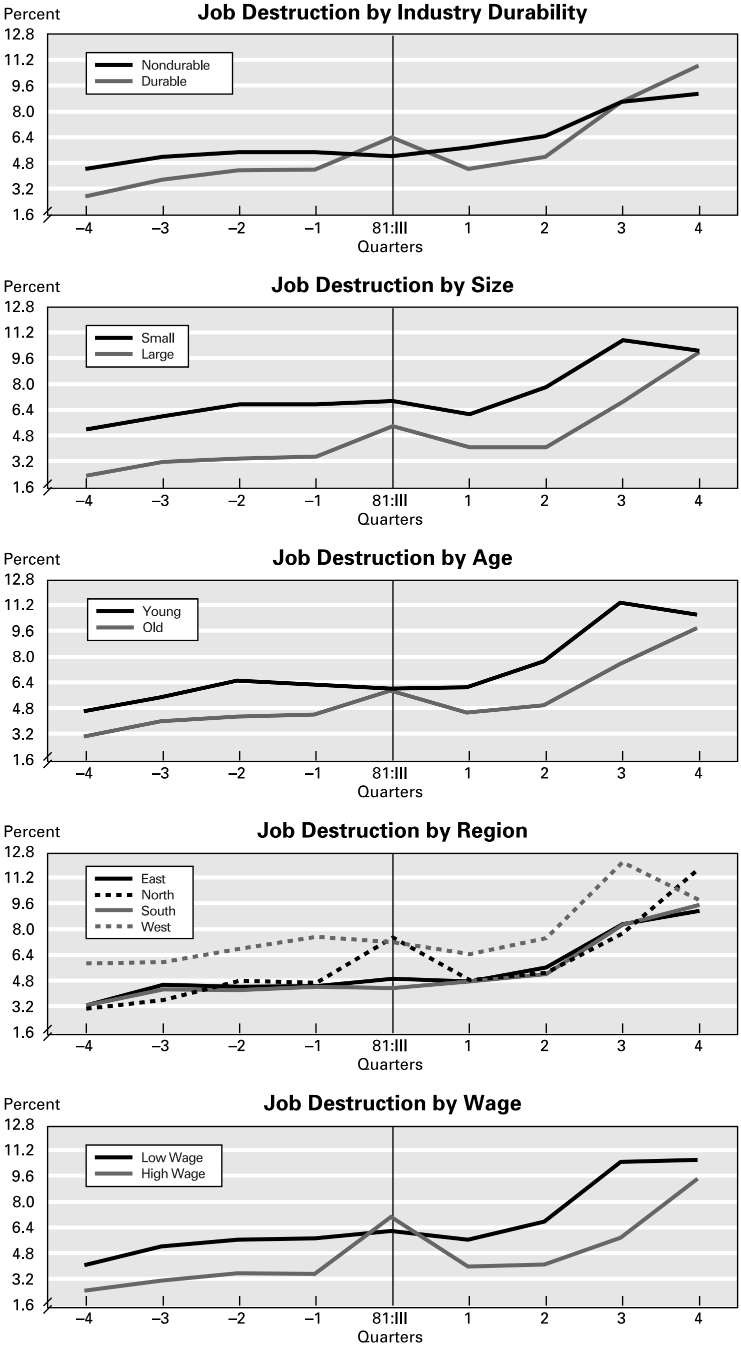

Source: Census / DHS job flows data.

Figure 13c, continued

Job Creation by Industry Durability

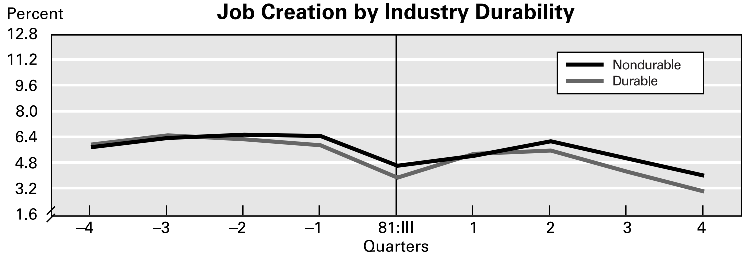

Job Creation by Size

Job Creation by Age

Job Creation by Region

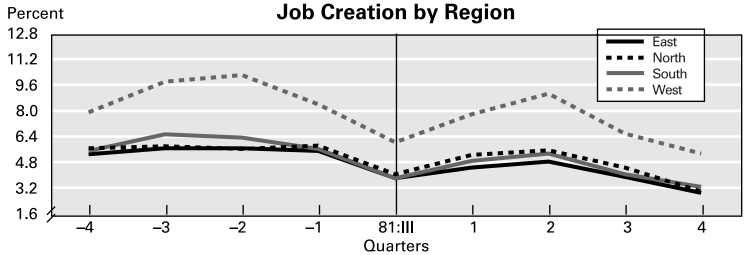

Job Creation by Wage

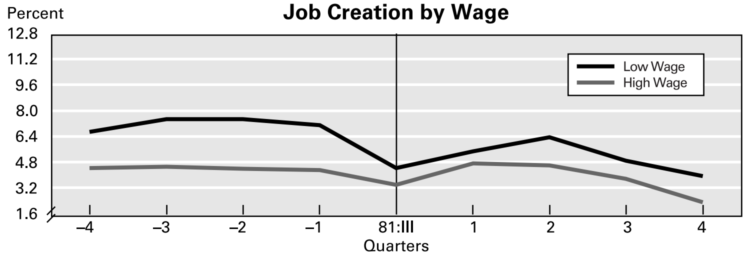

Source: Census / DHS job flows data.

Table 6
Causality Tests for Industrial Production (IP), Job Creation (C), and Job Destruction (D)[a]

Plant Classification	IP Equation			D Equation			C Equation		
	IP	C	D	IP	C	D	IP	C	D
Total Manufacturing	.32	.81	.05	.18	.57	.14	.95	.00	.08
Nondurable goods	.06	.57	.03	.04	.06	.00	.84	.00	.22
Durable goods	.37	.73	.06	.28	.94	.32	.88	.00	.03
East	.07	.51	.01	.04	.37	.00	.75	.00	.02
North	.18	.97	.52	.24	.99	.68	.45	.25	.01
South	.35	.74	.01	.45	.77	.00	.73	.00	.31
West	.07	.73	.01	.02	.00	.04	.34	.00	.06
Small	.12	.74	.02	.07	.04	.02	.91	.00	.30
Large	.32	.98	.11	.12	.81	.47	.84	.01	.03
Young	.15	.47	.01	.29	.11	.00	.75	.01	.02
Old	.37	.98	.13	.22	.88	.45	.64	.01	.04
Low-Wage	.27	.64	.01	.13	.14	.00	.72	.00	.17
High-Wage	.03	.92	.43	.04	.78	.84	.73	.00	.07
Specialized	.53	.21	.02	.37	.32	.05	.98	.00	.12
Diversified	.16	.63	.08	.08	.65	.23	.73	.00	.04

[a] Table entries are p-values from the hypothesis test that all lags of the variable (second row of headings) can be excluded from the equation (first row of headings).

Source: Authors' calculations using quarterly census/DHS job flows data for the period 1972 to 1988.

our case, growth of total manufacturing industrial production (IP).[21] But instead of trying to gauge the relative importance of contemporaneous allocative versus aggregate innovations, we simply examine what the VAR model lag structures imply about Granger-Sims causality. If IP growth causes job destruction, evidence accrues in favor of aggregate driving forces; if job destruction causes IP growth, evidence accrues in favor of allocative driving forces.

Table 6 reports the p-values from the causality tests. The third and fourth columns of numbers provide suggestive evidence that job destruction tends to cause IP growth. In all but four cases job destruction cannot

[21] We use manufacturing IP, rather than an economywide measure of aggregate activity such as GDP or unemployment, to prevent cyclical timing differences between manufacturing and nonmanufacturing from affecting the results. We use IP, rather than manufacturing net employment growth, to obtain a broader measure of economic activity. Two sensible extensions to the VARs are as follows: (1) include other aggregate variables, such as the federal funds rate and oil price growth, to account for the possibility that other aggregate variables cause job destruction and IP growth; (2) simultaneously include job flows from both sectors to examine intersectoral causality.

be excluded from the IP equation at the 10 percent level, and two of those cases are close (0.11 and 0.13). At the same time, IP lags are insignificant in most of the IP equations. In contrast, IP can be excluded from most sectoral job destruction equations, but job destruction generally cannot. In five sectors—durable goods, the south, young plants, low-wage plants, and specialized plants—the tests suggest destruction causes IP growth. In one sector—high-wage plants—the tests suggest that IP growth causes destruction. The remaining cases are statistically unclear.

Two other interesting conclusions emerge from these tests. First, lags of job creation are essentially irrelevant for either destruction or IP growth, though they are highly significant for creation itself. Only for nondurable goods, the west, and small plants does creation have any predictive power and that is only for destruction. Second, job destruction plays an important role in predicting job creation, while IP growth has no explanatory power for creation. Much more often than not, job destruction appears to cause job creation but not vice versa. To reiterate, aggregate activity tends to be irrelevant for predicting sectoral job flows, especially job creation.

Summarizing these causality tests without pushing them too hard, they offer modest evidence of allocative driving forces through lead-lag relationships. The data imply the following dynamic pattern. Job destruction in certain sectors reduces aggregate production. Both the decline in aggregate production and the allocative driving forces lead to job destruction in other sectors. Eventually, rising job destruction leads to higher subsequent job creation in most sectors.

This general pattern is difficult to reconcile with the view that an overall decline in aggregate activity causes an increase in job destruction. Also, the fact that job destruction in young, low-wage, specialized plants, rather than old, high-wage, diversified plants, leads IP growth is difficult to reconcile with vintage capital explanations of embodied technical change. Of course, these results could change significantly if an omitted aggregate driving force, such as monetary policy, leads sectoral job destruction. The results for young plants are somewhat suggestive of this, though one expects to see this causal pattern for small plants as a result of financial market imperfections, as argued by Gertler and Gilchrist (1994).[22]

Our investigation of time series causality between reallocation and business cycles yields two somewhat opposite conclusions. First, no sector clearly stands out as driving business cycles or recessions. Instead,

[22] The small-plant results may not be well-suited to testing for financial market imperfection effects because the small plant category includes small plants belonging to large companies, and the imperfections hypothesis is predicated on company-level financial conditions. Regressions on sectors defined by plant ownership type (single-plant companies versus multi-plant companies) and size would be more appropriate.

the cyclical experience is broadly similar for all types of plants, with differences in reallocation activity being primarily quantitative rather than qualitative in nature. These conclusions suggest that aggregate driving forces are at work. Second, the VAR models provide reasonably strong evidence that aggregate activity does not help predict sectoral job flow behavior, and modest evidence that sectoral job destruction helps predict aggregate activity. These causality results suggest that allocative driving forces are at work, although the VAR modeling and causality testing clearly require further development.

Reallocation, Investment, and Productivity

In this section, we examine the relationships among job reallocation, investment, and productivity growth. Two main factors motivate the investigation. First, as noted earlier, it is important to jointly consider decisions about all factors of production. Capital is an obvious factor to begin with given its relatively large share of output, plus the fact that it can affect job flows in different ways. Second, many reallocation theories based on cleansing, reorganizing, or other forms of "creative destruction" activities suggest that old, inefficient, and unproductive capital is destroyed and new, efficient, and productive capital is created. This idea implies that investment and productivity changes may be connected integrally to gross job flows, as workers matched with unproductive capital lose their jobs and productive new capital-worker matches are created. It also raises the general question of whether transitory declines in aggregate demand (recessions) are associated with permanent effects on productivity, presumably raising it. Productivity-enhancing creative destruction and embodied technological change could be internal to the plant, through restructuring, or internal to the industry, through shutdown, start-up, and other permanent reallocation across plants.

Direct evidence on the link between reallocation and productivity is limited, and most previous work has focused on the link at the plant level.[23] Baily, Bartelsman, and Haltiwanger (1996), for example, find that permanent job destruction in plants disproportionately accounts for the procyclicality of productivity. Other studies decompose industry productivity growth into the contributions of within-plant versus between-plant changes in productivity, including entry and exit. Several studies find that entry and exit of plants is a primary determinant of aggregate productivity growth, especially in the longer run. Foster, Haltiwanger,

[23] The relevant literature includes Dunne, Roberts, and Samuelson (1989); Baily, Hulten, and Campbell (1992); Dunne, Haltiwanger, and Troske (1996); Baily, Bartelsman, and Haltiwanger (1996); Olley and Pakes (1996); Liu and Tybout (1996); and Foster, Haltiwanger, and Krizan (1998).

and Krizan (1998) find that increases in nonproduction labor share are almost entirely driven by within-plant increases in the longer run. A serious drawback to these plant-level productivity studies, however, is the use of 4-digit industry deflators for output and input prices when there is likely to be tremendous heterogeneity in prices across plants.

In this section, we take a notably different approach to quantifying the relationship between job reallocation and productivity. First, we quantify the relationship between job reallocation and *trend*, rather than cyclical, productivity growth. This method is found in DHS (1996), which documents a jump in trend productivity growth in the steel industry after the massive job destruction during the 1981–82 recession (Figure 5.8, p. 117). One reason to consider trend growth is to mitigate the influence of capacity utilization. Another is that new investment may take several years to reach efficient operation and manifest itself in productivity. Second, we examine the relationship at the 4-digit industry level, rather than the plant level. The industry-level approach captures the cumulative effects of within-plant and between-plant productivity-enhancing reallocation. In addition, the price deflators required to construct quantities such as output and investment are industry-level prices. Thus, unlike the plant-level analyses, the industry analysis is immune to potential biases from heterogeneous plant-level prices.

The empirical exercise is as follows. For the three major recessionary periods in the sample, we sorted industries by the cumulative amount of industry reallocation during the recession. Then, for each quintile of industries, we calculated the changes in trend productivity growth (labor and total factor) and investment growth around recessions. Trend growth is measured as the average growth rate during an expansion, and is assumed to be able to change between expansions.[24] The key issue is to determine the extent to which cumulative reallocation in recessions is associated with increases in trend productivity and investment growth. If industries with the most cumulative reallocation during a recession also experience increased productivity and investment growth, in either absolute or relative terms, then the industries may be experiencing

[24] This exercise uses annual data from the Census/DHS job flows and the NBER Productivity Database described in Bartelsman and Gray (1996). Expansions are 1971–73, 1976–79, 1983–89, and 1992–94. The 1980 and 1981–82 recessions were combined because of their close proximity and relation. Cumulative industry reallocation is the sum of reallocation in all recession years, where reallocation is measured as the deviation from average industry reallocation during expansion years to control for cross-sectional variation in the level of reallocation (essentially a fixed effect). Labor productivity is real value added per production worker hour; total factor productivity (TFP) is preconstructed; and investment is total investment expenditures deflated by the investment deflator. Using average growth rates during expansions should uncover secular (trend), rather than cyclical, changes in productivity and investment. TFP growth for industry quintiles is the value-added weighted average of industry TFP growth.

Table 7
Productivity and Investment Growth around Recessions, by Cumulative Job Reallocation

Cumulative Job Reallocation (Percent)[a]	Productivity Growth (Percent)[b]						Investment Growth (Percent)[b]			
	Labor			Total Factor						
	Before	After	Chg.	Before	After	Chg.	Before	During	After	Change
1974–75 Recession										
Very low (<−3.4)	4.2	3.4	−.8	−.4	1.9	2.3	1.2	6.5	7.1	5.9
Low (−3.4 to 1.0)	4.4	1.3	−3.1	.2	.3	.1	12.7	−.9	13.1	.4
Average (1.0 to 5.9)	5.0	1.7	−3.3	.5	.7	.2	12.0	−8.0	6.7	−5.3
High (5.9 to 12.0)	1.1	.7	−.4	.8	1.3	.5	22.4	−19.0	10.4	−12.0
Very high (>12.0)	6.0	2.2	−3.8	1.7	2.8	1.1	22.7	−28.4	14.1	−8.6
1980–82 Recession										
Very low (<−3.7)	4.5	6.8	2.3	3.1	2.3	−.8	14.9	−2.7	6.1	−8.0
Low (−3.7 to 2.6)	1.9	3.5	1.6	1.0	.4	−.6	10.1	−2.0	6.5	−3.6
Average (2.6 to 9.1)	1.7	2.6	.9	1.5	.5	−1.0	11.2	−8.7	6.6	−4.6
High (9.1 to 18.2)	1.4	4.6	3.2	.9	1.4	.5	8.9	−15.9	5.9	−3.0
Very high (>18.2)	.6	3.3	2.7	.3	1.3	1.0	6.4	−22.8	3.3	−3.1
1990–91 Recession										
Very low (<−9.2)	4.2	9.9	5.7	1.0	5.5	4.5	2.8	−.1	8.2	5.4
Low (−9.2 to −5.2)	4.4	3.8	−.4	1.0	1.2	.2	4.3	−1.4	3.5	−.8
Average (−5.2 to −1.3)	3.1	1.9	−1.2	.9	.9	.0	8.9	−3.2	5.6	−3.3
High (−1.3 to 2.1)	4.5	5.7	1.2	1.8	2.8	1.0	6.0	−5.5	2.3	−3.7
Very high (>2.1)	4.2	2.9	−1.3	.7	.9	.2	5.4	−1.7	2.7	−2.7

[a] Cumulative reallocation is the sum of the rates of job reallocation in a 4-digit industry during the recession, where reallocation is measured at the deviation from mean industry reallocation during all expansion years.
[b] Annual average growth rates during the expansions before and after recessions.
Source: Authors' calculations using Census/DHS job flows data and the NBER Productivity Database for the period 1973 to 1994.

cleansing, reorganizing, or other forms of creative destruction. And, if so, it would then be appropriate to ask whether or not such creative destruction was the driving force behind the recession.

Table 7 reports the results. Total factor productivity (TFP) growth is the more pertinent measure for this exercise and yields the clearest results. In two of the three recessions (1974–75 and 1990–91), by far the greatest increase in trend productivity growth (column 6) occurred in industries with the *least* cumulative reallocation. In fact, productivity growth increased very little except in very low reallocation industries. In contrast, during the 1980–82 recessionary period, industries with the *most* cumulative reallocation experienced the largest increases in productivity growth. However, trend productivity growth did not increase much in any industry group after this recession, and the increase in very high reallocation industries was considerably smaller than the increase in very low reallocation industries around other recessions.

What accounts for these changes in trend productivity growth? The data suggest that investment growth likely played a key role. In all three recessions, industries with the largest increase in trend productivity growth also exhibited the largest increase in real investment growth. In the 1974–75 and 1990–91 recessions, industries with very low reallocation increased their investment growth markedly; the remaining industries actually reduced their investment growth collectively. In the 1980–82 recession, industries with very high reallocation showed the most improvement in investment (although investment in all industries declined, investment in very high reallocation industries declined the least). Not surprisingly, industries with the most reallocation exhibit the largest cyclical reductions in net employment growth and they also exhibit the largest cyclical reductions in investment growth. But these industries also exhibit the largest secular reductions in investment growth.

Figure 14 illustrates the basic point by plotting job creation and destruction in the very low and very high reallocation industries for each recession. The cleansing theories of Caballero and Hammour (1994, 1996), for example, postulate a sharp increase in job destruction followed by a lagged surge in job creation. Industries with very high reallocation exhibit precisely this postulated pattern in the first two recessions. However, these industries *reduced* their rates of investment—by 23 to 28 percent during the recessions and by 3 to 8 percent in trend—and did not generally see absolute or relative trend productivity growth rise much, if at all. In the third recession, this job flow pattern generally did not occur, at least not by 1993. In contrast, industries with very low reallocation, whose trend productivity growth often increased markedly, exhibit essentially acyclical job reallocation. The variances of job creation and destruction are about the same in each subsample (variance ratios of 1.0, 0.6, and 1.6, chronologically).

Table 7 also shows two other interesting developments related to the cross-sectional distributions of trend productivity and investment growth around recessions. First, the *levels* of cross-sectional trend productivity and investment growth were inversely related to the *changes* in trend growth. In other words, industries with relatively low trend growth before recessions tended to experience the largest increases in trend growth after recessions, and vice versa, *regardless of the amount of cumulative reallocation*. Second, dispersion in cross-sectional trend investment growth tended to decline dramatically after recessions. In other words, trend investment growth rates were quite different across industries before the recession but quite similar afterward. However, this feature was not observed generally in trend productivity growth.

Implications and Extensions. Industries experiencing unusually large bursts of job reallocation during recessions generally do *not* exhibit significant increases in trend productivity or investment growth. Instead, increases in productivity and investment tend to occur in industries with

Figure 14
Job Flows in High and Low Reallocation Industries in Three Time Periods

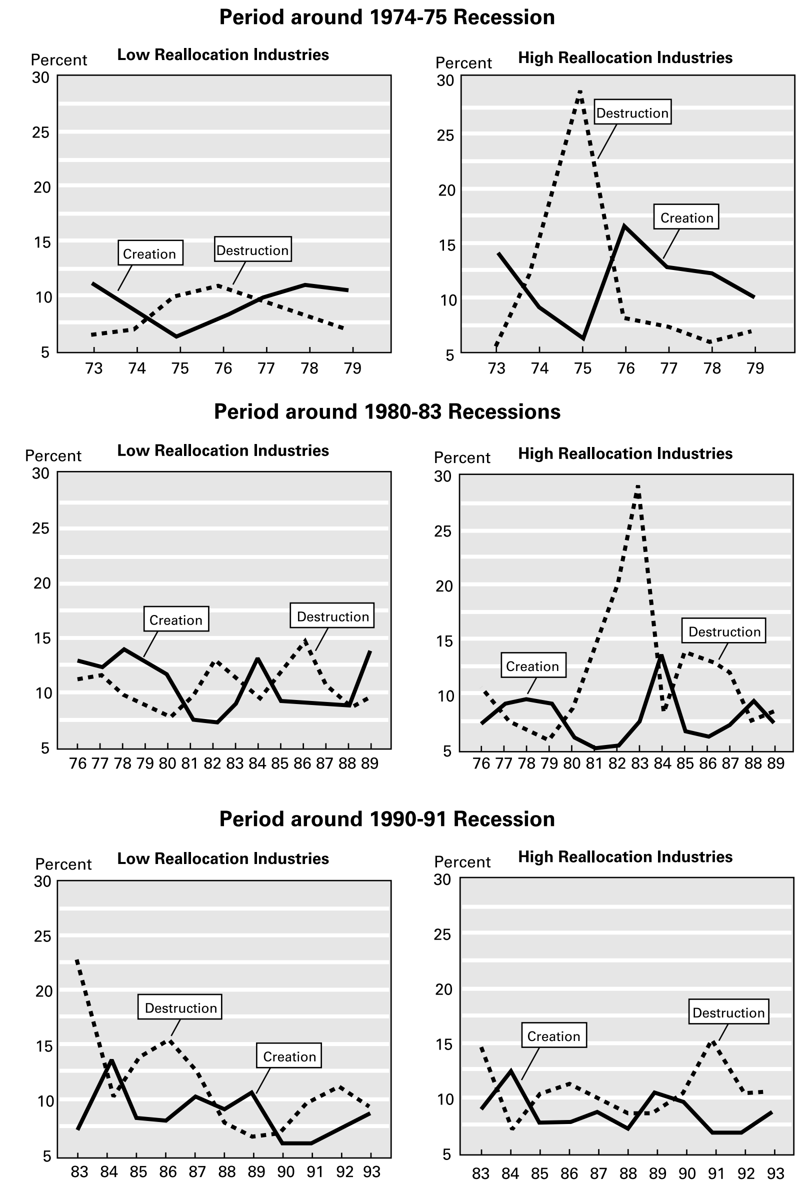

Source: Census/DHS job flows data.

very little job turnover and relatively low trend productivity and investment growth prior to the recession. The 1980–82 recessionary period shows some evidence for the conventional wisdom that this recession involved fundamental, productivity-enhancing restructuring. But the evidence is quite weak, and based on what we know about earlier postwar recessions, it appears that the 1980–82 recession was rather unusual in this regard.

This evidence on the connection between reallocation and trend productivity and investment growth does not conform well to theories that posit improvements in investment and productivity through creative destruction channels in either the short or long run. Whether the channel is through embodied technical change, intensified search and rematching, retooling, or some other mechanism, cleansing, reorganizing, and related theories imply some observable improvements in long-run productivity. Empirically, however, increased reallocation normally is not correlated with increased trend productivity and investment around recessions. Cyclical productivity fluctuations may be larger in the short run, but they are transitory and associated with fluctuations in utilization. Changes in the trend dominate over the long run.

Even if productivity improvements implied by creative destruction theories are very small, these theories have trouble explaining two main business cycle patterns observed in this data. First, trend productivity and investment growth often *fall* after bursts of job reallocation rather than rise, as predicted by the theories. It is possible, however, that the kinds of market inefficiencies stressed by Caballero and Hammour (1996, 1998), such as appropriability problems and financing constraints, could more than offset any productivity gains. But the second pattern is more problematic: Trend productivity and investment growth often rise substantially *without* any significant increase in reallocation. Thus, although these theories may explain how the process of reallocation works in *some* industries, they do not explain the connection among reallocation, investment, and productivity for the entire economy nor why the connection is linked to the business cycle. Because high reallocation industries normally do not exhibit observable productivity gains while low reallocation industries do, creative destruction theories do not appear to be complete explanations of business cycle fluctuations.

Our results suggest that understanding the role of investment—an old, familiar feature of recessions—is critical to understanding the nature and consequences of job reallocation. Investment appears to be a key determinant of productivity growth, but the bulk of investment growth does not always coincide with massive job reallocation over the business cycle or in the long run. Investment often is redirected toward industries that do not require major structural reallocation and that weather the storm of recession relatively well. For some reason, these industries also tend to have relatively low trend productivity and investment growth before the recession. These new facts are intriguing but puzzling, and certainly merit further consideration.

Sectoral Price and Productivity Dispersion

Early empirical efforts to identify allocative driving forces behind business cycles, such as Lilien (1982), focused on the sectoral dispersion in employment growth rates. Lilien found that dispersion in employment growth rates across sectors (defined as 1-digit industries) was countercyclical and helpful in explaining the time series behavior of the unemployment rate. Abraham and Katz (1986) disputed this finding by showing that heterogeneity in the cyclical response of industries to aggregate driving forces would produce sectoral growth rate dispersion as a consequence, rather than cause, of business cycles.

Loungani, Rush, and Tave (1990) and Loungani and Trehan (1997) contributed further evidence that allocative forces, termed "sectoral shifts," are important for business cycle fluctuations, using dispersion in stock prices. This measure is less susceptible to the Abraham-Katz criticism because, unlike employment dispersion, it appears to be econometrically exogenous: Stock price dispersion forecasts unemployment and output but not vice versa. Nevertheless, many observers remain unconvinced that allocative forces induce reallocation across sectors and cause business cycle fluctuations.

Even stock price dispersion may result from business cycles rather than cause them, though. If agents are rational, stock prices should equal the expected present value of future dividends. Thus, if firms' dividends fluctuate cyclically because of some link to aggregate activity, and firms' cyclical sensitivity is heterogeneous, then expected declines in aggregate activity will cause stock price dispersion as well. Empirically, stock price dispersion will lead employment dispersion because expectations and stock prices adjust quickly while real quantities, which may be subject to adjustment costs, change slowly. Thus the question remains: What causes dividends to fluctuate—the business cycle or firm- and sector-specific factors?

This section presents new evidence on allocative driving forces by examining two factors more likely to be a root *cause*, rather than *result*, of reallocation: relative prices and relative productivity. In a multisector general equilibrium model with fixed consumer preferences, relative prices and productivity are key determinants of firms' demand for factors of production and of consumers' demand for output.[25] Thus, changes in

---

[25] Determination of sectoral demand and production is complicated by numerous demand-side characteristics, such as differentiated products, imperfectly competitive industries, multi-product plants, corporate structure, and related issues. Size differentials among

relative prices and productivity should be key forces driving the reallocation of final demand and factors of production across sectors. The actual timing of the reallocation process (creation and destruction) will depend on the extent to which agents view the changes as permanent versus transitory, and on the relative flexibility of prices and productivity versus quantities. With uncertainty and real frictions, it seems reasonable to expect that the driving forces—changes in relative prices and productivity—precede the actual reallocation of quantities.

At least three types of relative prices determine the allocation of factors of production across sectors.[26] Output (finished goods) prices are a primary determinant of the relative level of demand for a sector's product, so relative output price changes will alter demand for output and factors of production across sectors. Input (raw) materials prices and investment prices are primary determinants of the mix of materials, capital, and other factors of production, so changes in relative materials and investment prices will change the mix of factors of production and demand for labor across sectors. Relative productivity, a key determinant of relative output prices, can also affect the demand for factors of production through many channels: relative output prices, investment, and so on. Cross-sectional changes in these relative prices and productivity should lead to reallocation of factors of production such as labor across sectors.

Figure 15 charts job destruction and dispersion in relative price changes and in total factor productivity (TFP) growth.[27] If sectoral relative prices and TFP growth converge to steady-state equilibrium cross-section distributions with a constant finite variance, dispersion would appear as a straight horizontal line in the figure. Actual dispersion reflects the extent to which sectoral rates deviate from that equilibrium. When dispersion increases, relative prices and productivity growth are changing across sectors. If the increase in dispersion is large enough and

plants or firms in industries raise questions about competitiveness, credit-market access, and scope for new technology. In short, the allocative role of relative prices and productivity is clearly more complex in practice.

[26] The term "sector" is used broadly enough to mean plants or group of plants with common characteristics, such as industry, geography, or any other observable plant characteristic.

[27] This exercise is also based on 4-digit industries and uses the Census/DHS job flows and NBER Productivity Database. Industry prices are the shipments deflators, materials price deflators, and investment price deflators from the NBER Productivity Database. Aggregate prices are the PPI for finished goods (shipments), PPI for crude materials (materials), and the GDP fixed-weight price index for nonresidential business fixed investment. Because the prices are indexes (all equal 100 in the base year), we cannot use dispersion in relative price *levels*; instead, our dispersion measures are the standard deviations of log *changes* of the relative prices. This approach is consistent with the relative price dispersion literature, typified by Debelle and Lamont (1997). TFP growth dispersion is the standard deviation of the gap between industry and manufacturing growth.

Figure 15
Job Destruction and Dispersion in Prices and Productivity across Industries

Real Shipments Price Changes

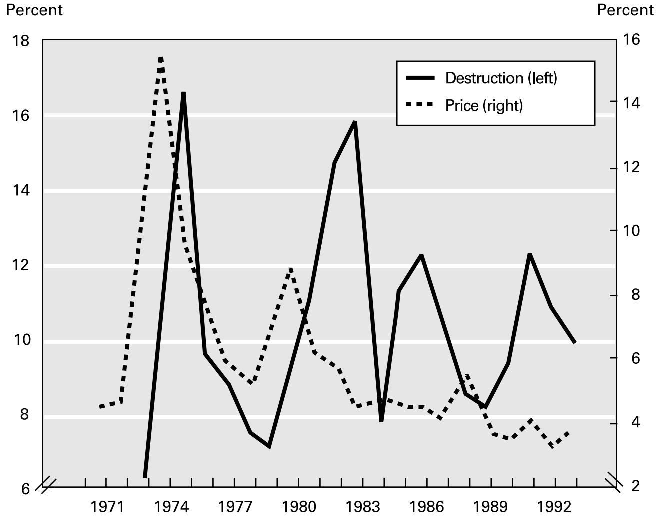

Real Materials Price Changes

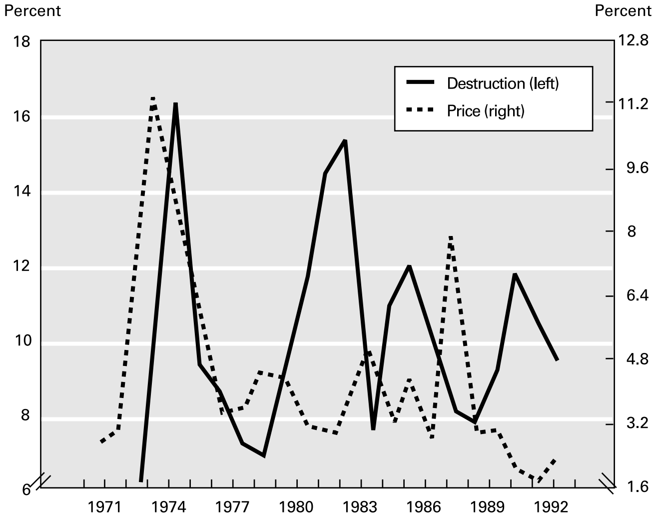

Source: Census/DHS job flows data.

Figure 15, continued

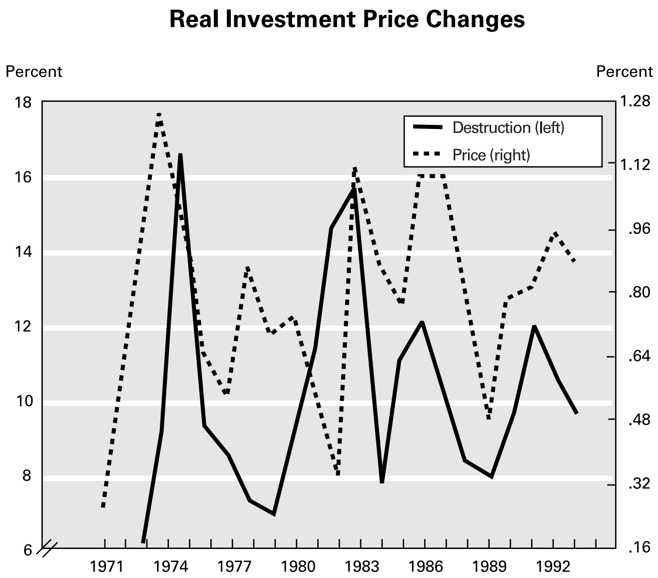

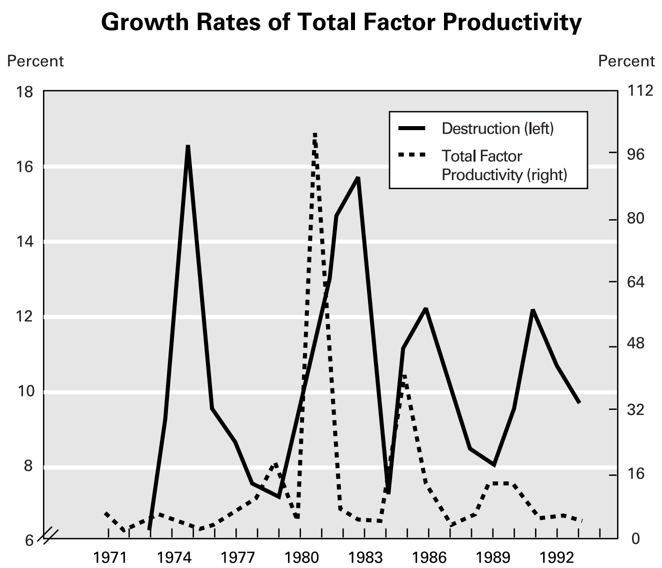

Source: Census/DHS job flows data.

permanent enough, it can cause permanent job destruction in adversely affected sectors and permanent job creation in positively affected sectors. The timing of job creation and destruction in response to the increased dispersion will depend on several factors. If productivity declines in one sector but is unchanged in another—an *unfavorable* sectoral driving force—job destruction is likely to precede creation. But if productivity increases in one sector and not the other—a *favorable* sectoral driving force—job creation may precede destruction. However, this latter scenario will depend on the availability of workers—unemployed and new labor force entrants—to fill the jobs.

Not surprisingly, recent decades have seen sharp increases in the dispersion of relative price changes and productivity growth. Perhaps surprising, however, is the timing: Dispersion increased sharply *prior* to each major increase in job destruction, which in turn was often followed by an increase in job creation. Cross-correlations (not reported) indicate that every dispersion measure leads job destruction by one or more years. This general pattern of increased dispersion, increased job destruction, and then increased job creation seems broadly consistent with the view that the U.S. economy suffered a series of persistent, unfavorable allocative driving forces during this period.

Prior to the 1974–75 recession, all price dispersion increased dramatically, and productivity dispersion rose somewhat too. In each case, dispersion peaked one year ahead of job destruction. Prior to the 1980–82 recessionary period, all four dispersion measures rose again but this time much earlier, with dispersion peaking as early as 1978 (relative investment price changes) and job destruction peaking in 1983. The increase in price dispersion was much smaller than in the previous recession, but the increase in productivity dispersion was much larger. Prior to the 1986 peak in destruction (not an official period of economywide recession), dispersion increased in all but the output price change measure. Prior to the 1990–91 recession, dispersion increased in all measures except investment price changes, which was contemporaneously correlated with destruction. The increase was especially evident in materials prices, corroborating the view that inflationary pressures during this time began "in the pipeline" of materials and supplies distribution.

Implications and Extensions. The data indicate that relative price change and relative productivity growth become significantly more dispersed prior to large increases in job destruction and hence prior to recessions. The leading nature of the dispersion is significant. Unlike employment growth dispersion, which is correlated contemporaneously with the business cycle, dispersion in relative prices and dispersion in TFP growth lead the business cycle. Hence, it is much less likely that the business cycle causes the dispersion in relative price changes and TFP growth. Unlike stock price dispersion, which is based on expectations of economic activity and appears to lead the business cycle for relatively short horizons, dispersion in relative price changes and TFP growth can lead the business cycle by several years. Whether such a long lead is reasonable depends on how quickly agents learn about the dispersion, its permanence, and the costliness of responses to the dispersion, such as the shutdown and start-up of new plants.

Although a picture can be worth a thousand words, the evidence presented here is merely suggestive and obviously does not establish causality from reallocation to business cycles. It does establish, however, that the variables most expected to determine the allocation of factors across sectors do change significantly prior to significant increases in job destruction and reallocation, and prior to decreases in aggregate economic activity. Note, however, that even if dispersion in relative price change or productivity growth induces reallocation and reallocation causes a reduction in aggregate economic activity, it would be incorrect, strictly speaking, to conclude that reallocation causes business cycles. Instead, relative prices and productivity would be the causes, or driving forces, behind business cycles. Of course, this conclusion leaves unanswered the question of what causes dispersion in relative price change and productivity growth.

In any event, the results in this section clearly motivate further investigation. The data suggest that models purporting to explain gross job flows at the plant or sector level should take relative prices and productivity into account. Furthermore, the models must also explain the significant lag between changes in the incentives to reallocate and the actual reallocation. It seems reasonable to suspect that expectations, learning, investment irreversibility, adjustment costs, and other frictions will be important components of successful models.

Our plan is to continue investigating these issues at the detailed industry and plant levels. The main unexplained issue is whether the dispersion is actually causing the fluctuations in gross job flows associated with deliberate reallocation. The only way to resolve this issue is to examine whether relative prices and productivity are important explanatory variables for the job flow behavior of particular plants and industries. Specifically, we need to know whether job destruction is occurring in plants and industries where relative prices (productivity) are rising (falling), and whether job creation is occurring where relative prices (productivity) are falling (rising). In future research, we plan to investigate this by estimating dynamic labor demand models with panel data econometric techniques using industry-level data bases and the Longitudinal Research Database (LRD).

The Outlook for Job Reallocation and Business Cycles

Undoubtedly, the primary drawback to incorporating reallocation into macroeconomic models and government policymaking is the lack of broad and timely gross job flow data. The best U.S. data source is the LRD, but the LRD data are several years behind (currently five) at best and available only for manufacturing. Although the Census Bureau is making good progress on releasing data on a more timely basis and in acquiring nonmanufacturing data, timely release of economywide gross job flows is years away. The U.S. Bureau of Labor Statistics is also making good progress on producing up-to-date gross job flow data, but release of these data is not imminent either.

New Proxies for Gross Job Flow Data

Fortunately, two proxies for gross job flows are available on a timely basis. One is a measure of job flows between 4-digit industries, as reported in Ritter (1993, 1994) and Haltiwanger and Schuh (1995). Between-industry job flows exhibit cyclical characteristics strikingly similar to those exhibited by within-industry and total gross job flows. Between-industry job flows still require fairly large data base maintenance and manipulation, but they provide job flows estimates within a few months of the current period and they cover the entire nonfarm U.S. economy.

A second, and previously unexploited, proxy for gross job flows comes from the National Association of Purchasing Managers (NAPM). The NAPM publishes data reflecting the assessments of about 350 purchasing managers in manufacturing and some nonmanufacturing industries about the qualitative change (higher, lower, or the same) in economic activity at their companies. Diffusion indexes are used to summarize the net change in employment, as well as production, inventories, deliveries, prices, and other variables for all NAPM companies. Underlying the diffusion indexes are "gross flow" measures representing the fraction of companies with higher or lower activity. Using the NAPM employment data, job creation (destruction) is proxied by the percentage of companies with higher (lower) employment.

Two key differences arise between the NAPM job flows and the LRD job flows. First, the NAPM job flows reflect the number of companies with increasing or decreasing employment, rather than the number of jobs actually being created or destroyed. Thus, if jobs are being created or destroyed disproportionately by certain types of companies (as the LRD data indicate they are), the NAPM job flows may be substantially biased. Only if all firms were the same size and employment adjustments the same magnitude would the NAPM job flows exactly mirror the actual gross job flows. Second, the unit of observation is a plant in the LRD but a company in the NAPM. Thus the NAPM proxy is an interfirm job flow measure that will understate the interplant measure from the LRD, but nothing is known yet about the relative cyclicality of interplant and interfirm job flows.

Despite potential measurement drawbacks, the NAPM job flows are readily available each month and for a long period of history, so it is worth seeing whether the NAPM and LRD job flows are closely correlated. Figure 16 plots the NAPM and LRD gross job flow data. Despite some differences, the NAPM job flows do a pretty decent job of proxying for the LRD job flows. The quarterly correlations for the seasonally adjusted data are 0.61 for creation, 0.77 for destruction, and 0.53 for reallocation. Visual inspection suggests the correlations are high enough to gauge the general pattern of recent reallocation behavior.

The NAPM data paint a somewhat different picture of gross job flow dynamics during the current expansion. Job destruction and reallocation have declined steadily since the last recession, as they typically do during expansions, but they are significantly higher than at this point in previous expansions. A particularly notable difference arises between the LRD and NAPM job destruction rates since the last recession, with NAPM destruction becoming increasingly larger than LRD destruction. A possible explanation for this wedge between LRD and NAPM destruction is that the anecdotes and speculation about nonmanufacturing experiencing greater churning this decade may be true (recall that NAPM includes some nonmanufacturing companies). Another interesting difference about this expansion is that job creation surged twice, peaking in 1994 and 1998. Unlike previous creation surges, these did *not* come on the heels of a large surge in job destruction.

One possible interpretation of this creation-led employment expansion, which has coincided with a surprising plunge in the unemployment rate, is that the economy has been experiencing *favorable* allocative driving forces rather than unfavorable ones. Recall that in allocative theories of fluctuations, the timing of creation and destruction may depend on whether allocative forces raise productivity and profitability in some sectors or lower productivity and profitability in some sectors. In the former case, job creation would rise—provided the available stock of workers was sufficient. And apparently it was. Not only did many unemployed workers become employed, but labor force growth surged simultaneously with the surges in creation. So why did labor force growth surge? Answers to this question, which require detailed investigation beyond the scope of this paper, may provide a better understanding of the nature of recent reallocation.

Predictions for Reallocation and Recessions

Although current rates of job reallocation are moderate, the outlook for job reallocation is less sanguine. Absent a formal model, we cannot provide rigorous forecasts of reallocation. Nevertheless, based on the analysis in this paper, we can provide a judgmental view of the likelihood that job reallocation will play a role in future economic fluctuations.

Figure 16
Comparison of LRD and NAPM Gross Job Flows

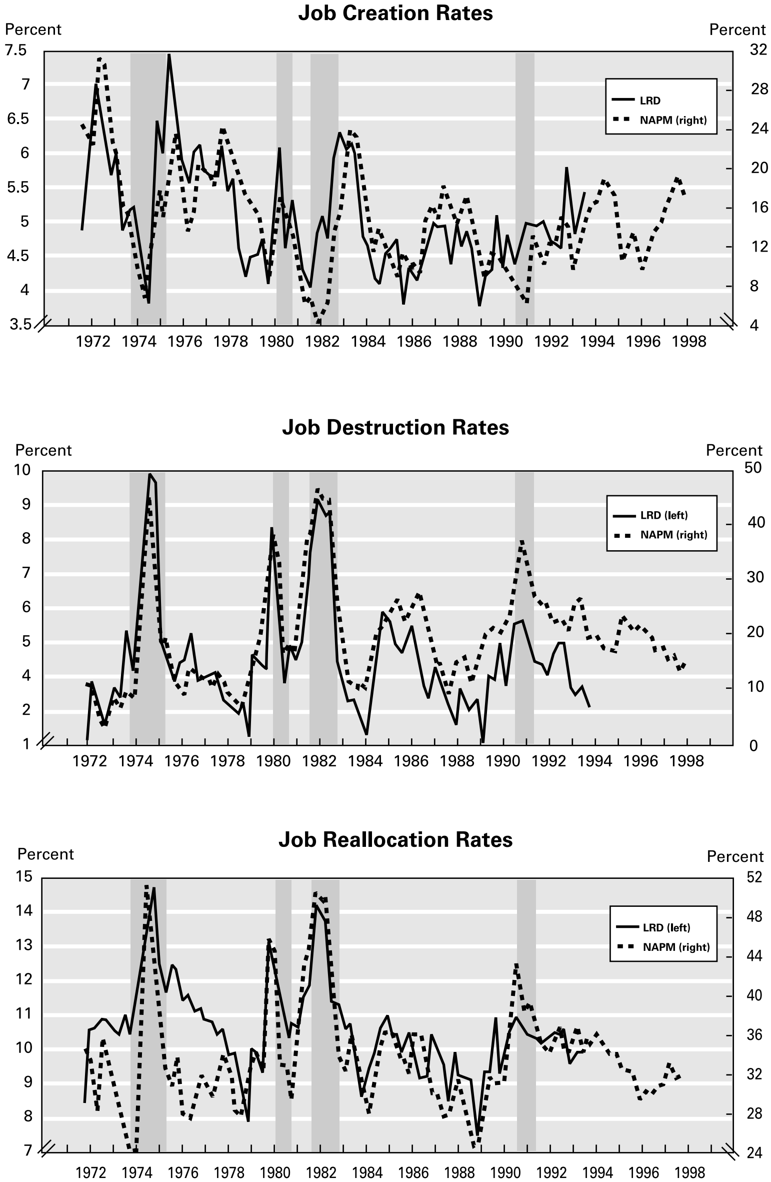

Source: Census Longitudinal Research Database and National Association of Purchasing Management.

Despite the relatively mild reallocation during the most recent recession, the propensity for reallocation to occur during recessions does not appear to have declined over the postwar period. Thus, it seems safe to predict that reallocation will rise again when the next recession occurs. Of course, if aggregate shocks cause business cycles and reallocation, then reallocation is irrelevant for predicting recessions, although the nature and extent of reallocation may be important for understanding the appropriate policy responses to the recession. But if reallocation causes business cycles, it is critical to be on the lookout for factors determining reallocation.

Several factors could induce job reallocation in the foreseeable future. Oil prices, now low and stable, always have the potential for disruptive increases. Regional conditions are currently quite evenly balanced throughout the country, but events could change that (although it is hard to predict what events). The Asian economic crisis could sow the seeds of reallocation in at least two ways. First, the violent swings in exchange rates affect the relative prices of goods imported from, and exported to, Asia. Furthermore, Asian demand for U.S. exports is already sagging along with Asia's output. Together, these developments directly affect plants and sectors that trade with Asia differentially from those that do not trade. Second, these direct effects, especially the price effects, are likely to indirectly impact plants that compete in the same industries with exporting and importing plants. Indeed, more generally, increasing globalization of the U.S. economy brings with it new dimensions for allocative driving forces.

Another source of potential reallocative activity is the considerable amount of retooling required in computer-intensive plants and sectors associated with the Year 2000 (Y2K) problem. Like the case of oil as a factor of production, the usage of computers in production is unevenly distributed across industries (see McGuckin and Stiroh 1998). Consequently, investment patterns and production disruptions—*if* they happen—will be quite uneven across industries. If work on this problem is spread out gradually, the reallocative activity may not cause fluctuations. But the point-in-time nature of the problem at least admits the possibility of significant trouble. Also, the recent humorous story about loads of personal pagers going down as a result of satellite trouble contains a serious strand as well. The increasing reliance on such high-tech devices in a wide range of telecommunications applications produces a vulnerability to concentrated problems. Recent government mandating of a changeover to high-definition TV falls into this category as well.

To summarize, we ask: Will the variables that induce high rates of job reallocation change again dramatically in the future? Probably. With the possible exception of monetary policy, which during the past 15 years may have become increasingly proficient in adjusting to economic conditions, nothing has occurred recently to lead us to believe that the factors determining reallocation have somehow become more stable or less likely to cause allocative fluctuations. Thus, proponents of the idea that reallocation causes business cycle fluctuations would almost surely agree that business cycles are *not* dead nor likely dampened. Instead, the relative calm since the early 1980s has been the result of relatively mild changes in the incentives to reallocate, incentives that generally cannot be controlled well by government policy. It is likely that, sooner or later, incentives to reallocate will arise and that the process of reallocation may cause or contribute to a recession.

References

Abel, Andrew B. and Janice C. Eberly. 1997. "The Mix and Scale of Factors with Irreversibility and Fixed Costs of Investment." NBER Working Paper No. 6148.

Abraham, Katharine G. and Lawrence F. Katz. 1986. "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *Journal of Political Economy*, 94(3), pp. 507–22.

Anderson, Patricia, M. and Bruce D. Meyer. 1994. "The Extent and Consequences of Job Turnover." *Brookings Papers on Economic Activity: Microeconomics*, pp. 177–236.

Baily, Martin Neil, Eric J. Bartelsman, and John Haltiwanger. 1996. "Labor Productivity: Structural Change and Cyclical Dynamics." NBER Working Paper No. 5503.

Baily, Martin Neil, Charles Hulten, and David Campbell. 1992. "Productivity Dynamics in Manufacturing Plants." *Brookings Papers on Economic Activity: Microeconomics,* pp. 187–249.

Bartelsman, Eric J. and Wayne Gray. 1996. "The NBER Manufacturing Productivity Database." NBER Technical Working Paper No. 205.

Bertola, Guiseppe and Ricardo J. Caballero. 1990. "Kinked Adjustment Costs and Aggregate Dynamics." In Olivier Blanchard and Stanley Fischer, eds., *NBER Macroeconomic Annual*. Cambridge, MA and London: The MIT Press.

Black, Fischer. 1987. *Business Cycles and Equilibrium*. New York and Oxford: Basil Blackwell.

Blanchard, Olivier and Peter Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity*, 1:1989, pp. 1–60.

_____. 1990. "The Cyclical Behavior of the Gross Flows of U.S. Workers." *Brookings Papers on Economic Activity*, 2:1990, pp. 85–143.

Blanchard, Olivier and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, 1, pp. 1–61.

Bresnahan, Timothy F. and Valerie A. Ramey. 1993. "Segment Shifts and Capacity Utilization in the U.S. Automobile Industry." *The American Economic Review*, 83(2), pp. 213–18.

Caballero, Ricardo J. and Eduardo M.R.A. Engel. 1993. "Microeconomic Adjustment Hazards and Aggregate Dynamics." *Quarterly Journal of Economics*, 108(2), pp. 359–83.

Caballero, Ricardo J., Eduardo M.R.A. Engel, and John C. Haltiwanger. 1995. "Plant-Level Adjustment and Aggregate Investment Dynamics." *Brookings Papers on Economic Activity*, 2:1995, pp. 1–39.

_____. 1997. "Aggregate Employment Dynamics: Building from Microeconomic Evidence." *The American Economic Review*, 87(1), pp. 115–37.

Caballero, Ricardo J. and Mohamad L. Hammour. 1994. "The Cleansing Effect of Recessions." *The American Economic Review*, 84(5), pp. 1350–68.

_____. 1996. "On the Timing and Efficiency of Creative Destruction." *Quarterly Journal of Economics*, 111(3), pp. 805–52.

_____. 1998. "Improper Churn: Social Costs and Macroeconomic Consequences." Unpublished paper.

Campbell, Jeffrey R. and Jonas D.M. Fisher. 1998. "Aggregate Employment Fluctuations with Microeconomic Asymmetries." Unpublished paper.

Campbell, Jeffrey R. and Kenneth Kuttner. 1996. "Macroeconomic Effects of Employment Reallocation." *Carnegie-Rochester Conference Series on Public Policy,* 44, pp. 87–116.

Cooper, Russell W. and John Haltiwanger. 1993. "The Aggregate Implications of Machine Replacement: Theory and Evidence." *The American Economic Review* 83(3), pp. 360–82.
_____. 1998. "On the Nature of Adjustment Costs for Capital and Labor." Unpublished paper.
Davis, Steven J. 1987. "Fluctuations in the Pace of Labor Reallocation." *Carnegie-Rochester Conference Series on Public Policy* 27, pp. 335–402.
Davis, Steven J. and John Haltiwanger. 1990. "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications." In Olivier Blanchard and Stanley Fisher, eds., *NBER Macroeconomics Annual*, pp. 123–68. Cambridge, MA and London: The MIT Press.
_____. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Quarterly Journal of Economics*, 107(3), pp. 819–63
_____. 1995. "Measuring Gross Worker and Job Flows." NBER Working Paper No. 5133.
_____. 1996. "Driving Forces and Employment Fluctuations." NBER Working Paper No. 5775.
_____. 1997. "Sectoral Job Creation and Destruction Responses to Oil Price Changes and Other Shocks." Unpublished paper.
_____. 1998. "Gross Job Flows." In Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics* (forthcoming).
Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1990. "Published versus sample statistics from the asm: Implications for the l & d." In *1990 Proceedings of the Business and Economic Statistics Section*, American Statistical Association, pp. 52–61.
_____. 1996. *Job Creation and Destruction*. Cambridge, MA: The MIT Press.
Davis, Steven J., Prakash Loungani, and Ramamohan Mahidhara. 1997. "Regional Labor Fluctuations: Oil Shocks, Military Spending, and Other Driving Forces." Unpublished paper.
Debelle, Guy and Owen Lamont. 1997. "Relative Price Variability and Inflation: Evidence from U.S. Cities." *Journal of Political Economy*, 105(1), pp. 132–52.
Diamond, Peter. 1981. "Mobility Costs, Frictional Unemployment, and Efficiency." *Journal of Political Economy*, 89(4), pp. 798–812.
Dunne, Timothy, John Haltiwanger, and Kenneth R. Troske. 1996. "Technology and Jobs: Secular Changes and Cyclical Dynamics." NBER Working Paper No. 5656.
Dunne, Timothy, Mark Roberts, and Larry Samuelson. 1989. "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, 104(4), pp. 671–98.
Foote, Christopher L. 1998. "Trend Employment Growth and the Bunching of Job Creation and Destruction." *Quarterly Journal of Economics* (forthcoming).
Foster, Lucia, John Haltiwanger, and C.J. Krizan. 1998. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." Unpublished paper.
Garibaldi, Pietro. 1997. "The Asymmetric Effects of Monetary Policy on Job Creation and Destruction." *IMF Staff Papers*, 44(4), pp. 557–84.
Gertler, Mark and Simon Gilchrist. 1994. "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms." *Quarterly Journal of Economics*, 109(2), pp. 309–40.
Gourinchas, Pierre-Olivier. 1998. "Exchange Rates, Job Creation and Destruction." In Olivier Blanchard and Stanley Fisher, eds., *NBER Macroeconomics Annual* (forthcoming).
Granger, C.W.J. 1969. "Investigating Causal Relations by Econometric Models and Cross-Spectral Models." *Econometrica*, 37, pp. 424–38.
Greenwood, Jeremy, Glenn M. MacDonald, and Guang-Jia Zhang. 1996. "The Cyclical Behavior of Job Creation and Job Destruction: A Sectoral Model." *Economic Theory*, 7(1), pp. 95–112.
Hall, Robert E. 1979. "A Theory of the Natural Unemployment Rate and the Duration of Employment." *Journal of Monetary Economics*, 5(2), pp. 153–69.
_____. 1991. "Labor Demand, Labor Supply, and Employment Volatility." In Olivier Blanchard and Stanley Fisher, eds., *NBER Macroeconomics Annual*, pp. 17–47. Cambridge, MA and London: The MIT Press.
_____. 1995. "Lost Jobs." *Brookings Papers on Economic Activity*, 1:1995, pp. 221–56.
_____. 1997a. "Impulses, Amplification, and Persistence." In John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics* (forthcoming).

_____. 1997b. "The Temporal Concentration of Job Destruction and Inventory Liquidation: A Theory of Recessions." Unpublished paper.
_____. 1998. "Labor Market Frictions and Employment Fluctuations." NBER Working Paper No. 6501.
Haltiwanger, John and Scott Schuh. 1998. "Macroeconomic Implications of the Relationship Between Plant and Industry Job Flows." Unpublished paper.
Hamermesh, Daniel S. 1989. "Labor Demand and the Structure of Adjustment Costs." *The American Economic Review*, 79(4), pp. 674–89.
Hamilton, James D. 1988. "A Neoclassical Model of Unemployment and the Business Cycle." *Journal of Political Economy*, 96(3), pp. 593–617.
Hopenhayn, Hugo and Richard Rogerson. 1993. "Job Turnover and Policy Evaluation: A General Equilibrium Analysis." *Journal of Political Economy*, 101(5), October, pp. 915–38.
Horvath, Michael. 1998a. "Business Cycles and the Failure of Marginal Firms." Unpublished paper.
_____. 1998b. "Cyclicality and Sectoral Linkages: Aggregate Fluctuations from Independent Sectoral Shocks." Unpublished paper.
_____. 1998c. "Sectoral Shocks and Aggregate Fluctuations." Unpublished paper.
Johnson, Harry. 1962. *The Two-Sector Model of General Equilibrium.* London: George Allen & Unwin Ltd.
Kydland, Finn and Edward Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica*, 50, pp. 1345–70.
Lane, Julia, Alan Isaac, and David Stevens. 1994. "Job Flows, Worker Flows, and Churning." Centre for Economic Policy Research Discussion Paper No. 1125.
Leonard, Jonathan S. 1987. "In the Wrong Place at the Wrong Time: The Extent of Frictional and Structural Unemployment." In Kevin Lang and Jonathan Leonard, eds., *Unemployment & the Structure of Labor Markets*. New York: Basil Blackwell.
Lilien, David. 1982. "Sectoral Shifts and Cyclical Unemployment. *Journal of Political Economy*, 90(4), pp. 777–93.
Liu, Lili and James R. Tybout. 1996. "Productivity Growth in Chile and Colombia: The Role of Entry, Exit, and Learning." In Mark J. Roberts and James R. Tybout, eds., *Industrial Evolution in Developing Countries: Micro Patterns of Turnover, Productivity, and Market Structure*. New York: Oxford University Press.
Long, John B. and Charles Plosser. 1983. "Real Business Cycles." *Journal of Political Economy*, 91(1), pp. 39–69.
Loungani, Prakash. 1986. "Oil Price Shocks and the Dispersion Hypothesis." *The Review of Economics and Statistics*, 62(3), pp. 536–39.
Loungani, Prakash, Mark Rush, and William Tave. 1990. "Stock Market Dispersion and Unemployment." *Journal of Monetary Economics*, 25(3), pp. 367–88.
Loungani, Prakash and Bharat Trehan. 1997. "Explaining Unemployment: Sectoral vs. Aggregate Shocks." Federal Reserve Bank of San Francisco *Economic Review*, (1), pp. 3–15.
Lucas, Robert E., Jr. and Edward Prescott. 1974. "Equilibrium Search and Unemployment." *Journal of Economic Theory*, 7(2), pp. 188–209.
McGuckin, Robert H. and Kevin J. Stiroh. 1998. "Computers, Productivity, and Growth: Explaining the Computer Productivity Paradox." The Conference Board Economic Research Report 1213-98-RR.
Mortensen, Dale T. 1982. "The Matching Process as a Non-Cooperative Bargaining Game." In John McCall, ed., *The Economics of Information and Uncertainty*. Chicago: University of Chicago Press.
_____. 1994. "The Cyclical Behavior of Job and Worker Flows." *Journal of Economic Dynamics and Control*, 18(6), pp. 1121–42.
Mortensen, Dale T. and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, 61(208), pp. 397–415.
Olley, G. Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6), pp. 1263–97.
Ritter, Joseph. 1993. "Measuring Labor Market Dynamics: Gross Flows of Workers and Jobs." Federal Reserve Bank of St. Louis *Review*, 75(6), pp. 39–57 .
_____. 1994. "Job Creation and Destruction: The Dominance of Manufacturing." Federal Reserve Bank of St. Louis *Review*, 76(5), pp. 39–57.

Rogerson, Richard. 1987. "An Equilibrium Model of Sectoral Reallocation." *Journal of Political Economy,* 95(4), pp. 824–34.
Schumpeter, J.A. 1942. *Capitalism and Democracy*. New York: Harper and Brothers.
Sims, Christopher A. 1972. "Money, Income, and Causality." *The American Economic Review*, 62(4), pp. 540–52.
Troske, Kenneth. 1993. "The Dynamic Adjustment Process of Firm Entry and Exit in Manufacturing and Finance, Insurance and Real Estate." *Journal of Law and Economics*, 39(2), pp. 705–35.
Veracierto, Marcelo. 1998. "Plant-Level Irreversible Investment and Equilibrium Business Cycles." Federal Reserve Bank of Chicago *Working Paper* WP-98–1.