REVIEW OF ECONOMIC DYNAMICS **1,** 409–429 (1998)
ARTICLE NO. RD980011

Capital Adjustment Patterns in Manufacturing Plants*

Mark Doms

Board of Governors of the Federal Reserve System

and

Timothy Dunne

University of Oklahoma and Office of the Chief Economist, U.S. Census Bureau

Received June 1, 1997

A common result from altering several fundamental assumptions of the neoclassical investment model with convex adjustment costs is that investment may occur in lumpy episodes. This paper takes a step back and asks ''How lumpy is investment?'' We answer this question by documenting the distributions of investment and capital adjustment for a sample of over 13,700 manufacturing plants drawn from over 300 four-digit industries. We find that many plants do undergo large investment episodes; however, there is tremendous variation across plants in their capital accumulation patterns. This paper explores how the variation in capital accumulation patterns vary by observable plant and firm characteristics, and how large investment episodes at the plant level transmit into fluctuations in aggregate investment. *Journal of Economic Literature* Classification Numbers: D24, L6, E22.

I. INTRODUCTION

Among Michael Gort's many contributions to economics is his early work using establishment-level data at the U.S. Census Bureau. Professor Gort realized early on that aggregate statistics mask important underlying

* We are grateful for comments from Ricardo Caballero, John Haltiwanger, seminar participants from the University of Maryland, Washington University, and the Census Bureau, and for helpful comments from three referees. We are particularly indebted to Jeffrey Campbell for providing detailed suggestions and help with the simulations presented in the paper. We also thank Lawrence Rutherford for excellent research assistance. Dunne thanks NSF for financial support. The opinions and conclusions expressed in this paper are those of the authors and do not necessarily represent those of the U.S. Bureau of the Census or of the Federal Reserve System.

1094-2025/98 $25.00

dynamics that belie the aggregate changes, and that to truly understand the dynamics within industries, one has to examine the underlying micro data. In his 1963 paper [23], ‘‘Analysis of Stability and Change in Market Shares,’’ Professor Gort explored the extent to which the market share of firms that make up published concentration ratios change over time. After all, if the concentration ratio for a particular industry remains high and stable over time, it does not necessarily imply that the industry is stagnant and controlled by a small handful of dominant firms. In fact, the industry could be extremely competitive with market share amongst the firms changing quite markedly, yet the published concentration ratios would not convey this information. To resolve this issue, which required access to firm-level data on market shares, Professor Gort utilized the raw micro data files at the U.S. Census Bureau. He was one of the first economists to exploit establishment-level data files at the U.S. Census Bureau for economic research. Moreover, 30 years later, Professor Gort returned to the U.S. Census Bureau to undertake a project that examined productivity growth and learning in new plants. In his 1993 paper [3], ‘‘Decomposing Learning by Doing in New Plants,’’ coauthored with B. H. Bahk, Professor Gort examined how productivity evolves in new plants as they age. Again, this is a paper that underscores the importance of understanding the underlying microeconomic dynamics as they relate to aggregate economic changes.

Continuing in the tradition that Professor Gort helped establish, this paper also uses U.S. Census Bureau establishment-level data to gain a better understanding of an aggregate phenomenon, in this case, investment activity. This paper examines the capital adjustment patterns for a large sample of manufacturing establishments. This is an important area to examine since accurately modeling new capital investment at the micro and macro levels has proved elusive. In standard neoclassical investment models, assumptions, such as convex adjustment costs and reversibility, dictate that firms continuously and smoothly adjust their capital stock over time. While theoretically tractable, these models generally fail to adequately explain investment fluctuations [1, 8]. The disappointing empirical performance of these investment models has caused economists to reexamine the potentially unrealistic assumptions of convex adjustment costs and reversibility. Rothschild [28] argued early on that adjustment costs faced by plants and firms possess nonconvexities for a variety of reasons.[1] Another

[1] The sources of speculated nonconvexities in the cost of capital adjustment include increasing returns, the cost of the equipment, costs associated with disruption, and installation costs.

assumption in the standard models that is unrealistic is reversibility, an area that has received a great deal of attention in recent years.[2] Models which assume nonconvex adjustment costs and irreversibility possess solutions where firms occasionally adjust their capital in discrete bursts when the capital stock falls (rises) below (above) a trigger level, solutions which differ markedly from those of standard neoclassical models.[3]

While a growing number of studies suggest that capital adjustments may occur in lumpy episodes, the theoretical literature is well ahead of its empirical counterpart.[4] This is largely due to the scarcity of data sets that follow the investment process for a large number of establishments. This situation is changing as access to microeconomic data, in particular plant-level data, increases. For instance, Caballero, Engel, and Haltiwanger [10], Power [27], and Cooper, Haltiwanger, and Power [14] investigate the lumpiness of plant-level investment and its relationship to aggregate investment fluctuations using the plant-level data on investment from U.S. Census Bureau micro data files.[5] There are two main findings. First, investment by manufacturing plants is characterized by periods of intense investment activity interspersed with periods of much lower investment activity. Second, episodes of intense investment activity are responsible for a significant fraction of aggregate investment fluctuations.

This paper also examines the patterns of investment spending at the plant level and relies on the Census Bureau micro data. As compared to the Census-based research discussed above, this paper is more descriptive. The goal of this paper is to present a series of stylized facts that will serve as benchmarks for investment models. In particular, the goal of this paper is not simply to show whether investment is lumpy or not, but instead to focus on how the distributions of investment and capital adjustment vary by plant characteristics (e.g., industry, size, age, and ownership) and by level of micro-unit aggregation (plant, line-of-business, and firm). Finally, the paper relates the evidence on micro-level lumpiness to aggregate investment fluctuations.

[2] Reviews of investment models with irreversibility include Pindyck [26], Dixit [17], and Dixit and Pindyck [18].

[3] Other underlying assumptions in neoclassical models are that capital is homogeneous and capital depreciates geometrically. Feldstein and Rothschild [22] discuss the unrealistic nature of homogeneous capital and geometric decay, and how changing these assumptions can result in lumpy investment patterns.

[4] The literature which examines labor adjustments is more mature. The importance of large proportional adjustments in employment at the establishment level has been documented by Hamermesh [24], Davis and Haltiwanger [16], and Caballero, Engel, and Haltiwanger [10].

[5] In addition to the Census Bureau microeconomic data studies, there are a number of other studies that examine machine replacement at the micro level. Rust [29] examines replacement investment with bus engines, and Cooper and Haltiwanger [13] model retooling in automobile assembly plants.

We first examine the patterns of capital accumulation within plants and focus on the magnitude of capital adjustments at annual frequencies. We find:

(1) Many plants occasionally alter their capital stocks in lumpy fashions. Over half of the plants in our sample experience a 1-year capital adjustment of at least 37%. While many manufacturing plants experience episodes of intense investment activity, 80% of the plants in a given year change their net capital stock by less than 10%. These relatively small changes account for 52% of total sample investment.

Whether or not capital adjustment is ‘‘lumpy’’ depends on to what it is compared. To help quantify what ‘‘lumpy’’ means, we compare the results from the sample of plants to those generated by simulated investment models, where the simulations include possible S, s behavior by including trigger and target levels. The larger are the estimated trigger levels in the simulations, the ‘‘lumpier’’ is capital adjustment. We find:

(2) The simulation models that best fit the observed capital adjustment patterns are those that possess trigger levels substantially above and below zero. That is, the simulation results that best fit the observed data are those in which plants mainly invest when the difference between the desired and actual capital stocks is substantially different. Otherwise, plants usually invest in small amounts, amounts that could be related to replacement and maintenance investment.

Although many plants do experience a large investment episode, perhaps our most striking finding is the tremendous variance across plants in their capital adjustment patterns. We find:

(3) With respect to plant characteristics, smaller plants, plants that undergo a change in organizational structure (e.g., ownership change), and plants that switch industries have lumpier investment patterns.

Although investment is conducted at the establishment level, investment decisions are made at the firm level. Hence, while investment may be relatively volatile at the establishment level, investment may be smoothed at the firm level, which may be consistent with the large literature on the role of firm finance constraints. In fact, we find:

(4) Plant-level capital accumulation patterns are considerably lumpier than those computed at the line-of-business level, and the line-of-business level capital accumulation patterns are noticeably more discrete than those at the firm level.

Whether or not investment is lumpy also influences models of aggregate investment. Increasing attention has recently been placed on unraveling aggregate fluctuations by examining the distribution of micro changes (e.g., [7, 9, 10, 14, 16]). Bertola and Caballero [5] model firms making investment decisions in an uncertain environment and when investment is irreversible. In this model, firms do not continually invest, but invest in lumps; hence, aggregate fluctuations in investment are partially attributable to changing proportions of the population undergoing large investment episodes. To shed light on this issue, we examine how plant-level changes in capital and investment transmit to aggregate fluctuations in investment, focusing particularly on the role of investment spikes. We find:

(5) Large investment projects in a small number of plants greatly impact aggregate investment. For our sample, 25% of expenditures on new equipment and structures goes into plants that are increasing their real capital stock by more than 30%. However, these plants make up only 8% of the sample. For the population as a whole, investment is highly skewed. In 1977 and 1987 the 500 largest investment projects accounted for 35.7 and 32.1% of total manufacturing investment.

(6) Periods of large aggregate investment are due, in part, to changes in the frequency of plants undergoing large investment episodes, though not necessarily large percentage changes in capital adjustments.

The paper proceeds as follows. Section II describes the data and the patterns of capital adjustment observed in our data sets, and provides results of simulations used to benchmark the empirical patterns. This section also examines how capital adjustment patterns vary by producer characteristics and by level of aggregation. Section III discusses the correlation between large capital adjustments and fluctuations in aggregate investment. Section IV provides summary analysis.

II. PLANT-LEVEL CAPITAL ACCUMULATION PATTERNS

In this section we examine the patterns of plant-level investment and capital growth, focusing especially on those periods when plants undergo large changes in their capital stocks. The section presents some basic statistics on capital growth rates and investment, and examines how these patterns vary by plant characteristics such as industry, plant size, and unit of aggregation (e.g., plant, line of business, firm). Before proceeding with a description of the basic patterns, we briefly describe the data. A more thorough discussion of the data set can be found in Doms and Dunne [20].

The information on annual investment and capital growth is constructed from a panel data set of manufacturing plants for the period 1972–1988. The establishment-level data are drawn from the Longitudinal Research

Database (LRD), which is maintained at the U.S. Census Bureau and contains establishment-level production data from the Annual Survey of Manufacturers (ASM). The main data set contains a balanced panel of establishments from the LRD and covers the period 1972–1988. The balanced nature of the panel ensures that capital stocks for plants can be constructed using the perpetual inventory method. The resulting data set includes 13,702 manufacturing establishments. This sample is small relative to the manufacturing population, which ranges from 312,000 to 360,000 plants over the sample period. However, while the sample coverage in terms of number of establishments is relatively small, these establishment are on average quite large and account for a significant fraction of manufacturing investment, production, and employment. Table I presents some basic characteristics of this data sample. The establishments in the sample averaged over 500 employees and accounted for 55.4–61.1% of manufacturing investment, employed 39.3–44% of manufacturing workers, and produced 47.4–53.8% of manufacturing output over the 1972–1988 period. While not reported in this paper, we have also constructed a larger data set that allows for establishment births and deaths. In general, the results reported below hold qualitatively for plants that do not span the entire time period. These results are reported in Doms and Dunne [20].

In order to measure plant-level capital growth rates, a capital series must be developed for each plant. In this paper we use the perpetual

TABLE I
Sample Coverage by Year

Year	Investment coverage (%)	Labor coverage (%)	Production coverage (%)	Average employment
1973	58.6	43.5	53.0	598.7
1974	60.1	44.0	53.8	600.4
1975	58.8	44.0	52.7	551.8
1976	56.9	44.2	54.0	570.3
1977	57.1	43.5	53.1	588.2
1978	55.4	43.3	53.2	608.1
1979	59.3	43.2	53.7	622.2
1980	60.5	42.7	52.7	601.9
1981	60.5	43.2	52.6	595.8
1982	57.7	42.2	50.3	548.7
1983	61.1	41.9	51.1	534.7
1984	57.3	42.8	51.8	558.7
1985	60.8	42.9	50.6	547.6
1986	58.2	42.5	50.1	530.5
1987	56.7	40.2	48.3	520.7
1988	58.1	39.3	47.4	514.3

inventory method. The capital stock in period t for plant i, $K_{i,t}$, is defined as

$$K_{i,t} = K_{i,t-1}(1 - \delta) + I_{i,t}, \tag{1}$$

where δ represents the depreciation rate and $I_{i,t}$ is current period investment. The rate of depreciation, δ, is estimated for each three-digit industry by imbedding the depreciation parameter within a production function. The parameters of the production function are estimated simultaneously with the parameters of the investment stream (see Doms [19] for details). Utilizing the above measure for the capital stock we construct net capital growth rates analogous to the employment growth rates of Davis and Haltiwanger [16]. The growth rate of capital for plant i at time t is computed as

$$GK_{i,t} = \frac{I_{i,t} - \delta K_{i,t-1}}{0.5 \cdot (K_{i,t-1} + K_{i,t})}. \tag{2}$$

For each plant in our sample, we compute $GK_{i,t}$ for every year from 1973 to 1988.[6,7]

Figure 1 presents two distributions, the density of $GK_{i,t}$ and the density of $GK_{i,t}$ weighted by $I_{i,t}$. The figure shows that 51.9% of plants in a year increase their capital stock by less than 2.5%, while 11% of plants in a year increase their capital stock by more than 20%. However, the few plants that do undergo large changes contribute significantly to the level of aggregate investment. The weighted distribution shows that 25% of investment is in plants increasing their capital stock by more than 25%. At the other end of the distribution, 19.2% of investment is occurring in plants changing their capital stock by less than 2.5%.

Figure 1 shows that the distribution of investment is skewed, with a small number of plants accounting for a relatively large share of investment. While this is present in our subsample of data, it is also true in the establishment population as a whole. Table II gives the share of total investment in 1977 and 1987 accounted for by the top 100 investing plants, top 500 investing plants, top 1000 investing plants, etc. Also given in Table II are the analogous figures for ranked employment and output. This table

[6] The Annual Survey of Manufacturers stopped collecting the book value of capital data in 1989, as well as other investment related data, making it difficult to compute capital stocks after 1988.

[7] Unfortunately, the above expression ignores early retirements in the construction of the capital stock. The LRD does contain some data on retirements, but these data appear to contain significant errors. The constructed growth rate is therefore a relative measure of new capital accumulation net of depreciation.

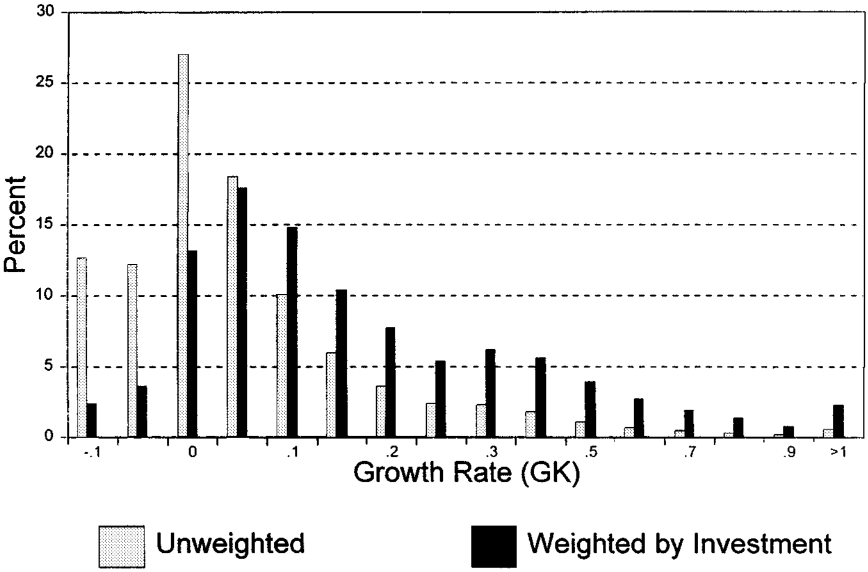

FIG. 1. Capital growth rate (*GK*) distributions: Unweighted and weighted by investment.

is based on the *entire* manufacturing establishment population. The overall message is relatively clear. A small number of plants account for a large fraction of investment. In 1977 and 1987, 18.2 and 16.2% of total manufacturing investment was accounted for by the top 100 plants, respectively. In contrast, the top 100 plants accounted for a substantially smaller fraction of ouput (9.0%) and employment (5.9%). Note that 100 plants make up only 0.028% of the entire population. The bottom line is that in a cross section a small number of investment "projects" account for a substantial fraction of aggregate investment. While this cross-sectional result is suggestive of "lumpy" investment, it does not provide any information on the within-plant investment patterns over time. It is a description of these within-plant patterns of investment and capital adjustment that we turn to next.

To examine the within-plant capital accumulation patterns, we construct two sets of ranks to describe the distributions of capital growth and investment at the plant level. The first measure constructs a ranked distribution of capital growth rates for a plant. For each plant in the balanced panel, we rank their capital growth rates from highest to lowest, so that their maximum growth rate is rank 1 and their lowest growth rate is rank 16. Throughout this paper, the rank 1 growth rate is denoted by MAXGK. Figure 2a presents the means and medians of these ranked

TABLE II
Share of Investment, Employment, Shipments, and Capital Accounted for by the Top Plants in Each Category

	1987 Census of manufactures: 358,567 plants			
	Investment	Employment	Output	Capital stock
Top 100 plants	.16204	.06344	.10077	.11888
Top 500 plants	.32154	.14057	.23031	.28882
Top 1000 plants	.41268	.18982	.30819	.38497
Top 5000 plants	.64769	.36233	.52581	.60963
Top 10000 plants	.74987	.47020	.62994	.70622
Top 25000 plants	.86863	.64043	.77045	.83002
Top 50000 plants	.93531	.77445	.86831	.90761
	1977 Census of manufactures: 350,648 plants			
	Investment	Employment	Output	Capital stock
Top 100 plants	.18172	.05932	.09005	.12883
Top 500 plants	.35657	.14584	.21638	.29359
Top 1000 plants	.44948	.20269	.29398	.39090
Top 5000 plants	.67240	.38958	.51551	.62407
Top 10000 plants	.76821	.50301	.62389	.72395
Top 25000 plants	.87931	.67753	.77172	.84548
Top 50000 plants	.94131	.80945	.87187	.91819

growth rates, so the first set of bars in Fig. 2a shows the mean and median MAXGK. The next set of bars shows the means and medians of the second largest growth rates, and so on. These bars indicate that the mean MAXGK slightly exceeds 46%, while the median is 36%. The means and medians drop off significantly after rank 1. Figure 2a illustrates that many plants experience a few periods of intense capital growth and many periods of relatively small capital adjustment: of the 16 capital growth rate ranks, 12 possess means or medians between -10 and $+10\%$.

Besides the growth rates of the capital stock, we are also concerned with episodes of investment that account for a large share of a plant's total investments. Figure 2b plots the mean proportion of total 16-year investment that occurs in each year. For instance, the leftmost bar represents that the average plant experiences a 1-year investment episode that accounts for 24.5% of its total real investment spending over the 16-year interval. The secondary growth rate accounts for 14.7%, and the third highest accounts for 10.9% of investment. This implies that, on average, half of a plant's total investment over the 1973–1988 period was performed in just three years. An important point is that while a significant portion of investment occurs in a relatively small number of episodes, plants still invest in every period.

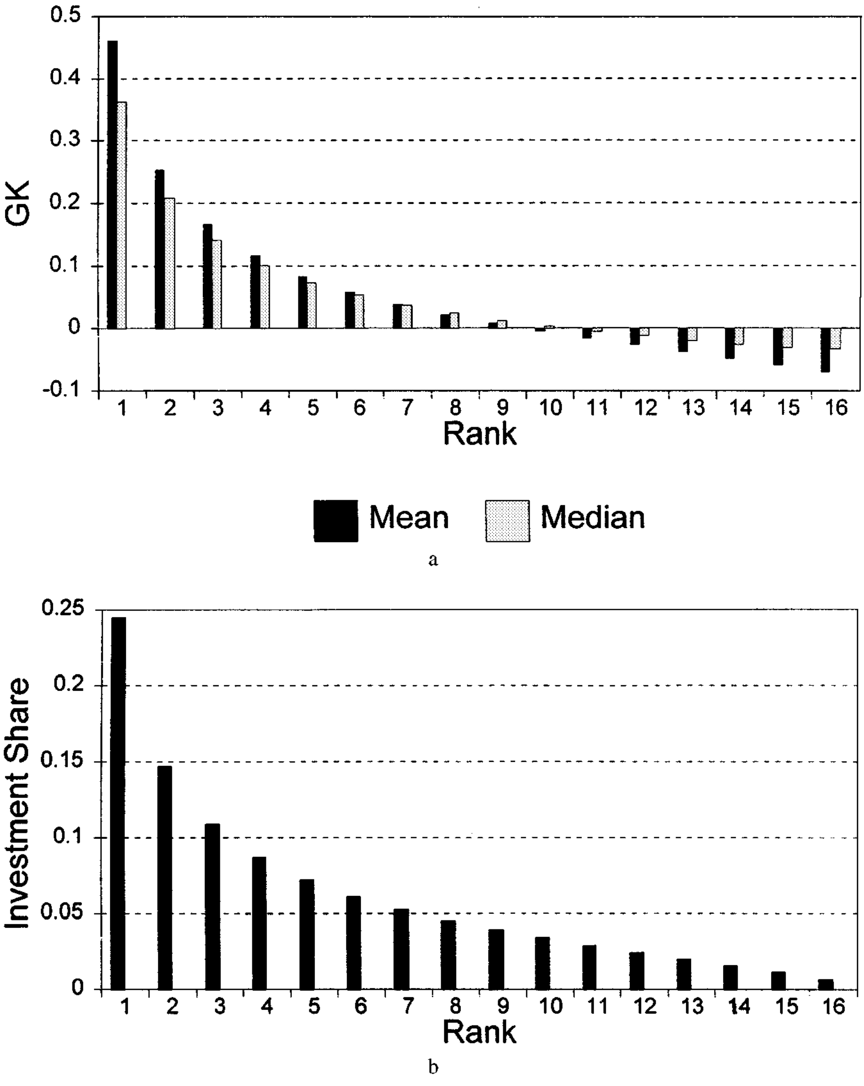

FIG. 2. Capital growth rates (GK) by rank, means, and medians. (b) Mean investment shares by capital growth rate rank.

What does Fig. 2a and b say about whether investment is "lumpy" or not? By construction, these figures slope down, and it is difficult to tell if, for instance, the data generating the figures come from something as simple as a Gaussian white noise process or whether the data are truly representative of a "lumpy" process. To benchmark our results, we compare our empirical results to simulations of simple capital investment models that include the possibility of lumpy adjustment episodes. Although the simulations do not formally test particular investment models, the simulations do provide a convenient benchmark to view our results. The following model was kindly provided by Jeffrey Campbell.

Let $k^*_{i,t}$ denote the optimum level of the logarithm of the capital stock of plant i at time t if the plant faced no frictions in adjustments, frictions that might arise from nonconvex adjustment costs or irreversibilities. Let $k_{i,t}$ be the actual capital stock. In the simulations that follow, we assume that the optimum level of capital, $k^*_{i,t}$, follows a random walk of the form

$$k^*_{i,t} = k^*_{i,t-1} + \varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim N(\mu, \sigma^2). \tag{3}$$

The disturbance $\varepsilon_{i,t}$ is i.i.d. across time and plants. Let $z_{i,t} = k_{i,t} - k^*_{i,t}$ be the difference between the frictionless optimum and the actual capital stock. The investment decision for a plant follows that of a general S, s model, where the trigger levels are denoted by U and L and target levels by u and l, such that $L \leq l \leq u \leq U$. For instance, if $U = u = l = L = 0$, then there are no frictions and plants will always invest to their optimal frictionless level of capital and capital adjustment would be normally distributed. However, if the target levels do differ from zero, then there will be periods when no investment takes place, that is, periods in which $z_{i,t-1} + \varepsilon_{i,t}$ lies within the U, L band. If $z_{i,t-1} + \varepsilon_{i,t} < L$, then the plant will invest up to l. Likewise, if the optimum level of capital falls sufficiently, $z_{i,t-1} + \varepsilon_{i,t} > U$, then the plant will disinvest to u. We modify this basic friction model by adding replacement investment since some investment always takes place in our sample of establishments. Replacement investment, $r_{i,t}$, is uniformly distributed and is independent across time and plants.

The simulations are performed with 1000 plants and are run for 300 periods. The last 16 observations for each plant are taken and the capital adjustments are ranked, just as they are ranked with the real data. The parameters of the simulations are calibrated to minimize the mean squared error between the simulated values and the real values of the ranked capital adjustments. For nearly all of the simulations presented in this paper, the values of the replacement and innovation parameters are nearly identical; for the innovation parameters, $\mu = 0.05$ and $\sigma = 0.18$. The mean value of replacement investment is 0.05 with a standard deviation of 0.005–0.02.

Figure 3 reproduces Fig. 2a with the results of two simulations. The first simulation is the best fitting simulation with frictionless adjustment, $U = u = l = L = 0$. What is perhaps most striking about the frictionless adjustment simulation is its symmetry, which stands in stark contrast to the asymmetry in the real data. Additionally, the frictionless simulation does not drop as quickly or have as many periods with low capital accumulation activity as in the real data. The second simulation presented in Fig. 3 introduces frictions, that is, the target and trigger levels are allowed to deviate from zero. The friction parameters that produce the best results are $L = -0.34$, $l = -0.05$, $u = 0.20$, and $U = 0.22$. What is most striking about this simulation is how the simulated values sharply fall after the first rank and then stay much closer to 0, as in the real data. In fact, the mean squared error between the simulations and the actual data falls from 0.129 for the frictionless model to 0.012 for the friction model.

The results in Fig. 3 are for the entire sample of establishments that span the sample period. What is also striking is how the results in Fig. 2a vary by other observable plant characteristics, such as size. Figure 4 presents mean ranks by size quartile, where plants are ranked by their mean employment over the sample period. The basic result is that smaller plants have higher maximum growth rates than the largest plants. We again perform the simulations for these four plant size categories, and the parameter that changes the most is the trigger level L, which goes from -0.37 for the smallest quartile to -0.20 for the largest quartile, a significant difference. One of many possible reasons why smaller establish-

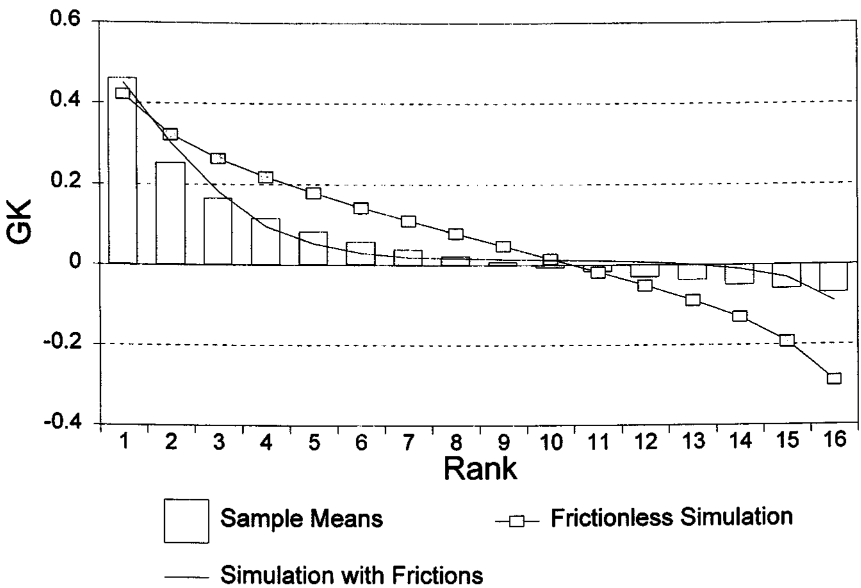

FIG. 3. Mean capital growth rates (GK) by rank, sample means, and simulated values.

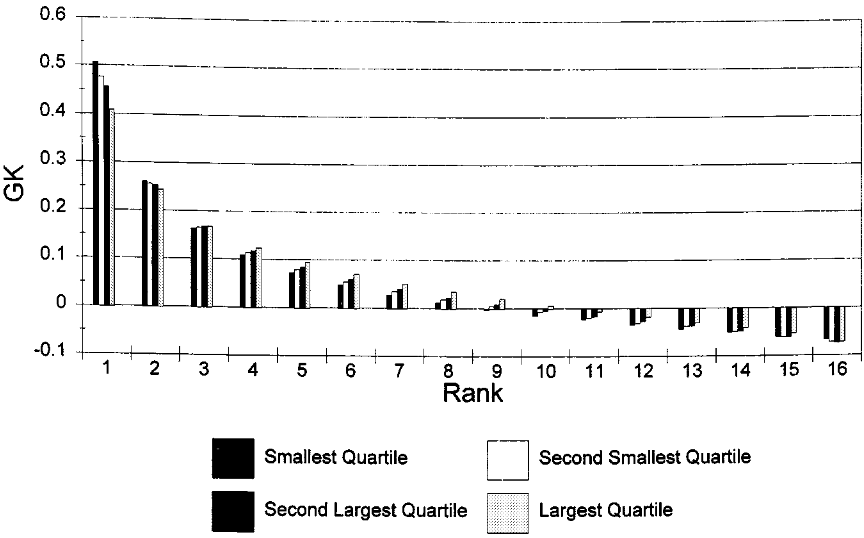

FIG. 4. Mean capital growth rates (GK) by rank and by size quartile.

ments may have higher trigger levels than larger establishments may be the indivisible nature of capital equipment; buying a single new machine at a smaller plant may represent a large share of its capital stock, so that its investment pattern may appear "lumpy." Large plants employing many machines may have smoother investment patterns because a single machine is a very small share of its capital stock. Additionally, one could view a large plant as a collection of smaller operations producing a range of products. These multiproduct operations may face less variable sales due to the fact they produce a number of different products, and hence their investment may be smoother as well.

Up to this point the unit of observation has been the plant; however, there are many arguments which suggest that the investment decisions of a plant are made at the divisional or firm level. Additionally, there are reasons why firms may smooth investment across plants. To examine how capital adjustment patterns vary by plant and firm, we construct capital adjustment ranks at the plant, the two-digit industry line-of-business level, and the firm level. The sample used to construct the plant, line-of-business, and firm statistics is a subset of the balanced panel. First, only those plants that remain with a single firm for at least 14 out of the 16 years are used. Second, only those plants that belong to firms with at least three plants are kept. Given these requirements, only 5822 plants out of the 13,702 plants in the balanced panel remain, representing 648 firms and 955

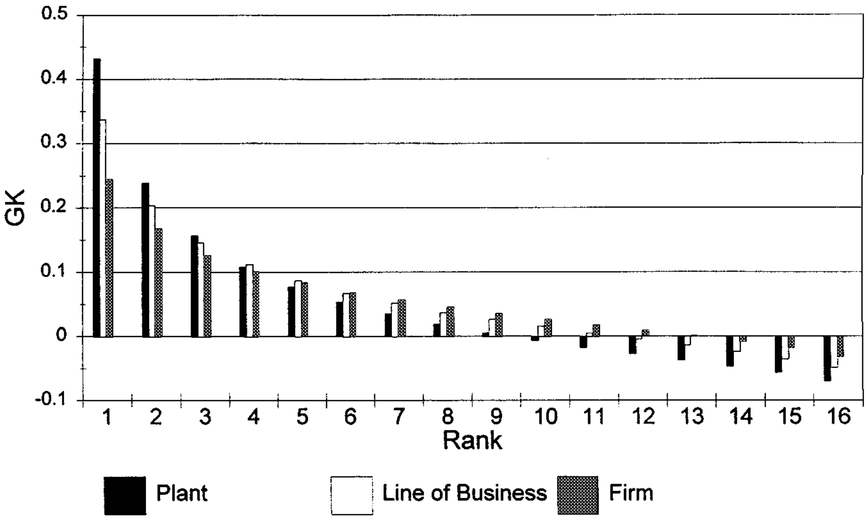

FIG. 5. Mean capital growth rates (GK) by plant, line of business, and firm.

lines of business. Note, however, that these plants make up 72.5% of the balanced panel investment.

The results of this exercise are presented in Fig. 5.[8] Basically, the higher is the level of aggregation, the smoother is the capital adjustment rank distribution. Examining the height of the largest capital adjustment episode, the mean MAXGK for plants is 0.432, and it falls to 0.336 for the line of business, and falls even further to 0.245 for firms. Again, simulations for these three distributions are run, and the estimate for the trigger level L changes from -0.35 for plants, to -0.19 for the line of business, and to -0.10 for the firm. This finding of smoother investment at the firm level may be consistent with the results reported by Cummins, Hassett, and Hubbard [15]. Note, however, that the asymmetry of the capital growth rate distribution still persists even at the firm level.

The analysis, so far, shows considerable across plant variation in capital growth rates, suggesting some plants experience relatively smooth changes in their capital stocks while other plants undergo sizable jumps in their capital stocks. A possible explanation of the differences in size is that for

[8] The basic results in Fig. 5 also hold for the investments distribution. Examining the height of the largest investment spike episode, the mean plant maximum investment share is 24%. This is quite close to that reported in Fig. 2b for the entire balanced sample. The mean maximum plant investment share drops to 17.1% at the line-of-business level and to 15.8% at the firm level. The bottom line is that firm-level investment patterns appear to be considerably smoother than plant-level investment patterns.

some industries investment is inherently lumpy because of the nature of the capital goods (which could arise due to the indivisibility of large machines), while for other industries it may be easier to adjust capital more smoothly.[9] To examine this possibility, we model MAXGK for a plant as a function of size, controlling for industry and other effects.

We estimate a regression model using all plants in our balanced panel. Our plant-level measure of capital lumpiness is the maximum single year capital growth rate (MAXGK), which our simulations show to be closely related to the magnitude of the trigger levels. Also, we have constructed other variables that characterize a plant's capital adjustment patterns, and arrive at the same qualitative results. The regressions include controls for both plant and firm size. Plant size is modeled using a set of dummy variables representing plant-size quintiles. The quintiles go from smallest to largest, with the quintile representing the largest plants omitted. The firm size variables are similarly defined. Two variables are included to capture potential changes in organizational structure and production mix that may affect capital accumulation patterns. The first variable is a dummy variable indicating whether a plant has changed ownership during the sample period. The second variable is a dummy variable which indicates whether the plant changes the two-digit industry in which the plant operates. Two age variables are included to capture differences in the age of plants that entered the panel in 1972. Finally, the regressions are all run with four-digit industry dummy variables. To conserve on space, the industry coefficients are not reported in the tables.

The second column of Table III reports the regression results. The starkest result is the strong inverse relationship between plant size and MAXGK. Smaller plants have considerably larger spikes, even after controlling for industry and other plant characteristics. Alternatively, there is no discernible pattern in the firm size coefficients. The two variables which capture change in ownership and change in industry indicate that plants which undergo ownership changes or switch industries experience somewhat larger MAXGKs. This is consistent with the view that organizational and industry changes lead to changes in plant-level operations which affect capital accumulation decisions. In terms of the simulation models, changes in ownership structure and industry may be indicators of discrete changes

[9] Doms and Dunne [20] report considerably more industry-level detail. For example, in the case of investment spikes, we find that 10% of industries (four-digit SIC) have maximum investment spikes under 0.20, 80% have maximum investment spikes between 0.20 and 0.30, and the remaining 10% have maximum investment spikes exceeding 0.30. Hence, the investment spike patterns observed in Fig. 2b are also present in a wide range of four-digit SIC industries. The same finding would be true for the capital growth rate distributions. Figure 2a is qualitatively similar to the growth rate distributions for a large number of industries.

TABLE III
Capital Growth Rate Regression: MAXGK Is the Dependent Variable

440 4-digit Industry Controls	Included
16 Year Dummies in which MAXGK occurs	Included
Plant Size Quintiles (Smallest to Largest)	
1st Quintile	.319 (.012)
2nd Quintile	.178 (.011)
3rd Quintile	.109 (.010)
4th Quintile	.056 (.010)
5th Quintile	omitted
Firm Size Quintiles (Smallest to Largest)	
1st Quintile	.007 (.010)
2nd Quintile	.037 (.009)
3rd Quintile	.013 (.009)
4th Quintile	.014 (.014)
5th Quintile	omitted
Industry Change Indicator	.031 (.011)
Ownership Structure Change Indicator	.040 (.006)
1963 Age Dummy	−.086 (.009)
1967 Age Dummy	−.049 (.011)
Mean of Dependent Variable	.461
Number of Observations	13072
R^2	.211

in the desired capital level. On the other hand, older plants have generally smaller than average capital growth rate spikes. This last result is consistent with the Jovanovic's [25] model of industry evolution that predicts that the variance of growth should decline as firms age.

The regression coefficients provide basic evidence of how capital growth varies with observable plant characteristics. However, on the whole, the plant and industry characteristics explain relatively little of variation in the standard deviation of capital growth or in the size of MAXGK. The amount of variation explained by plant and industry controls is about 20%. In general this lines up with the results reported by Davis and Haltiwanger [16], who report R^2s of similar magnitudes for employment growth regressions.

III. AGGREGATE INVESTMENT FLUCTUATIONS

This paper has so far focused on the predominance of large capital adjustments in plants and the variation across plants in their capital adjustment patterns. Increasing attention has recently been placed on unraveling aggregate fluctuations by examining the distribution of micro changes (e.g., [7, 9, 10, 12, 14, 16]). In this section, we present some basic summary statistics on the relationship between aggregate fluctuations in investment, the uniformity of changes in capital, and the frequency of large capital adjustments.

Using the balanced panel, which annually accounts for approximately 58% of aggregate investment, we compute the frequency of plants that have their MAXGK and MAXI (the maximum investment share) in a given year. Figure 6 presents these frequencies in addition to aggregate real investment over the period 1973−1988. There are several items to note. The first is that the correlation between MAXI and aggregate investment is 0.59, which is significant at the 99% level. The correlation between MAXGK and aggregate investment, however, is not statistically significant. This is due primarily to the high frequency of MAXGKs in 1973 and 1974 which is not reflected in the aggregate data.

Figure 6 conveys that aggregate fluctuations are correlated with the frequency of plants undergoing large investment episodes. An alternative way to summarize the relationship between aggregate investment and lumpy episodes is to see if investment is more skewed or concentrated in high investment periods. To address this issue, we compute a Herfindahl index for investment in each year and plot this series in Fig. 7 along with

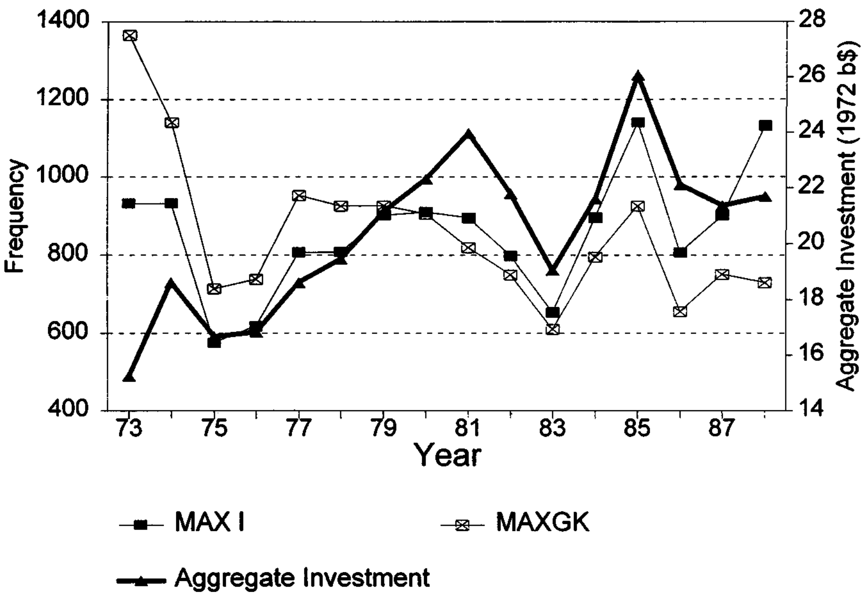

FIG. 6. Aggregate investment and frequency of plant spikes.

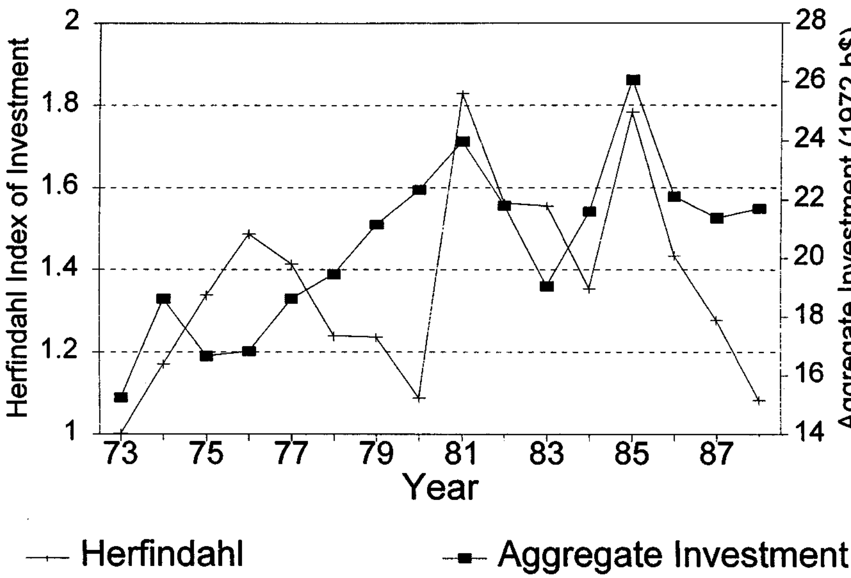

FIG. 7. Aggregate investment and the Herfindahl index of investment.

the aggregate investment series for the period 1973–1988.[10] In general, the series move together. The correlation between the two series is 0.450 and is significant at the 90% level. An interesting feature to note in Fig. 7 is that in 1980 and 1988 there are periods of relatively high aggregate investment in which there are relatively low Herfindahls. However, the two highest Herfindahls are in the 2 years with the highest aggregate investment.

IV. CONCLUSION

The objective of this paper is to present a series of stylized facts concerning the capital accumulation patterns for a large set of manufacturing plants. Although this paper is primarily descriptive in its examination of plant-level investment behavior, the facts presented here are quite striking and raise a number of issues. We have shown that many manufacturing plants do indeed alter their capital stocks in lumpy fashions, and these large adjustments do account for a significant portion of a plant's

[10] The Herfindahl index for investment is constructed as $\sum (I_i/TI)^2$, where I_i is investment in plant i and TI is aggregate investment. The Herfindahl is just the sum across all plants of the squared investment shares.

total capital expenditures and aggregate investment. However, we also find tremendous heterogeneity in the capital accumulation patterns across plants, finding that the degree of lumpiness of capital adjustment varies considerably across plants. These facts certainly raise the question of whether traditional representative agent models based on convex costs of adjustment are adequate enough to examine the dynamics of investment and capital accumulation.

That said, there are many features of the capital accumulation process that have not been addressed in the paper. One key aspect we have not examined is the within-plant timing pattern of investment. In particular, we have said little about what happens to a plant before a spike and, more importantly, what happens to a plant after a spike. To shed some light on this issue, Fig. 8 presents the mean growth rates of capital over a 5-year period surrounding the maximum capital growth spike, MAXGK. One can see that both before and after a spike, plants return to a much lower level of investment spending. This confirms the view that large capital growth episodes are interspersed with periods of relatively modest capital growth at the plant level. The specifics of investment timing are addressed more fully in papers by Caballero, Engel, and Haltiwanger [10], Cooper, Haltiwanger, and Power [14], and Power [27]. Importantly, Cooper, Haltiwanger, and Power [14] show that the probability of an establishment undergoing an investment spike increases in the time since the last investment spike. This line of research lends support to the notion that plants wait until their

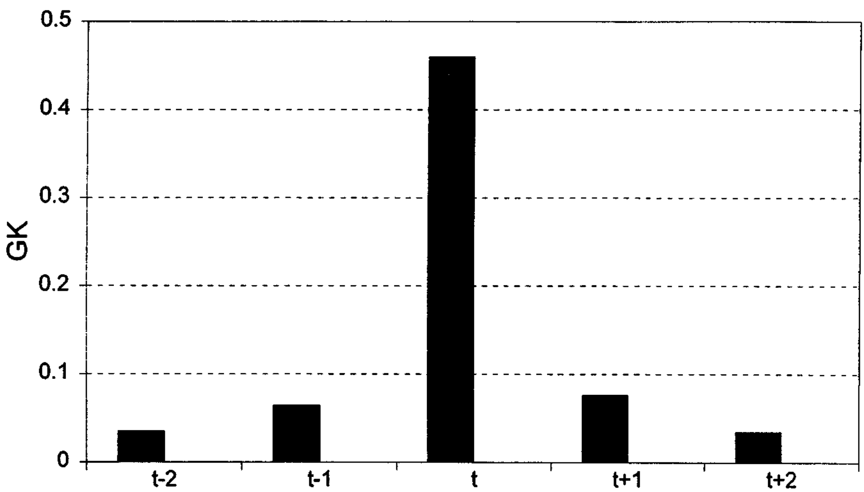

FIG. 8. Mean pre- and post-spike capital growth rates (GK).

actual capital stock deviates from the desired stock by a threshold before they invest.

In closing, this paper has described the patterns of capital accumulation using micro data on manufacturing establishments from the U.S. Census Bureau. This type of work builds on the tradition which Michael Gort helped establish almost three and half decades ago in his research using firm-level Census data from the 1940s and 1950s, and it highlights the importance of access to and development of micro data resources in understanding the underlying micro dynamics of aggregate data fluctuations.

REFERENCES

1. A. Abel and O. Blanchard, The present value of profits and cyclical movements in investment, *Econometrica* **54** (1986), 249–273.
2. A. Abel and J. Eberly, A unified model of investment under uncertainty, *Amer. Econom. Rev.* **84** (1994), 1369–1384.
3. B. H. Bahk and M. Gort, Decomposing learning by doing in new plants, *J. Polit. Econ.* **101** (1993), 561–583.
4. M. Baily, D. Campbell, and C. Hulten, Productivity dynamics in manufacturing plants, *in* ''Brookings Papers on Economic Activity: Microeconomics,'' pp. 187–249, Brookings Institute, Washington, D.C., 1992.
5. G. Bertola and R. Caballero, Kinked adjustment costs and aggregate dynamics, *in* ''NBER Macroeconomics Annual 1990'' (O. J. Blanchard and S. Fischer, Eds.), MIT Press, Cambridge, MA, 1990.
6. G. Bertola and R. Caballero, Irreversibility and aggregate investment, *Rev. Econom. Stud*. **61** (1994), 223–246.
7. O. Blanchard and P. Diamond, The Beveridge curve, *in* ''Brookings Papers on Economic Activity: Microeconomics,'' Vol. 1, pp. 1–60, Brookings Institute, Washington, D.C., 1989.
8. R. Caballero, Aggregate investment: A 90's view, *in* ''Handbook of Macroeconomics,'' (J. Taylor and M. Woodford, Eds.), North Holland, Amsterdam, 1997.
9. R. Caballero and E. Engel, Microeconomics adjustment hazards and aggregate dynamics, *Quart. J. Econom.* **108** (1993), 359–384.
10. R. Caballero, E. Engel, and J. Haltiwanger, Plant-level adjustment and aggregate investment dynamics, *in* ''Brookings Papers on Economic Activity: Macroeconomics,'' Vol. 2, pp. 1–54, Brookings Institute, Washington, D.C., 1995.
11. R. Caballero, E. Engel, and J. Haltiwanger, Aggregate employment dynamics: Building from microeconomic evidence, *Amer. Econom. Rev.* **87** (1997), 115–137.
12. J. Campbell and J. Fisher, ''Aggregate Employment Fluctuations with Microeconomic Asymmetries,'' mimeo, 1996.
13. R. Cooper and J. Haltiwanger, The aggregate implications of machine replacement: Theory and evidence, *Amer. Econom. Rev.* **84** (1993), 360–382.
14. R. Cooper, J. Haltiwanger, and L. Power, ''Machine Replacement and the Business Cycle: Lumps and Bumps,'' Working Paper 5260, National Bureau of Economic Research, 1995.
15. J. G. Cummins, K. A. Hassett, and R. G. Hubbard, A reconsideration of investment behavior using tax reforms as natural experiments, *in* ''Brookings Papers on Economic Activity: Microeconomics,'' Vol. 2, pp. 1–59, Brookings Institute, Washington, D.C., 1994.

16. S. Davis and J. Haltiwanger, Gross job creation, gross job destruction, and employment reallocation, *Quart. J. Econom.* **107** (1992), 819–864.
17. A. Dixit, Investment and hysteresis, *J. Econom. Perspectives* **6** (1992), 107–132.
18. A. Dixit and R. Pindyck, ''Investment under Uncertainty,'' Princeton University Press, Princeton, N.J., 1994.
19. M. Doms, ''Essays on Capital Equipment and Energy Technology in the Manufacturing Sector,'' Ph.D. dissertation, University of Wisconsin, Madison, 1992.
20. M. Doms and T. Dunne, ''Capital Adjustment Patterns in U.S. Manufacturing Plants,'' Working Paper 94-11, Center for Economic Studies, U.S. Census Bureau.
21. R. Eisner and R. Strotz, Determinants of business investment, *in* ''Impacts of Monetary Policy,'' pp. 59–236, Research Studies Prepared for the Commission on Money and Credit, Prentice-Hall, Englewood Cliffs, N.J., 1963.
22. M. Feldstein and M. Rothschild, Towards an economic theory of replacement investment, *Econometrica* **42** (1974), 393–423.
23. M. Gort, Analysis of stability and change in market shares, *J. Polit. Econ.* **71** (1963), 51–63.
24. D. Hamermesh, ''Labor Demand,'' Princeton University Press, Princeton, N.J., 1993.
25. B. Jovanovic, Selection and the evolution of industry, *Econometrica* **50** (1982), 649–670.
26. R. Pindyck, Irreversibility, uncertainty, and investment, *J. Econom. Lit.* **26** (1991), 1110–1148.
27. L. Power, ''Causes and Consequences of Investment Spikes in U.S. Manufacturing Plants,'' Ph.D. thesis, University of Maryland, 1994.
28. M. Rothschild, On the cost of adjustment, *Quart. J. Econom.* **85** (1971), 605–622.
29. J. Rust, Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher, *Econometrica* **55** (1987), 999–1034.