Understanding the Life-Cycle of a Manufacturing Plant

BY AUBHIK KHAN

What determines whether a manufacturing plant survives? Is it access to credit markets? Or does learning about plants' profitability over time determine survival? Should government policy play a role in helping plants survive? In this article, Aubhik Khan discusses the collateral and learning views as two possible explanations for a typical plant's life-cycle. He concludes that although it remains unclear as to which explanation is the more relevant, the two views have very different implications for what government *can* do and what it *should* do.

Local news reports tend to highlight major plant closings whenever they occur. Sometimes, when the number of workers displaced is large, such events are even reported in the national news.[1] Such reports indicate the importance of plant openings and closings to the general public.

Actually, such events are of interest to economists also because they help economists understand changes in employment and investment.[2] In recent years, a large quantity of information on the behavior of U.S. manufacturing plants, or factories, has become available through the census. These data have created an opportunity for economists to improve their understanding of both the distribution of production across manufacturing plants and the evolution of individual plants over time.

The census data indicate that large changes in economic activity are common within plants. Moreover, as some plants are growing, others are declining, and production is reallocated from declining to growing plants. This reallocation of production requires an associated reallocation of both employment and investment. Hence, this simultaneous growth and decline of plants implies a concomitant investment in and scrapping of equipment and hiring and firing of workers. When economists first used census data to study changes in plants, they were surprised to learn how often plants that are undergoing growth and decline are not only in the same region but also in the same industry. In some cases, they are even owned by the same firm.

Census data also indicate a typical pattern to a plant's life-cycle. Most new plants are small, and they begin with relatively low levels of production. New plants also appear to be riskier. They are very volatile — swings in economic activity are large and frequent — and they tend to have unusually high failure rates. Over time, those that survive grow larger. The survivors increase both their number of employees and stocks of equipment. Survivors also become less risky, and their production exhibits less month-to-month volatility.

This article summarizes some of what we know about plants and explores two explanations of the typical plant's life-cycle: the collateral view and the learning view. The collateral view stresses the importance of a firm's access to credit markets, which is generally believed to be more constrained in

[1] In the 1980s, General Motors closed its production plant in Flint, Michigan. In a town of 150,000 people, this closing, which was even reported in the *New York Times*, resulted in the loss of 40,000 jobs and Michael Moore's 1989 documentary film, "Roger and Me."

[2] Economists refer to a place where production takes place, for example, a factory, as a plant. Firms own plants, and there exist both single- and multi-plant firms. For example, Lockheed Martin operates many plants, including several located in Fort Worth, Texas; Marietta, Georgia; and Palmdale, California. The musical instrument company Mid-East Manufacturing operates at a single location in West Melbourne, Florida. Small plants are ones with fewer than 250 employees; large plants are those with more than 250 employees.

Aubhik Khan is a senior economist in the Research Department of the Philadelphia Fed.

smaller, newer firms. The learning view suggests that as managers learn which plants can be run more profitably than others, this influences employment patterns across plants and may be responsible for the observed life-cycle of the typical plant.

While both of these views are able to explain the same data, they suggest very different approaches for monetary policy. If some firms have trouble obtaining loans during downturns in demand, for example, in a recession, a central bank may stimulate economic activity by cutting interest rates, thereby facilitating firms' borrowing. However, if the differences across plants are not primarily caused by differences in the collateral of the firms that operate them, then the plant-level data do not, in and of themselves, provide evidence of a channel for monetary policy to affect the economy.

WHAT THE JOB DATA REVEAL

To discuss differences across plants and how individual plants change over time, we need to measure their output. This is difficult because of the formidable problem of finding direct measures of plant output that can be usefully compared across plants. (See *Difficulties in Measuring Output at the Plant Level.*) However, recent work at the U.S. Census Bureau offers a very comprehensive look at a large majority of U.S. manufacturing plants. The Census Bureau collects data on the number of jobs in a plant. As a plant increases in size and produces more, it will hire more workers and buy more equipment. Hence, job creation should be a good measure of a plant's output growth. Alternatively, when a plant reduces production, the number of jobs destroyed serves as a measure of its output decline.

These data on plant-level employment reveal startling insights about the manufacturing sector. There is constant, simultaneous expansion and contraction within industries; that is, some plants are hiring workers while others are laying workers off.

The following example may help illustrate the importance of this simultaneous job creation and destruction. In 1973, net U.S. manufacturing employment growth was 5.7 percent.[3] Now consider two alternatives, both of which might have led to this growth in employment. In the first case, all plants engage in job creation. As a result, the total number of jobs in the manufacturing sector rises by 5.7 percent because each plant increases its employment by this amount. In the second case, some plants engage in job creation, but others engage in job destruction. Assume that the job creation driven by plants' expanding employment is 11.8 percent while the level of job destruction at declining plants is 6.1 percent. In this instance, net employment growth — the difference between job creation and destruction — was also 5.7 percent.

Excess job reallocation — the sum of creation plus destruction minus net employment growth — was 12.2 percent.[4] This is one measure of churning in the labor market — that is, the amount of job reallocation beyond that explained by job growth alone.

While the net effect on total employment is the same in either case, the first case involves far less social disruption and far less short-term unemployment than the second one. (Indeed, there is no excess job reallocation in the first case.) Nonetheless, the second case describes what actually occurred in 1973. This type of simultaneous expansion and contraction across plants and the resultant reallocation of employment often impose large adjustments on workers, local schools, and housing markets.[5] The greater the employment reallocation across plants, the larger the size of these effects on local communities. The data indicate substantial reallocation.

Between 1973 and 1993, the last year for which data are available, job creation averaged about 9 percent while job destruction averaged 10 percent. In their book, *Job Creation and Destruction*, Steven J. Davis, John C. Haltiwanger, and Scott Schuh show that this level of excess job reallocation is amazingly consistent across very diverse industries such as food, rubber, and electric machinery.[6] While there is some inter-industry variation, relatively little of the

[3] All percentages reported here are fractions of total U.S. manufacturing employment.

[4] When net employment growth is negative, excess job reallocation is the sum of job creation and job destruction plus net employment growth.

[5] See Davis, Haltiwanger, and Schuh (1996), page 30, for additional discussion.

[6] See Table 3.1, page 39, in Davis, Haltiwanger, and Schuh.

excess job reallocation we observe can be explained by structural adjustments across industries as they respond to shifting patterns of supply and demand. In other words, job creation in plants in the rubber industry and job destruction in plants in the paper industry do not explain most of the simultaneous expansion and contraction we observe. Much of the concurrent job creation and destruction occurs within the same industry.

Simultaneous job creation and destruction occurs not only within the same industry but also within the same region. While manufacturing employment shrank most in the middle Atlantic region and grew fastest in the Mountain region over the sample period of 1972–88, all regions exhibited excess job reallocation of 12 percent or more.[7] Therefore, the concurrent job creation and destruction we observe cannot be explained by differences in economic growth across regions. While workers may well be moving from one region of the country to another, this movement is not driving most of the excess job reallocation. Indeed, regional differences in rates of job creation and destruction are smaller than cross-industry differences.[8]

Davis, Haltiwanger, and Schuh argue that most of the excess job reallocation, which averages over 15 percent of employment in a typical year, must be caused by differences between plants that are not the result of changes in the industries or regions in which they operate.[9] Moreover, much of the simultaneous job creation and destruction we observe involves large changes in plant size. More than two-thirds of total creation and destruction occurred in plants that adjusted employment by more than 25 percent. *Sharp changes in employment are the rule, not the exception.*

What lies behind this simultaneous job creation and destruction? Two explanatory factors are the plant's size — that is, the number of employees — and the plant's age.[10] New plants tend to be small. The census data show that in the first decade of their existence, plants exhibit substantial growth. Start-ups — which account for about 7 percent of all existing plants in a typical year — are roughly one-tenth the size of the average plant in their industry. Over the next 10 years, should they survive, they reach average size. However, failure rates are much higher for younger and smaller plants than for larger and more mature plants.

We can examine the rates of job creation, job destruction, and excess reallocation over plants of different sizes (Table 1). The table illustrates that small plants have higher rates of job creation, which explains why politicians and

Difficulties in Measuring Output At the Plant Level

To examine differences across plants or changes within a plant over time, we need to measure plants' size. One possibility is to measure the output of each plant in the economy. However, different plants produce different goods, and it is not always easy to compare them. You cannot determine if a tire factory is larger than a chemical plant by comparing the number of tires produced by the first to the drums of paint shipped from the second.

Of course, you could use the prices of tires and paint to compare the value of their different total outputs. This requires access to data on the value of sales and costs incurred by each plant. Changes in plants would then be defined to be changes in the value of their production. Unfortunately, the Longitudinal Research Datafile, which is the Census Bureau's most comprehensive data on U.S. manufacturing plants, does not contain plant-level sales data.

Plants are secretive about the prices at which they sell goods. One reason may be that often the same plant sells the same commodity at different prices; for example, a plant may offer quantity discounts to large customers. In other instances, some plants do not sell their output in a market that determines prices. Instead, they produce specialized intermediate goods for a single firm that owns both the producing and the consuming plants. As these commodities are neither bought nor sold outside the firm, it is particularly difficult to judge their price. While it is true that a firm will declare a value to its intermediate inputs, these values are unlikely to be accurate. The prices of these internal inputs are determined by many factors, including tax considerations.

The alternative that many economists have adopted is to use the rich information on employment changes in the Longitudinal Research Datafile. We define changes in plants to be changes in their employment. This would be a perfect alternative if all plants used the same type and amount of labor for each unit of output. But, of course, they do not; therefore, our approach is a compromise based on the availability of data.

[7] See Table 3.3, page 42, in Davis, Haltiwanger, and Schuh.

[8] Davis, Haltiwanger, and Schuh, page 42.

[9] In their 1992 paper, Stephen Davis and John Haltiwanger argue that cross-industry flows explain less than 1 percent of excess job reallocation, and that cross-regional flows explain even less.

[10] As discussed by Thomas Cooley and Vincenzo Quadrini in their forthcoming paper, both size and age are independently important, in a statistical sense, as factors when trying to explain differences in growth or failure rates for plants or firms. The 1989 paper by Timothy Dunne, Mark Robert, and Larry Samuelson finds that both size and age are independently important when examining U.S. manufacturing plant failure and growth rates. Similar evidence for U.S. manufacturing firms is presented in the 1987 papers by David Evans and Bronwyn Hall and also the 1989 paper by David Evans.

lobbyists often tout small businesses as engines of employment growth. However, the table also indicates that small plants have higher rates of job destruction. As a result, *net* employment growth — the difference between job creation and job destruction — shows no evidence that small plants grow faster than large plants. This may raise questions about policies that either directly subsidize small businesses or offer them regulatory relief in an effort to promote employment.

The most striking finding in Table 1 is that excess job reallocation — the level of job reallocation across plants above that attributable to net employment growth — falls with plant size. In other words, small plants undergo a lot more day-to-day fluctuations in employment that don't affect the total number of jobs available. Individual small plants clearly exhibit more fluctuations in employment and, presumably, in their investment and production as well.[11] In this sense, they are riskier than large plants.

More evidence that small plants are riskier is apparent from the finding that jobs created in small plants are more likely to be permanently destroyed than those in large plants. Davis, Haltiwanger, and Schuh calculate that in firms with less than 250 employees, 50 percent of the jobs created in any typical year still exist two years later. For large firms — those employing more than 250 employees — 60 percent of the jobs created in a typical year survive for two or more years.[12]

Plants also show systematic differences due to their age (Table 2). The table shows that young and middle-aged plants have higher rates of job creation and destruction than do mature plants. As a result, they have higher rates of excess job reallocation. Thus, younger plants are more volatile, exhibiting, on average, more dramatic changes in the number of employees than mature plants. Moreover, this is not simply the result of rapid transition to a larger, more stable size. Younger plants exhibit higher rates of both job creation and job destruction. In this sense, they are riskier than mature plants.

[11] Unfortunately no data are available to measure directly investment and production at the plant level.

[12] See Davis, Haltiwanger and Schuh, Table 4.6, page 79.

TABLE 1

Average Levels of Job Creation and Destruction Across Plants of Different Sizes between 1973 and 1988

Employees	Job Creation	Job Destruction	Net Growth	Excess Reallocation
0 – 9	18.7	23.3	-4.5	34.6
20 – 49	13.2	15.3	-2.1	23.6
50 – 99	12.2	13.5	-1.3	21.5
100 – 249	9.6	10.7	-1.1	16.1
250 – 499	7.7	8.7	-1.0	12.5
500 – 999	7.0	7.6	-0.6	10.7
1000 – 2499	6.3	7.3	-1.0	10.2
2500 – 4999	6.1	7.5	-1.3	9.7
5000 and more	5.4	5.6	-0.2	7.7

The numbers reported here are the averages over the years 1973–88. For any year, job creation for plants in a particular size class is the average number of new jobs created at each plant in the size class that created jobs as a fraction of the plant's average employment in the year and the prior year. For any year, job destruction in a particular size class is the average number of jobs destroyed at each plant in the size class that destroyed jobs as a fraction of the plant's average employment in the year and the prior year. Net growth is the difference between job creation and job destruction. For any year, excess reallocation is the sum of job creation and destruction minus the absolute value of net growth across all plants in the size class. All numbers are expressed as percentages. Since each column is an average over time and excess reallocation involves an absolute value calculation, the formula for excess job reallocation does not directly apply to the averages displayed in the other columns. From Steven J. Davis, John C. Haltiwanger, and Scott Schuh, *Job Creation and Destruction*. Cambridge, MA: MIT Press, 1996, Table 4.1, page 61. Reprinted with permission.

We have additional evidence that younger plants are inherently more risky (Table 3). The table shows the percentage of plants of each age group that do not survive to the next age group. We observe that 41 percent of plants shut down before their sixth birthday, while the probability of failure falls as a plant grows older.

In summary, plant-level data lead us to three conclusions. First, the U.S. manufacturing sector experiences a lot of excess job reallocation — far more jobs shift between plants than can be explained by the net change in employment growth. Second, smaller plants exhibit more employment volatility than larger plants. Third, younger plants exhibit more employment volatility and have higher failure rates than older plants.

We now turn to two explanations for these phenomena: the collateral view and the learning view. The collateral view emphasizes the role of a firm's *size* in determining its access to credit. To the extent that small firms mainly operate small plants, we can explain the observed relationship between plant size and employment volatility. The learning view explains the relationship between a plant's *age* and employment volatility as the result of a gradual process of learning how a plant can be operated most efficiently.

THE IMPORTANCE OF COLLATERAL

Anyone who has ever applied for a loan knows the value of collateral. If a lender can secure a potential loan, he is more likely to make the loan. Home-equity loans are routinely offered to borrowers who might not otherwise get a loan without offering their homes as collateral. Moreover, the more collateral a borrower has, the lower the interest rate on the loan is likely to be.

By their very nature, small firms have fewer assets. Since they don't have as much marketable property to use for collateral, they may find borrowing more difficult. This situation has several implications. First, a small firm is unable to borrow sufficiently to invest as much as it would like. Instead, it has to rely on sales to finance investment; therefore, most firm-level investment is financed through retained earnings. Over time, as small firms reinvest their profits in the enterprise, they accumulate more equity and they grow. Note that the collateral view can easily explain why new firms tend to be small: they have to generate profits to finance investment, and this takes time. As a result, they cannot begin production at their ideal size, but instead reach it slowly over time.[13]

Another implication of the collateral view is that any short-term fall in sales will tend to have a much larger impact on small firms because these firms are either unable to obtain or cannot afford sufficient loans to sustain their employment levels. Workers have to be laid off, machines sold off, and operating hours shortened. So a small firm's lack of access to short-term credit means that a temporary fall in sales will lead to a disproportionate decline in its operations. Moreover, subsequent increases in sales will generate a sharp rise in employment at a small firm.

By contrast, a large firm, which is not subject to severe borrowing constraints, is able to sustain a short-term drop in demand and to continue operating most of its plants and retain its

[13] There are many papers that examine the role of collateral, for example, the paper by Mark Gertler and Simon Gilchrist; the paper by Ben Bernanke, Mark Gertler, and Simon Gilchrist; and the one by Thomas Cooley and Vincenzo Quadrini.

TABLE 2

Average Job Creation and Destruction Across Plants Of Different Ages between 1973 and 1988

Plant Age in Years	Job Creation	Job Destruction	Net Growth	Excess Reallocation
Young (0-1)	45.8	12.5	33.3	25.1
Middle-Aged (2–10)	12.3	13.3	-1.0	21.0
Mature (10+)	6.9	9.4	-2.5	12.4

See Table 1 for definitions of job creation, destruction, net growth, and excess reallocation. All ages are in years. From Steven J. Davis, John C. Haltiwanger, and Scott Schuh, *Job Creation and Destruction*. Cambridge, MA: MIT Press, 1996, Table 4.5 page 77. Reprinted with permission.

TABLE 3

Failure Rate by Plant Age

Plant age	1-5	6-10	11-15
Exit rate	0.41	0.35	0.30

From Timothy Dunne, Mark J. Roberts, and Larry Samuelson, "The Growth and Failure of U.S. Manufacturing Plants," *Quarterly Journal of Economics*, 104(4), 1989, pp. 671-98. Reprinted with permission.

workers through a temporary fall in demand. When demand recovers, returning to full production is a simple matter, since a large firm does not have to bear the costs of first reducing, then increasing its scale of production. Hence, the collateral view can explain the higher volatility of job creation and destruction at small firms.

LEARNING YOUR TRUE VALUE

Another explanation for the life-cycle of a plant emphasizes differences in plants' long-run profitability. This explanation assumes that managers are able to discern long-run profitability only through experience gained by operating the plant. This long-run profitability, which determines the plant's value and its most efficient scale of operation, is difficult to discern because shocks — such as unexpected fluctuations in input prices, delays in supplier deliveries, workers quitting, unexpected changes in demand, changes in regulation, and entry of a new competitor — constantly buffet plants. All of these shocks temporarily influence profits and thus make long-run profitability harder to determine.[14] Over time, however, as managers gain experience with operating a plant, they can better determine long-run profitability because temporary shocks will tend to offset one another and managers will learn what to expect.

This slow process of learning can explain why younger plants are riskier. When there is little information about long-run profitability, managers are very responsive to new information. In a plant's early years, any observed change in profits is likely to have a large influence on managers' assessment of the plant's value. Their lack of experience leads them to place considerable weight on each new piece of evidence. This can lead to sharp adjustments in employment as managers re-evaluate the plant's long-run profitability. However, as more and more evidence of the plant's profitability accumulates, the influence of additional evidence lessens. Managers will adjust employment far less in response to new information.

In addition to predicting that older plants will have less employment volatility, the learning view can also explain why older plants are larger. Over time, profitable plants are more likely to continue operating, while less profitable plants shut down. Being more productive, more profitable plants typically hire more workers, that is, they are larger according to our measure of plant size.

EVALUATING THE COLLATERAL AND LEARNING VIEWS

Some Difficulties with the Collateral View. Whether we consider the role of collateral or the importance of learning, we gain insights into why a plant changes over time. We are able to explain differences across plants of different ages and sizes. However, neither explanation really addresses the question of why plants are risky or why plant employment exhibits so much variation in the first place. Both views make assumptions about the magnitudes of the shocks that hit plants, but do not explain why the shocks occur.

The collateral view assumes that large firms have more collateral, relative to their borrowing needs, than small firms. However, while large firms may possess far more physical capital than small ones, most of the equipment and structures may be specialized and thus not very useful as collateral. If the borrower defaults on a loan, the lender can sell such capital only at a large discount because it is not very useful to other firms.

Whether we consider the role of collateral or the importance of learning, we gain insights into why a plant changes over time.

Again, the collateral view was developed to explain *firm* growth, not *plant* growth: a plant is a physical location where production occurs; a firm is a collection of property under common ownership; and a large firm may well operate multiple small plants. However, to the extent that small firms *typically* operate small plants and large firms *typically* operate large plants, this may not present a serious problem.

Empirical Evidence. The main difference between the collateral and learning views is their prediction about borrowing by firms. The learning view suggests that if you gave $500 to a small firm's manager, he would be as likely to invest it elsewhere as to purchase more plant and equipment for his firm. Remember that according to this view, the small firm's manager faces no difficulty in meeting his firm's borrowing needs. Therefore, the return on investing in his own firm won't necessarily exceed the return he could expect from investing in another firm. In contrast, the collateral theory predicts that he would invest the extra funds in his own firm because, under this view, a small firm's ability to borrow is constrained.

Empirical evidence helps us to distinguish between these views. However, the evidence is controversial. If firms are competitive and don't face borrowing constraints due to inadequate collateral, the profitability of their potential investment projects is measured by the market value of the firm

[14] Boyan Jovanovic's article offers an excellent explanation of the learning view.

per unit of its capital stock, and that is the sole determinant of their investment spending.[15] As Robert Chirinko describes in his survey, there is a large literature that finds that firms' investment spending is affected by their cash flow as well as by the market valuation of their investment opportunities. Thus, if cash flow matters, this would imply that firms' borrowing is constrained by a lack of sufficient collateral, and consequently, when managers have excess cash flow, they invest it in their firms.

However, recent work by economists Andrew Abel and Janice Eberly and Russell Cooper and Joao Ejarque argues against this conclusion, suggesting instead that the results indicate greater market power on the part of firms than previously assumed. When firms have market power to set their own prices, these economists show that a firm's market value – even if we assume that the market correctly values the firm's investment prospects – is not a complete determinant of investment spending. In such cases, cash flow, given its relationship to other important determinants of investment such as profitability, will be important for explaining investment even without borrowing constraints.

In their forthcoming paper, economists Thomas Cooley and Vincenzo Quadrini argue that properly matching the data on employment changes across plants requires a model with elements of both the learning and collateral theories. While the learning view can explain how employment volatility depends on a firm's age, it cannot address why employment volatility depends on a firm's size. They note that the data present evidence that both age and size are independently important in determining a plant's riskiness.

In sum, although recent research has called into question earlier empirical evidence for the collateral view, it appears the collateral view and the learning view may be complementary explanations of plant-level data.

POLICY IMPLICATIONS OF THE TWO VIEWS

It remains unclear as to which explanation is more relevant to the actual evolution of plants. While these theories help us understand plant dynamics, they also differ in important ways with respect to their implications for government policy. They have very different implications about what government *can* do and what it *should* do. A proponent of the collateral view might argue that since government *can* do quite a lot during recessions, it *should* help small firms by lowering the cost of borrowing.[16] This argument supports the collateral view's contention that when lending declines, as it does in recessions, the incidence of the reduction in credit falls disproportionately on smaller firms with less collateral. However, by reducing interest rates, the monetary authority can ease the costs of borrowing and help small firms survive.

A subscriber to the learning theory will argue that the evolution of plants is not related to financial market characteristics that make borrowing more difficult for small establishments. As such, it provides evidence of neither a role nor a channel for monetary policy to affect the economy. Given the important differences in policy implications, it will be important to develop

[15] Strictly speaking, the value of the next dollar invested in the firm determines its level of investment in competitive markets.

[16] Recall that financial markets discriminate against small firms. So the collateral view implies that government can help small firms by lowering the cost of financing.

more empirical evidence to better distinguish between these two theories.

SUMMARY

Plants evolve over their lives. Typically, they start out as relatively small factories, and employment fluctuates sharply in the early years. In contrast, older plants are not only larger, they are also far less volatile.

There are two leading explanations of plant evolution. The first emphasizes the impact of the lack of collateral and, thus, the difficulties of obtaining credit for new plants. This lack of collateral makes new plants less able to borrow, to weather a temporary decline in earnings. And this inability to withstand a temporary decline makes new plants riskier.

The second view emphasizes the importance of learning about the profitability of a plant slowly over time. Managers operate new plants with less confidence, and they are quick to make large changes in the scale of their operation. These frequent, large changes make new plants riskier.

While either the collateral or learning view may explain the observed differences between plants of different ages and sizes, the theories differ sharply in their implications for monetary policy. The collateral view provides support for the existence of a channel through which changes in interest rates could affect a small firm's ability to borrow and, therefore, its chances to survive. The learning view sees no role for monetary policy to affect smaller plants. These differences indicate the need for further empirical tests of the two views.

REFERENCES

Abel, Andrew, and Janice C. Eberly. "Q Theory Without Adjustment Costs and Cash Flow Effects Without Financing Constraints," unpublished research available from http://finance.wharton.upenn.edu/~abel/, 2001.

Ben S. Bernanke, Mark Gertler, and Simon Gilchrist. "The Financial Accelerator in a Quantitative Business Cycle Framework, " in John B. Taylor and Michael Woodford (eds.): *Handbook of Macroeconomics* Volume 1c. New York: Elsevier Science, North-Holland, 1999.

Chirinko, Robert S. "Business Fixed Investment Spending: Modeling Strategies, Empirical Results, and Policy Implications," *Journal of Economic Literature*, 3, 1993, pp. 1875-1911.

Cooley, Thomas F., and Vincenzo Quadrini. "Financial Markets and Firm Dynamics," forthcoming in *American Economic Review*.

Cooper, Russell, and Joao Ejarque. "Exhuming Q: Market Power vs. Capital Market Imperfections," Boston University Working Paper, 1999.

Davis, Steven J., and John C. Haltiwanger. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107(5), 1989, pp. 819-63.

Davis, Steven J., John C. Haltiwanger, and Scott Schuh. *Job Creation and Destruction.* Cambridge, MA: The MIT Press, 1996.

Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. "The Growth and Failure of U.S. Manufacturing Plants," *Quarterly Journal of Economics*, 104(4), 1989, pp. 671-98.

Evans, David S. "Tests of Alternative Theories of Firm Growth," *Journal of Political Economy*, 95(4), 1989, pp. 657-74.

Evans, David S. "The Relation Between Firm Growth, Size, and Age: Estimates for 100 Manufacturing Industries," *Journal of Industrial Economics*, 35(4), pp. 567-81.

Gertler, Mark, and Simon Gilchrist. "The Role of Credit Market Imperfections in the Monetary Transmission Mechanism: Arguments and Evidence," *Scandinavian Journal of Economics,* 95(1), 1993, pp. 43-64.

Hall, Bronwyn. "The Relation Between Firm Size and Firm Growth in the U.S. Manufacturing Sector," *Journal of Industrial Economics,* 35(4), 1989, pp. 583-606.

Jovanovic, Boyan. "Selection and the Evolution of Industry," *Econometrica,* 50(3), pp. 649-70.