

RICARDO J. CABALLERO

Massachusetts Institute of Technology

EDUARDO M. R. A. ENGEL

Universidad de Chile

JOHN C. HALTIWANGER

University of Maryland and the

Center for Economic Studies, Bureau of the Census

Plant-Level Adjustment and Aggregate Investment Dynamics

THE EMPIRICAL INVESTMENT literature is full of disappointments. From time to time waves of new ideas challenge the aggregate investment equation, but these challenges are rarely successful, and progress is, at best, slow. There are serious theoretical obstacles, stemming mostly from the richness of the cross-sectional and time-series scenarios faced by actual investors, from the complexity of the investment technologies available to them, and from the myriad incentive problems that these economic agents face. There are at least as complex, and perhaps insurmountable, data problems. Both right- and left-hand side variables are seldom measured properly.¹

In spite of this pessimistic picture, this paper makes a new attempt. By building from the microeconomic actions of individual plants, it

We thank Summer Institute 1995 (EFCCCL) participants Kevin Hassett and Austan Goolsbee for providing us with data, and Lucia Foster and Catherine Buffington for excellent research assistance. We are grateful to Robert Hall and Michael Woodford, our discussants on the Brookings Panel, who went far beyond their duties and provided many helpful expositional and substantive suggestions. Ricardo Caballero acknowledges financial support from the National Science Foundation and the Alfred P. Sloan Foundation, and thanks the hospitality of the Inter-American Development Bank, where part of this research was conducted. Eduardo Engel acknowledges financial support from FONDECYT (Chile) grant 195-520. John Haltiwanger acknowledges financial support from the Census Fellow Program and the National Science Foundation.

1. See Chirinko (1993) for a comprehensive survey of the empirical investment literature.

aims to highlight some of the landmarks on the path leading from microeconomic investment decisions to aggregate investment dynamics. The story that emerges from this experiment does not seem to challenge common sense as much as do most contributions to this literature.

The starting point of the paper is motivated, but not limited, by the literature on aggregation of lumpy and discontinuous microeconomic investment. It is difficult to argue with the claim that constraints to adjustment faced by individual production units depart significantly from the constraints implicit in the workhorse quadratic adjustment cost model. Technological and market-induced irreversibilities, as well as indivisibilities and other forms of increasing returns (nonconvexities) in the adjustment technology, are more likely the norm than the exception.² The pattern of microeconomic investment that emerges from such constraints, contrary to the implications of the quadratic adjustment model, is highly nonlinear. Periods of more or less passive investment response to shocks are followed by feverish reactions not only to present but also to accumulated shocks.³

In reality, the fortunes and actions of individual units are very imperfectly synchronized. When combined with the nonlinear nature of the investment policies described above, this leads to a nontrivial aggregation problem; the representative agent model is not suitable for the task. Giuseppe Bertola and Ricardo Caballero model the aggregation problem of firms facing idiosyncratic and aggregate shocks when investment is irreversible. They estimate the model and find broad conformity with aggregate U.S. manufacturing data. More recently, Caballero and Eduardo Engel have generalized this to a broader set of nonlinear microeconomic policies which, at the aggregate level, nest

2. Sources of convexity do exist, however, especially at the large firm and market levels. Some of the most important ones are time-to-build, constraints to credit, and market equilibrium forces; but these are complementary to, rather than substitutes of, the nonconvexities emphasized here.

3. There have been extensive recent developments in the microeconomic literature on (S,s) models. See Harrison, Sellke, and Taylor (1983) for a technical discussion of impulse control problems. For a good survey of the economics literature on this problem (although with an emphasis on models where investment is infrequent but not lumpy) see Dixit and Pindyck (1994). See also Abel and Eberly (1994) for an elegant characterization of models that combine infrequent (but not lumpy) adjustment with convex adjustment costs.

the linear dynamics of the standard quadratic adjustment cost model. Using two-digit (aggregate) postwar U.S. manufacturing data, they find a dramatic improvement in the performance of (S,s) -type models over linear ones. The main reason behind this gain is that the aggregate counterpart of the microeconomic nonlinearities described above is a time-varying elasticity of aggregate investment with respect to shocks. The histogram of accumulated shocks across production units becomes an important state variable. If the history of shocks and microeconomic actions is such that many production units are about to enter a feverish state of investment, then aggregate investment becomes very responsive to further shocks. This additional flexibility is needed to explain the high skewness and kurtosis of aggregate investment time series.⁴

Microeconomic evidence also seems supportive of this view of aggregate investment dynamics. Mark Doms and Timothy Dunne study a large group of manufacturing plants from the Longitudinal Research Database (LRD) and document that investment spikes account for a large fraction of the investment of these plants. Furthermore, they show that the number of plants undergoing primary spikes exhibits strong positive correlation with aggregate investment cycles. Russell Cooper, John Haltiwanger, and Laura Power go one step further and find that the probability of an investment spike for a plant increases with the time that has elapsed since the previous spike, lending additional support to the view of a microeconomic environment characterized by nonconvexities in the adjustment technology.⁵

This paper integrates the organizing framework of the aggregation literature and the microeconomic data of the LRD. It organizes data from a large sample of continuously operating plants in the U.S. manufacturing sector for the period 1972–88 (that is very similar to the sample used by Doms and Dunne) in a way that is useful for understanding aggregate investment dynamics. Indeed, the final product is an aggregate investment equation that has aggregate equipment investment on the left-hand side and not only aggregate, but also microeconomic

4. See Bertola and Caballero (1994) and Caballero and Engel (1994). Bar-Ilan and Blinder (1992) make a similar point in a different context, observing that the number of consumers purchasing cars is an important explanatory variable for aggregate durable purchases.

5. See Doms and Dunne (1993) and Cooper, Haltiwanger, and Power (1995).

variables on the right-hand side. Throughout the paper we attempt to explain equipment investment. For purposes of brevity, we often omit the term “equipment” and simply refer to “investment.”

This equation, and the process of building it, leads to several conclusions:

—Measures of shocks faced by plants reveal large long-run elasticities of investment with respect to cost of capital. These elasticities vary from about -0.01 for transportation to -2.0 for textiles, with an average of around -1.0 , the neoclassical benchmark.

—Over the short run the responses are substantially smaller and variable over time; they range from 0.07 to 0.12 of the corresponding long-run elasticities.

—The variability of this aggregate elasticity comes from the nonlinearity in microeconomic policies, which indicates that plants adjust substantially more to large shortages of capital than to small ones, and that they are more likely to tolerate excesses of capital than shortages; this behavior at the plant level is not exclusive of, but is quite consistent with irreversibility plus increasing returns in the adjustment technology.

—The tax reforms of the 1980s had substantial impact on equipment investment, positive in 1981 and large and negative from 1986.

—The effect of the 1986 reform was exacerbated by a large elasticity with respect to shocks that had been brought about by the preceding expansion.

—Fluctuations in microeconomic policies also played a role in explaining aggregate investment during the sample period. This was particularly true during election years. Other things equal, investors seemed overly reluctant when Presidents Carter and Bush were elected and optimistic when President Reagan was first elected.

Organizing Framework

Working with large microeconomic data sets can be overwhelming, especially for macroeconomists. These data sets contain so much information, so many details that could be of interest, that the number of ways in which the data can be explored is almost limitless. There is a constant danger of the analysis going astray. In order to reduce this risk, this paper follows an organizing framework that proved useful in

an earlier study that analyzed the closely related problem of aggregate employment dynamics, building from the actions of, more or less, the same plants.⁶

The approach rests on the simple observation that capital is seldom at its “desired” level when adjustment costs are of any importance. The organizing framework has two basic elements, both of which are functions of an index, x , that measures the deviation between desired and actual (from hereon, natural log of) capital stock at the plant level. That is, x is the investment rate “mandated” by investment theory if adjustment costs are momentarily removed (*mandated investment*).⁷ Positive values of x reflect capital shortages, while negative values correspond to excess capital. This section takes mandated investment as given and describes the accounting part of the organizing framework, which relates microeconomic actions to aggregate investment. The theory and measurement of mandated investment are explained in the next section.

One basic element of the organizing framework captures locations: The cross-section of plants’ mandated investment immediately before the capital stock adjustments of period t is denoted by $f(x,t)$, so that the fraction of plants with mandated investment between x and $(x + \Delta x)$ is approximately equal to $f(x,t)\Delta x$. The other basic element captures actions: For every time period, we group together plants with similar mandated investment before adjustment and calculate the fraction of mandated investment that is, on average, actually undertaken by the plants within each of these groups. The resulting function is called the

6. See Caballero, Engel, and Haltiwanger (1995). For our purposes investment and labor demand have in common that adjustment costs play an important role in shaping individual plant dynamics, and that there is substantial heterogeneity in the history of individual plants.

The frequencies at which adjustment costs matter differs for labor and capital. These costs seem to be important for labor demand at a quarterly frequency, but play only a secondary role at lower frequencies, such as the annual frequencies considered in this paper. Thus the analysis of investment decisions in the present paper assumes that labor is largely flexible, and the analysis of labor demand in the other paper, to a large extent, does not consider capital stock fluctuations. Similarly, this paper analyzes equipment investment assuming that investment in structures, at the level of a continuing plant, is significantly less frequent.

7. We owe this terminology to Robert Hall. Previous versions of this paper followed the usual (S,s) terminology and sign convention, defining x as actual minus desired capital, or the “capital deviation.”

adjustment rate function (or simply, *adjustment function*) and is denoted by $A(x,t)$. It follows that firms with mandated investment x have an average investment rate during period t equal to $xA(x,t)$.⁸

Two equations are central to the approach used in this paper. The first relates individual actions to aggregate equipment investment:

$$(1) \quad I_t \equiv \int xA(x,t)f(x,t)dx.$$

The right-hand side of the above expression sums average investment rates for plants over all possible mandated investment; the sum is weighted by the cross-section of plants' mandated investment immediately before the adjustments of period t . If investment rates are independent of the stock of capital before adjustment, so that the average and aggregate investment rates coincide, then I_t is equal to the standard aggregate investment-to-capital ratio. We make the independence assumption and refer to I_t as the aggregate equipment investment rate.⁹

This accounting framework is quite general. In the simplest case, when the adjustment function, $A(x,t)$, does not depend on x , the partial adjustment model results. At the aggregate level, this has the same implications as a quadratic adjustment costs model applied to a representative agent.¹⁰ In this case, and only in this case, knowing the first moment (that is, the aggregate) of mandated investment suffices to determine the actual aggregate investment rate; higher moments (other aspects) of the cross-section of mandated investment are not needed.

Yet a departure from this basic case will imply that other characteristics of the cross-section of mandated investment than its first moment are needed to describe aggregate investment. For example, if there is a difference between the adjustment costs of increasing and decreasing the capital stock of plants, it might be expected that $A(x,t)$ would behave differently for negative and positive values of x . A simple adjustment

8. Again, we owe this terminology to Robert Hall. We used to refer to the adjustment rate function as the "hazard." The concept of the adjustment function, in isolation, is silent with respect to the way in which the average adjustment of plants at x takes place. For example, it could represent all firms adjusting by the same small fraction, or fewer plants adjusting by a larger fraction while the rest remain inactive.

9. This assumption is roughly validated by the data. A comparison of the aggregate investment rate implied by figure 3 and the mean investment rate depicted in figure 10 shows that they exhibit similar variation over time, with coincident peaks and troughs.

10. See Rotemberg (1987).

function capturing this asymmetry lets $A(x,t)$ equal a constant, A^- , for negative x and another (possibly larger) constant, A^+ , for positive x . The first panel of figure 1 shows this adjustment function and a possible cross-section of mandated investment. From equation 1 it follows that, in this case, aggregate investment rates can be calculated only if it is known what are the average investment rates mandated both for plants with negative x and for plants with positive x , and what is the fraction of plants in each group.¹¹ Knowing the first moment of mandated investment is not enough to calculate aggregate investment.

Another important model of adjustment costs is obtained if $A(x,t)$ is equal to zero when x takes values between S and s , and otherwise equal to one, where S and s denote a negative and a positive constant, respectively. The second panel of figure 1 shows this other adjustment function and a possible cross-section of mandated investment. This corresponds to a two-sided (S,s) policy with two thresholds and a common target.¹² Calculating aggregate investment requires knowing what fraction of plants are to the left and right of the inaction range, and the average mandated investment of plants in each of these groups.

More generally, it seems unlikely that plants tolerate small and large deviations in their capital stock equally well, although the degree of intolerance need not change as dramatically as it does in the (S,s) model. A simple example of an *increasing* adjustment function is $\lambda_2 x^2$.¹³ It follows from equation 1 that for this adjustment function, aggregate investment depends on the third moment of the cross-section of mandated investment. Similarly, if the adjustment function is approximated by a k -th degree polynomial, aggregate investment will be a function of the first $(k + 1)$ moments of the cross-section of mandated investment. Characterizations of aggregate investment dynamics in terms of the evolution of the higher moments of the cross-sectional distribution of mandated investment are discussed further, below.

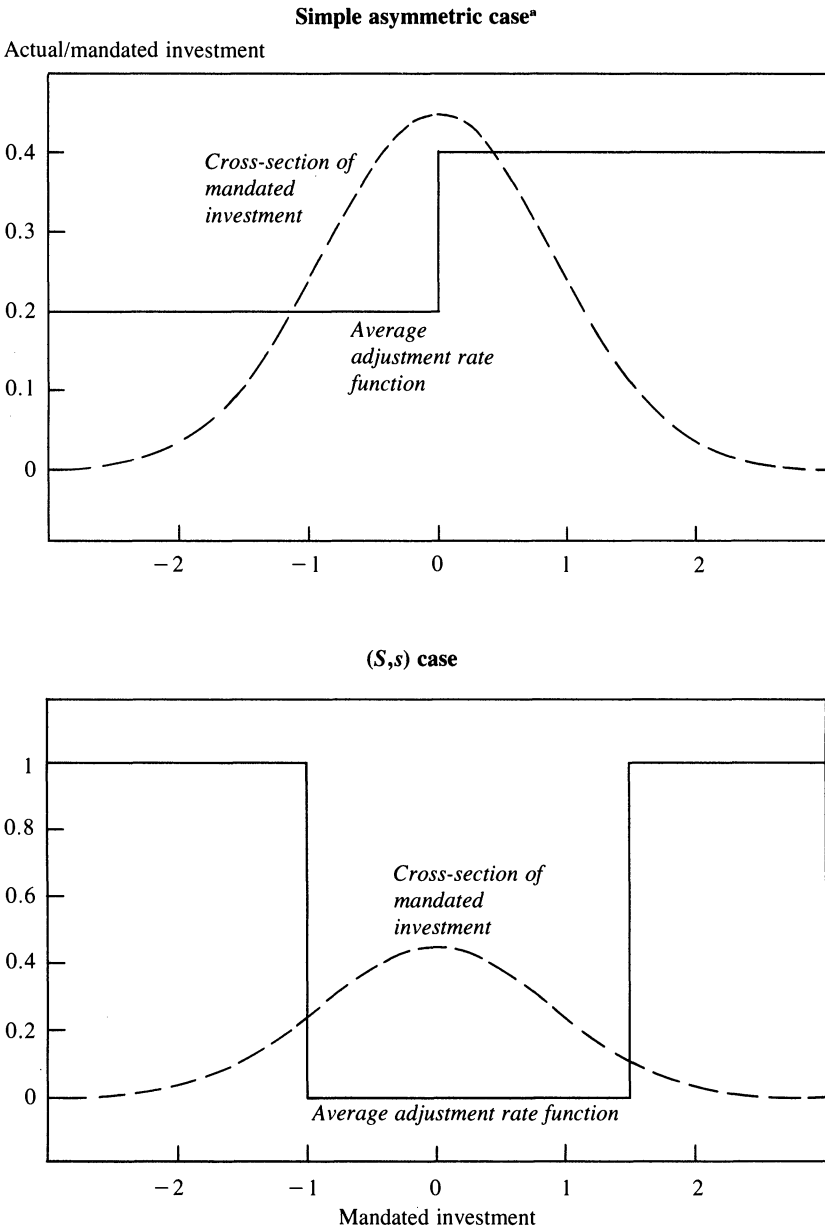
The second equation that is central to the approach taken in this paper is more involved than equation 1 because it describes the evolu-

11. Note that the composition of each group varies over time.

12. See, for example, Harrison, Sellke, and Taylor (1983).

13. More generally, an increasing adjustment rate function grows with the magnitude of plants' mandated investment rates, so that $A(x,t)$ is decreasing for negative values of x and increasing for positive values of x . In earlier work we referred to this property as the "increasing hazard."

Figure 1. Adjustment Rate Function and Cross-Sectional Distribution



Source: Authors' model as described in text.

a. The adjustment rate is the ratio of actual investment to mandated investment, where mandated investment represents the difference between the natural log of desired capital and the natural log of actual capital.

tion of the cross-sectional density of mandated investment as a function of the history of aggregate and idiosyncratic shocks, as well as of the history of individual responses to these shocks. The timing of events is described here, while the corresponding expressions are presented in appendix A.

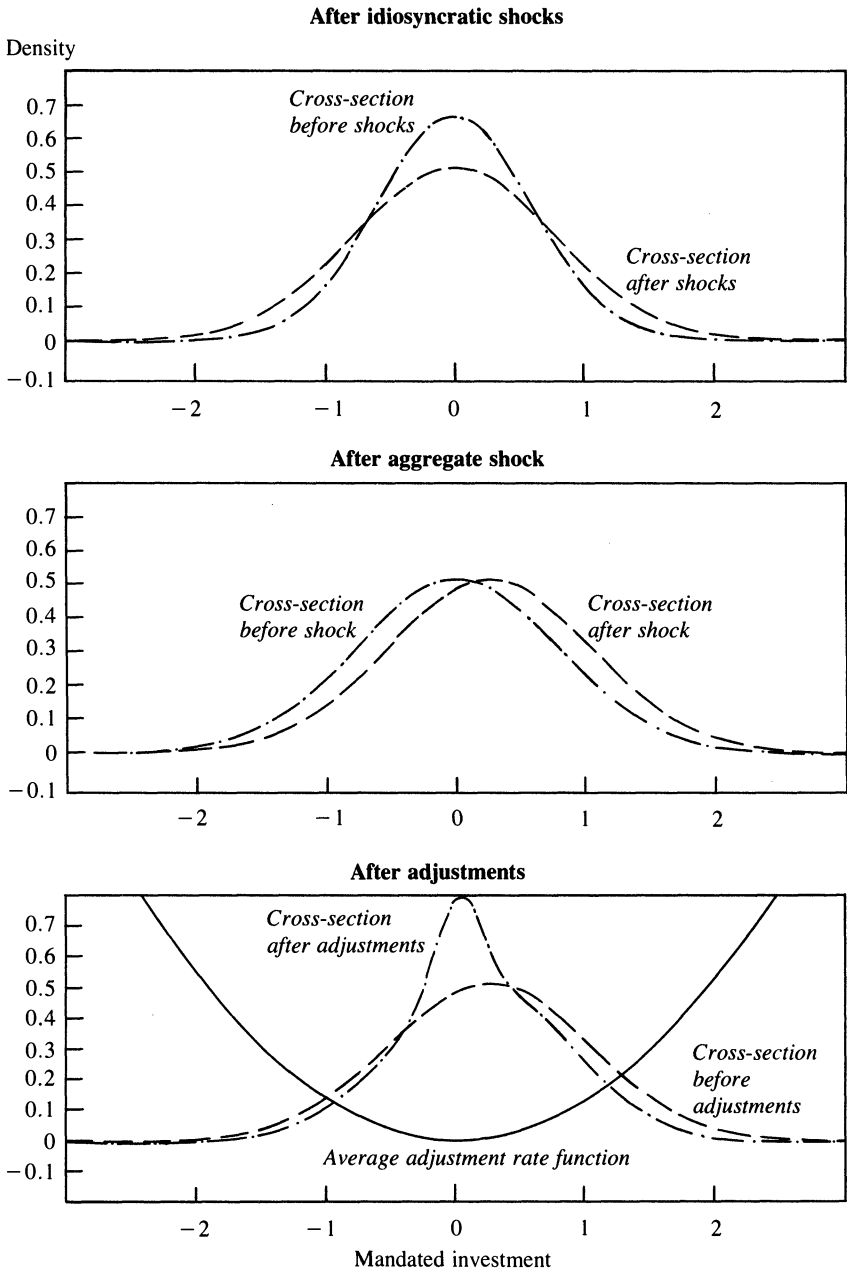
Figure 2 illustrates the evolution of the cross-section density of mandated investment during period t . Idiosyncratic shocks hit first, followed by the aggregate shock. The period concludes with adjustments by the plants. The sum of idiosyncratic and aggregate shocks reflects a combination of plant-specific and economywide depreciation, demand shocks, productivity shocks, and cost of capital shocks. By construction, the aggregate shock leads to the same change in x for all plants, while idiosyncratic shocks average to zero across plants.

The first panel of figure 2 shows how the cross-section density of mandated investment at the beginning of period t evolves after the idiosyncratic shocks have taken place. In this panel idiosyncratic shocks are assumed normal, and independent of plants' mandated investment before the shock. The second panel shows how the aggregate shock shifts the cross-section of plants' mandated investment. The cross-sectional density that results immediately after the aggregate shock (and before plants adjust) is the $f(x,t)$ of equation 1. Finally, the third panel shows the effect of capital stock adjustments on the complete cross-section of mandated investment. For illustrative purposes, in this panel it is assumed that plants either adjust (almost) fully or not at all, and that the adjustment function is quadratic. The dashed line denotes the cross-section of mandated investment before adjustments, and the dash-dotted line is the cross-section after plants complete their capital stock adjustments. The "spike" in the neighborhood of $x = 0$ reflects the plants that adjusted their capital stocks, thus leaving their mandated investment rates at approximately zero.

Data

The analysis in this paper is conducted using a balanced panel of approximately seven thousand plants in the U.S. manufacturing sector for the period 1972–88. The data are a subset of the Longitudinal Research Database, representing all large, continuously operating man-

Figure 2. Evolution of Distribution of Mandated Investment



Source: Authors' model as described in text.

ufacturing plants over the sample. The LRD is housed at the Center for Economic Studies at the Bureau of the Census and was created by longitudinally linking the establishment-level data from the Annual Survey of Manufactures (ASM). The data set includes information on shipments, materials, inventories, employment, wages and salaries, fringe benefits, energy use, cost of contract work, expenditures on equipment and structures, retirements of equipment and structures, and book value of equipment and structures. A novel aspect of the present analysis of plant-level investment dynamics is the use of data on both expenditures and retirements in order to study plant-level investment dynamics. This makes it possible to examine both positive and negative adjustments at the plant level and, in turn, to consider the aggregate implications of the plant-level adjustments.¹⁴

Investment and Capital

Our explicit treatment of both positive and negative capital stock adjustments poses interesting measurement issues beyond the typical difficulties of estimating real capital stocks and associated investment rates. A detailed description of our methodology is presented in appendix D, and a brief outline follows here. The standard procedure is to initialize the capital stock for each plant in some year and then generate a capital stock series by the perpetual inventory method, using the deflated new expenditures and depreciation rates published by the Bureau of Economic Analysis (BEA). That is, capital stocks are measured as

$$(2) \quad K_t = (1 - \delta_t)K_{t-1} + NI_t,$$

where K_t is the real end-of-period capital stock, δ_t is the depreciation rate, and NI_t is real capital expenditures. However, in generating plant-level capital stocks and investment rates that exploit the retirements data, it must be recognized that the published depreciation rates already

14. Retirements in the ASM reflect the gross value of assets sold, retired, scrapped, destroyed, and so forth. Retirements data for equipment and structures are available annually up to 1988, when the ASM terminated collection due to the development of the new Annual Capital Expenditure Survey, and this determines the end of our sample. Our use of the retirements data contrasts with the recent work of Doms and Dunne (1993) and Cooper, Haltiwanger, and Power (1995), who also study plant-level investment dynamics using the LRD, but focus only on new expenditures on equipment and structures.

incorporate retirements. Specifically, the published depreciation rates reflect two components: the loss of efficiency of an asset over its service life, while the asset is in use, and the retirement of the asset at the end of its service life.¹⁵ For example, for a one-hoss-shay pattern, depreciation entirely reflects retirements. The following discussion denotes the first component of depreciation the in-use depreciation and the second component, retirements.

The procedure used to estimate in-use depreciation rates at the two-digit level is outlined in appendix D. These estimates are used in an appropriately modified perpetual inventory method to estimate capital stocks and investment rates. Specifically, for each asset the capital accumulation equation is given by

$$(3) \quad K_t = (1 - \delta^a)K_{t-1} + NI_t - R_t,$$

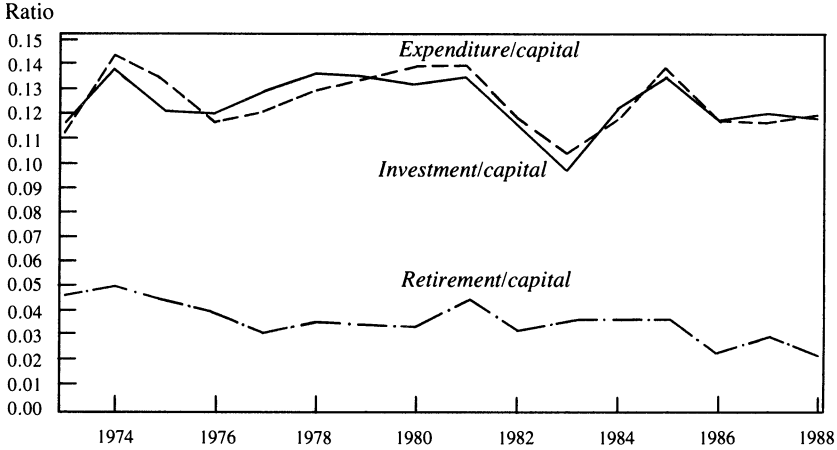
where δ^a is the in-use depreciation rate and R_t is real retirements. Estimating real retirements is complicated, since the available data are based upon the gross book value of the disposition of the assets. We estimate real retirements assuming a FIFO retirements pattern and using the appropriate deflators and adjusted depreciation rates for each vintage.

Examination of equation 3 makes clear that treating retirements explicitly is important for a number of reasons. First, it permits separation of plant-level depreciation into in-use depreciation and retirements. Second, the actual retirements pattern at the plant level will, typically, exhibit patterns quite different from the average service life distributions used by the BEA to construct depreciation rates. Put differently, using equation 2 with plant-level data on expenditures and two-digit data on depreciation rates, instead of equation 3, can yield potentially large measurement error in the evolution of the capital stock at the plant level, because the average service life distributions are applied to all plants in the same industry. Third, without a measure of retirements it would be impossible to estimate the left-hand side of the adjustment function.

Figure 3 depicts the aggregate properties of the measured rates of equipment investment. For purposes of comparison, the figure shows

15. See Hulten and Wykoff (1981) for extensive further discussion of the construction of BEA depreciation rates.

Figure 3. Aggregate Investment Rates in the Manufacturing Sector, 1973–88



Source: The investment rate is based on authors' calculations using data from the ASM. The expenditure and retirement rates are based on authors' calculations using data from the LRD.

the new equipment investment rate using published aggregate data for total manufacturing, as well as the aggregate expenditure rate (expenditures over capital stock) and aggregate retirement rate (retirements over capital stock). The expenditure rate mimics the time-series pattern of the published rate quite closely. The retirement rate averages about 4 percent and is relatively constant at the total manufacturing level.¹⁶ The following analysis focuses on a plant-level investment rate of $I_t = (NI_t - R_t)/K_{t-1}$. The rate can thus be either positive or negative.

Mandated Investment

Each element of the right-hand side of equation 1 depends on the deviations between desired and actual capital at the plant level, x_{it} , such that

$$x_{it} \equiv \tilde{k}_{it} - k_{it-1},$$

where \tilde{k}_{it} and k_{it-1} represent the natural log of desired and actual capital

16. Thus one early result is that fluctuations in aggregate net investment are dominated by fluctuations in aggregate expenditures, rather than in aggregate retirements. The in-use depreciation rate component of net investment is not reported in figure 3, but it is essentially constant; using our measurement procedures, it only varies over time as a result of changes in industry composition.

in plant i at time t (before adjustment). It should be apparent from this expression that mandated investment is not easily measured. Besides the standard problems of measuring capital (the second term in the definition of x), it requires constructing a measure of desired capital. The construction of mandated investment is summarized here and described in detail in appendix B.

First, it is assumed that desired capital is proportional to the stock of capital that the plant would hold in the absence of any frictions to adjustment, k_{it}^* , such that

$$(4) \quad \tilde{k}_{it} = k_{it}^* + d_i,$$

where d_i is a plant-specific constant to be determined later on.¹⁷ Desired capital refers to the stock of capital that the firm would hold if adjustment costs were momentarily removed; frictionless capital, on the other hand, refers to the stock of capital that the firm would hold if it never faced adjustment costs. Conceptually, the latter is a simpler construct.

Second, we let the frictionless stock of capital, k_{it}^* , be determined by the standard neoclassical expression, modified to relax the constraint that the elasticity with respect to cost of capital be equal to the output elasticity (θ_i unconstrained, below). Aside from constants, this expression can be written as

$$(5) \quad k_{it}^* - k_{it-1} = \eta_i \{y_{it} - k_{it-1}\} - \theta_i c_{it},$$

where y_{it} and c_{it} , respectively, represent the natural log of the value of output and the cost of capital in plant i at time t .¹⁸ The parameter η_i is

17. Bertola and Caballero (1994) show that this assumption is consistent with the behavior of a rational plant whose profit function is isoelastic and that faces shocks that have independent increments. In this paper we take this as a reasonable approximation.

18. The cost of capital measure we use is given by $(r + \delta_i)T_i(p_i/p_i)/(1 - \tau_i)$, where r is the real interest rate, δ_i is the depreciation rate, T_i is one minus the sum of the investment tax credit and the present-discounted value of depreciation allowances, p_i is the new capital expenditures deflator for equipment, p_i is the industry-output deflator, and τ_i is the corporate tax rate. For this measure we assume a constant real interest equal to 0.06, and use the BEA two-digit depreciation rates and the output and new capital deflators from the ASM published data compiled by Wayne Gray and Eric Bartelsman (see Bartelsman and Gray, 1995). We thank Kevin Hassett and Austan Goolsbee for providing, respectively, perfect-foresight and myopic-expectations versions of T_i on a detailed asset basis. We construct two-digit analogues using an industry-asset weighting matrix provided to us by the BEA. In empirically implementing equation 5, we use the projection of the perfect-foresight-based cost of capital measure on the myopic-expectations-based cost of capital measure and a time trend. Details of the data construction for the other terms and variables in equation 5 are discussed in appendix D.

a decreasing function of the curvature of the profit function with respect to capital, once flexible factors have been optimized over. Under a wide range of reasonable assumptions about the cost share of flexible factors and the elasticity of product demand (see appendix B for further discussion), η_i can be approximated by $1/(1 - \alpha_i)$, where α_i is the cost share of equipment capital.


Third, we estimate the cost of capital elasticities using the fact that the left-hand side of equation 5, which is equal to $(x_{it} - d_t)$, is likely to be stationary (that is, deviations are not likely to persist indefinitely), while the elements on the right-hand side of equation 5 are, individually, very persistent. Thus θ_i can be estimated from a cointegrating regression of the natural log of the capital-to-output ratio on cost of capital, and the resulting coefficient can be interpreted as the long-run elasticity of capital with respect to its cost. We do this imposing the constraints that these elasticities be equal across plants within each two-digit sector. Figure 4 reports the estimates obtained for each sector. These range from -0.01 to -2.0 , with an average at about -1.0 , the long-run elasticity in the neoclassical model.¹⁹ The fourth and last step is to estimate the plant-specific constant d_i . We let this constant be the average gap between k_{it-1} and k_{it}^* for the five points with investment closest to median investment (broadly interpreted as maintenance investment).²⁰

Thus all the necessary ingredients are available to construct estimates of desired capital and mandated investment at the plant level.

19. Cointegration regressions in models of adjustment costs are subject to small sample biases that are increasing with respect to the size of adjustment costs. These biases are reduced by adding lagged differences of the right-hand side variables to the cointegrating regression (see Caballero, 1994, for a discussion of this issue). With this purpose, we include five lags of changes in cost of capital on the right-hand side of our regressions. Our results are fairly robust to modifications in the number of lags; perhaps this is because there is enough variation in the four-digit price deflator used in the cost of capital measure that two-digit samples are, effectively, quite large.

Estimates of the long-run cost of capital elasticity are quite robust to changes in particular measures of cost of capital and levels of aggregation. For example, using the same procedure on aggregate manufacturing data alone, Bertola and Caballero (1994) and Caballero (1994) also obtain estimates of the long-run elasticity of equipment capital to cost of capital shocks of around minus one.

20. The results reported below are robust to a variety of procedures for estimating d_i , including estimating it as an industry-specific (as opposed to a plant-specific) constant. This allowed us to use procedures similar to those described in the text, or alternatively, regression-based procedures with the pooled data.

Figure 4. Long-Run Response to the Cost of Capital^a

Source: Authors' calculations using LRD data. Industry code numbers are from the Standard Industrial Classification (SIC).

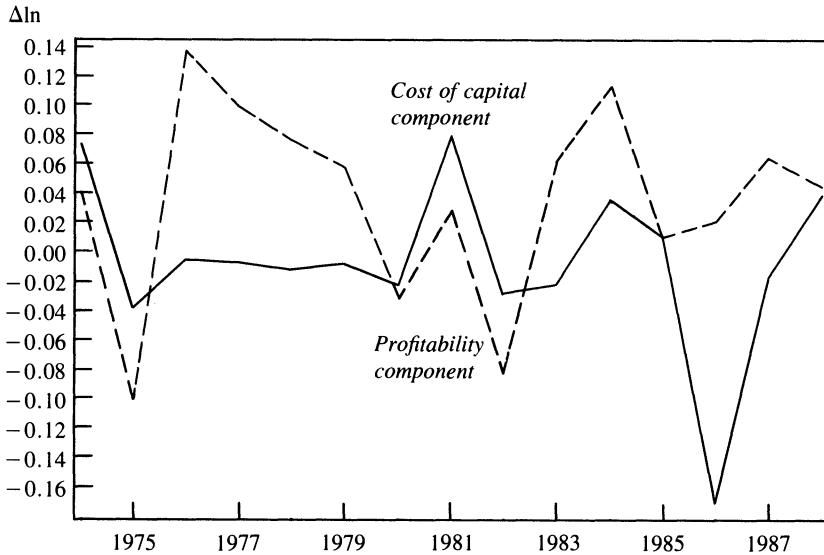
a. Plant-level data for industries 20 and 21 are pooled in estimation, constraining the elasticities for these two industries to be the same.

Shocks

For the purposes of this paper, shocks to plants correspond to changes in productivity, demand, or cost of capital that lead to changes in desired capital. From equations 4 and 5,

$$\Delta \tilde{k}_{it} = \eta_i \{ \Delta(y_{it} - \alpha k_{it-1}) - \theta_i \Delta c_{it} \}.$$

With some abuse of terminology, the first term on the right-hand side of this expression can be referred to as the profitability component of the shock and the last term, as the cost of capital component. The aggregates corresponding to these shocks are defined as the average across all plants at each point in time. Figure 5 plots these aggregate shocks. The solid line corresponds to the cost of capital component, which clearly reflects the favorable effect of the tax reform during the

Figure 5. Components of Aggregate Shocks to the Mandated Investment-to-Capital Ratio, 1974–88

Source: Authors' calculations using data from the LRD.

early 1980s, as well as the large adverse cost of capital shock brought about by the tax reform of 1986.²¹ The dotted line portrays the procyclical pattern of the path of the profitability component of the shock. Figure 5 also shows that, at the aggregate level, both components are about equally important. This contrasts with microeconomic-level data in which the profitability component is an order of magnitude more volatile than the cost of capital component.²²

21. As described above, the cost of capital measure is based on the projection of the perfect-foresight cost of capital measure on a myopic measure. We have also examined the results using the perfect-foresight measure. To a large extent, our results are robust to this alternative. The main exception is for the mid-1980s. In terms of figure 5, the less myopic version exhibits more of an increase in the cost of capital in 1984 and 1985 which, for subsequent results, implies that the 1986 tax reform has a somewhat (but not dramatically) smoother impact on aggregate investment.

22. The ratio of the variance of the aggregate profitability component to the variance of the aggregate cost of capital component is close to one. The median of the same ratio at the plant level is about ten (we only have four-digit data for cost of capital, but there is not much margin to increase the variance of cost of capital at more disaggregated levels).

Diagnostic

Before plunging into the analysis suggested by the organizing framework described above, it is useful to describe a few features of the data that are suggestive of the relevance of that framework. For this purpose we construct standardized measures of plant-level investment-to-capital ratios, shocks, and mandated investment rates. For each of the variables, we subtract from the original observations the corresponding plant-level mean, and divide this difference by the corresponding plant-level standard deviation.

The first panel of figure 6 depicts the histogram of plant-level standardized investment-to-capital ratios. It is apparent from this figure that microeconomic investment has both skewness and high kurtosis. This observation has been made before, for a similar sample; it holds true at more aggregate levels; and it has often been considered as suggestive of the presence of nonconvexities in the adjustment technologies (fat tails indicate the presence of a large fraction of large adjustments).²³

An alternative explanation is that skewness and large kurtosis are inherited directly from the shocks affecting plants. The second panel of figure 6 invalidates the latter claim by showing that the distribution of the plant-level standardized shocks does not exhibit nearly as much skewness and kurtosis as the distribution of investment-to-capital ratios.




Finally, the third panel shows that the standardized distribution of mandated investment rates exhibits negative, rather than positive, excess kurtosis. Since mandated investments are a form of storage for shocks to which plants have not yet adjusted, and since actions have positive kurtosis while shocks do not (or have less), the "leftover" component of these shocks must have negative kurtosis. The significant differences in the distributions described in the three panels, together with the near-normality of distribution of shocks, invalidate the standard quadratic adjustment cost model, in which both investment-to-capital ratios and mandated investments are linear combinations of previous shocks.²⁴

23. In regard to a similar sample, see Doms and Dunne (1993) and Cooper, Haltiwanger, and Power (1995); at a more aggregate level see, for example, Caballero and Engel (1994) for two-digit manufacturing evidence.

24. These simple observations do not deny the possibility that microeconomic units are following a probabilistic partial adjustment model which also has linear aggregate


Figure 6. Histograms of Standardized Variables^a

Fraction of plant-year observations



Source: Authors' calculations using data from the LRD.

a. For each of the variables, we standardize by subtracting a plant-level mean and dividing by the plant-level standard deviation.

Figure 7. Excess Kurtosis, 1973–88

Source: Authors' calculations using data from the LRD.

a. The difference between the sample kurtosis and 3, the kurtosis of the normal distribution.

Figure 7 provides the dynamic counterpart of figure 6 by portraying the path of the excess kurtosis coefficient of the standardized variables. It is apparent from this figure that the standardized investment-to-capital ratio has larger cross-sectional kurtosis than do shocks and deviations and that, except for a mild trend, the movements in kurtosis of investment and shocks are mostly uncorrelated (a correlation of 0.24).²⁵

Main Results

Understanding the forces behind aggregate investment dynamics, with their occasional bursts and busts, requires characterizing each of the elements of equation 1. Of particular interest are the shape of the

dynamics (see Calvo, 1983). However, when we go on to impose the structure of the organizing framework, we show that the partial adjustment model is not consistent with the data either.

25. It is also apparent from comparing figures 6 and 7 that at each year depicted in figure 7, mandated investment has less excess kurtosis than the overall excess kurtosis in figure 6. This difference is due to the effects of aggregate shocks, which are removed for figure 7 but not for figure 6 (on the distribution of standardized mandated investment).

adjustment function, its interaction with fluctuations in the cross-sectional density, and the origins of these fluctuations.

Microeconomic Adjustment

The first panel of figure 8 illustrates the average adjustment function over time as a function of mandated investment.²⁶ Bearing in mind that the right-hand side of the figure illustrates situations of capital shortage, it is apparent that, on average, plants with relatively large shortages adjust proportionally more than do plants with small shortages. This feature of an increasing adjustment rate is akin to—although not exclusive of—what is found in models where adjustment is optimally bunched due to a nonconvexity in the adjustment technology, such as a fixed cost. This seems largely realistic, and is quite consistent with the findings of skewness and fat tails for the cross-section of investment-to-capital ratios reported above. The left-hand side of the figure, which shows the region where capital is in excess, portrays a fairly flat and small adjustment rate that is reminiscent of irreversibility of investment.²⁷

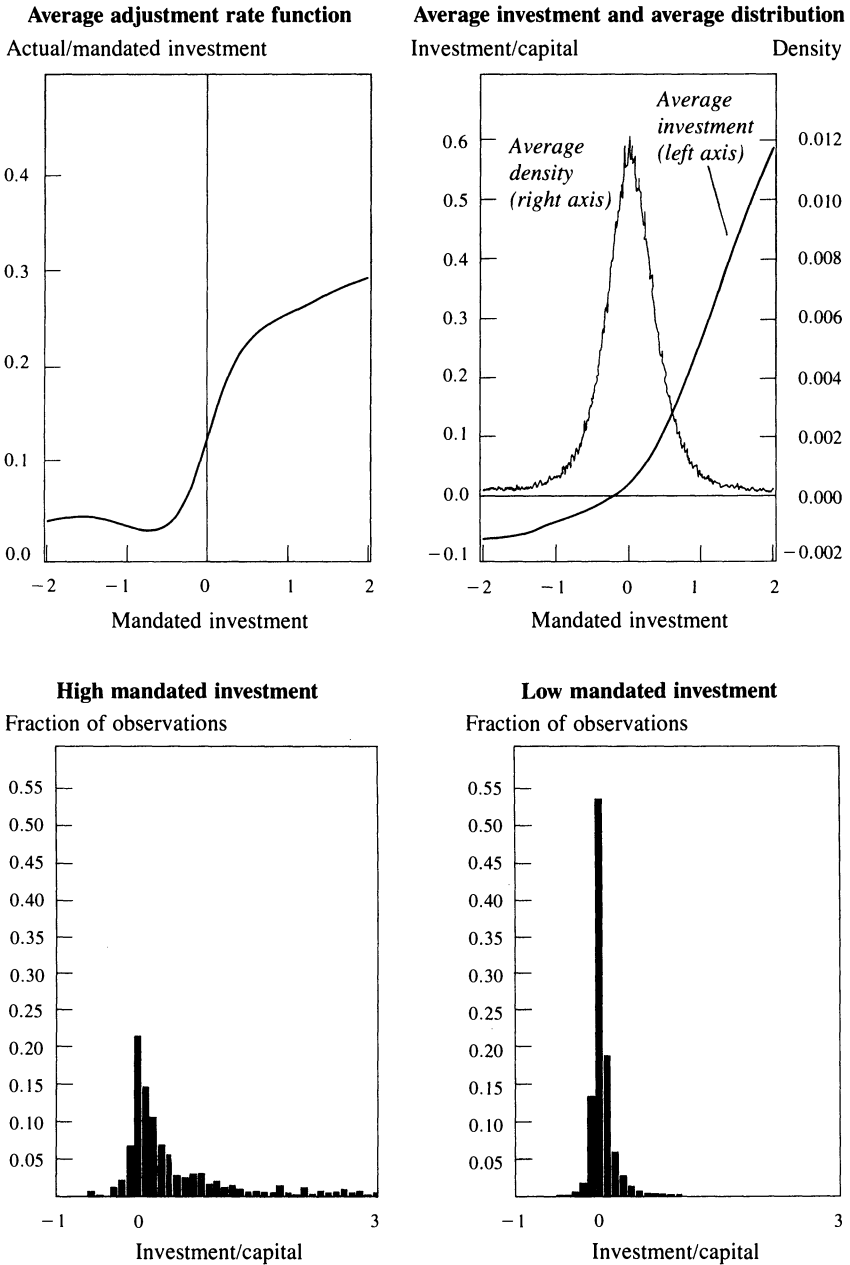
The second panel contains supportive material. The dashed line portrays the average cross-sectional density of mandated investment. It clearly establishes the fact that the nonconstant segments of the adjustment function occur in ranges where there is a significant number of observations. The solid line, on the other hand, represents average expected investment for plants with any given level of mandated investment. The nonlinear features of the adjustment function shown in the first panel are now reflected in the convex nature of expected investment.

For each level of mandated investment, the adjustment function is an average across adjustments of different sizes. The third and fourth panels of figure 8 show histograms of conditional adjustments for high (over 80 percent of desired capital) and low (between 5 and 20 percent

26. The depicted adjustment function corresponds to a cubic spline fitted over a fine grid ($\Delta x = 0.01$). In depicting the adjustment rate function in the first panel, values of x equal, or very close, to zero (that is, between -0.02 and 0.02) are excluded, since calculating the adjustment rate for a given x involves dividing the rate of investment by x .

27. It is also possible that, despite our efforts, we still have not captured most of actual capital retirements, and that this bias is increasing with the size of capital retirements.

Figure 8. Relationship between Investment and Mandated Investment



Source: Authors' calculations using data from the LRD.

of desired capital) levels of mandated investment, respectively. Although there is substantial noise in the construction of these figures, it is still possible to see clear differences between them. At high levels of mandated investment, most plants invest; for many of them, investment is well over 50 percent of capital stocks. At low levels of mandated investment, on the other hand, a large number of firms do not invest at all. Conditional histograms for negative mandated investment look very much like the fourth panel of figure 8.

Fluctuations

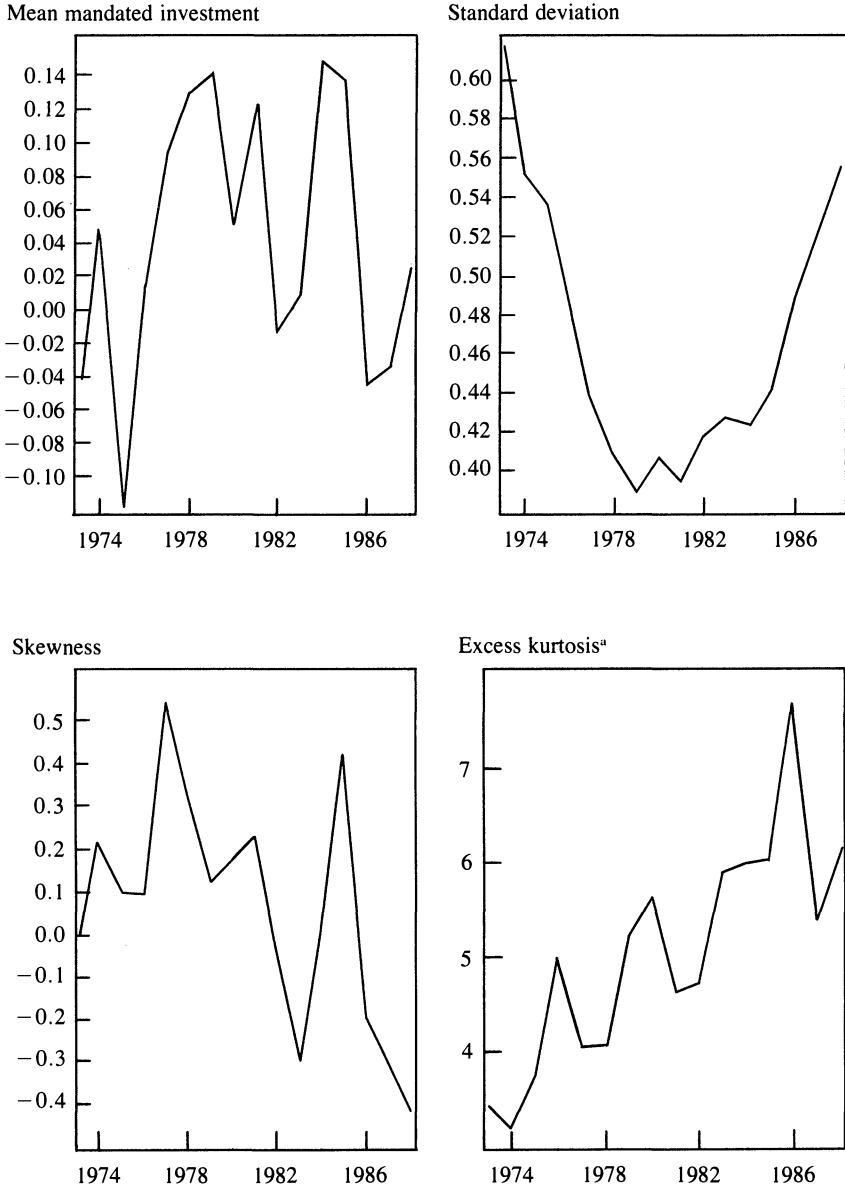
Figure 9 describes the time paths of the first four moments of the cross-sectional distribution of mandated investment. They are all quite volatile and are far from perfectly correlated. In conjunction with the nonconstant (with respect to deviations) adjustment function illustrated in the first panel of figure 8, they all shape aggregate investment dynamics.²⁸

Figure 10 shows the actual path of aggregate (average across plants) investment-to-capital ratios in the sample, as well as the counterfactual of replacing the actual adjustment function in equation 1 by the average adjustment function over time (from the first panel of figure 8) while preserving the actual cross-section distributions. The difference between the two paths is entirely due to fluctuations in the adjustment function. From the proximity of the two lines, it is apparent that fluctuations in the cross-sectional density, driven by the history of plant-level shocks and the responses of plants to these shocks (that is, adjustments), account for an important fraction of aggregate investment fluctuations.

It should be mentioned, however, that even though the difference between the two lines is not large, it is not negligible. The shifts in the adjustment function over time that lie behind this difference may largely reflect the omission of certain variables. They may, for example, represent a perception that current shocks are more or less persistent than usual, or they may reflect the closely related effect of a generalized

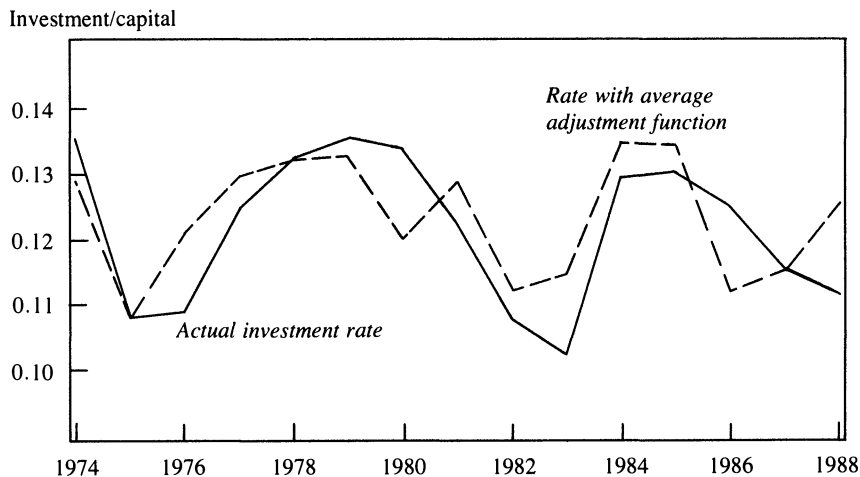
28. The careful reader may note that, contrary to figures 6 and 7, figure 9 shows excess kurtosis to be positive. The reason is that plant-level mandated investment is not standardized in figure 9. The relative comparison with figures 6 and 7 still holds, however: the nonstandardized investment-to-capital ratio exhibits excess kurtosis of more than one hundred.

Figure 9. Moments of Mandated Investment, 1973–88



Source: Authors' calculations using data from the LRD.

a. The difference between the sample kurtosis and 3, the kurtosis of the normal distribution.

Figure 10. Investment Rate with Average Adjustment Rate, 1974–88

Source: Authors' calculations using data from the LRD.

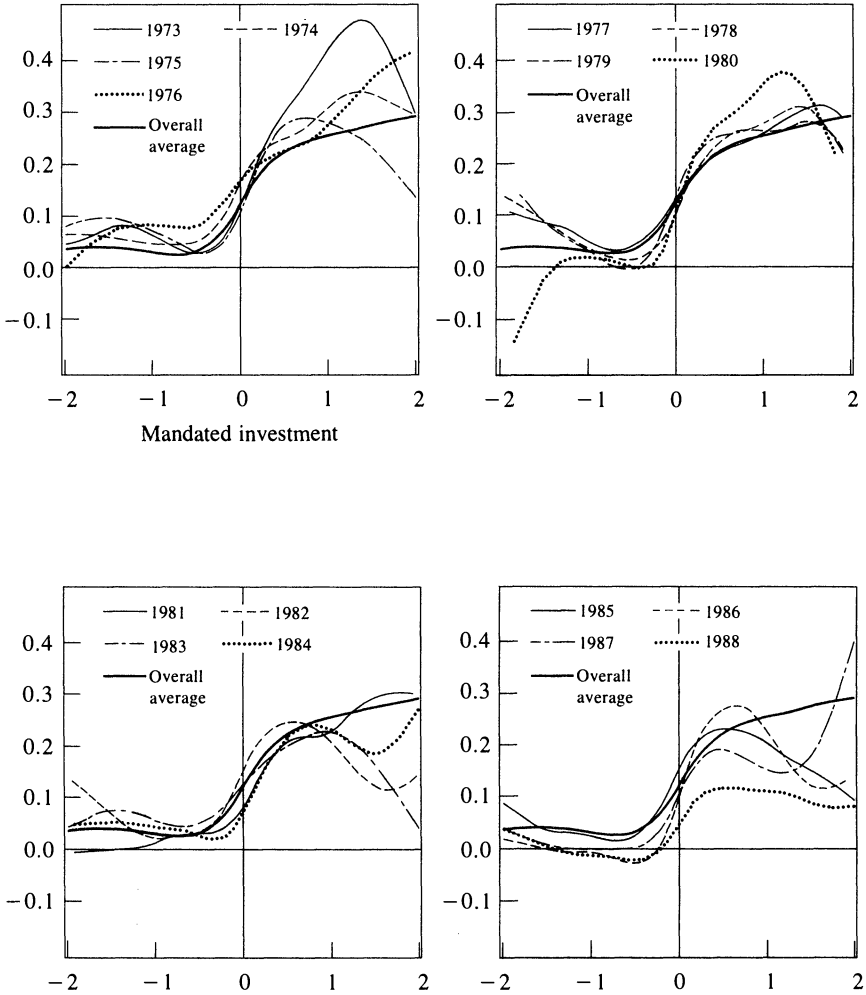
change in the confidence of investors. Whatever the structural interpretation, it is interesting to note that during election years changes in the adjustment function had a significant impact on aggregate investment fluctuations: *Given mandated investment*, both Presidents Carter and Bush's elections coincide with periods when investors were more reluctant to invest, while President Reagan's first election coincides with a period of higher investment than usual.²⁹ Other years when shifts in the adjustment function seem to have contributed significantly to aggregate investment dynamics are 1983 (negative) and 1986 (positive). The latter observation perhaps reflects the fact that we have overestimated the size of the negative surprise in the permanent component of cost of capital.

Figure 11 illustrates the actual path of the adjustment functions; it also reproduces the average adjustment function of the first panel of figure 8. *Given mandated investment*, it is apparent that 1976 was a year of abnormally high capital retirements (or, investment was below

29. There was a large positive profitability shock in 1976 (see figure 5) which investors may have interpreted, correctly, as mainly transitory. It is, however, puzzling that the bulk of the shift in the adjustment rates does not correspond to less new investment for given mandated investment, but instead is due to an increase in retirements (or, investment below maintenance levels), as figure 11 shows.

Figure 11. Adjustment Rates, by Year


Actual/mandated investment



Source: Authors' calculations using data from the LRD.

Figure 12. Investment Rate with No Current Cost of Capital Shock, 1974–88

Investment/capital



Source: Authors' calculations using data from the LRD.

maintenance level), 1980 was a year of high new investment, and 1988 was a year of low new investment. It is also apparent that the adjustment function became highly unstable following the 1986 tax reform.

Short-Run Effect of Changes in Cost of Capital

From figure 5, which shows the large size of cost of capital shocks, and the large long-run elasticities obtained above, it is apparent that cost of capital shocks have played a significant role in investment fluctuations over the long run. One elusive issue is whether these shocks have had a significant role at higher frequencies. To assess this, figure 12 compares the actual path of the investment-to-capital ratio with what it would have been in the absence of aggregate cost of capital shocks. The latter path is obtained by making the “unwarranted” assumption that neither the adjustment function nor history is affected by the experiment. This static counterfactual, therefore, merely amounts to shifting the cross-sectional density in equation 1 by the current cost of capital shocks or, equivalently, by substituting $f(x,t)$ in equation 1 by the cross-section that would have resulted if the cost of capital in period t had been equal to the actual cost of capital in the preceding period, plus the average change in this cost.

Figure 12 reveals that the 1981 tax reform played a significant role in boosting investment, but the 1986 tax reform had a much more dramatic detrimental effect on investment during the second half of the 1980s. Indeed, in the absence of the cost of capital shock of 1986, the investment rate would have been more than 15 percent higher than it actually was.³⁰

Combining Plant- and Aggregate-Level Data: A Parametric Approach

So far, this paper has used an extreme form of nonparametric approach to estimate the adjustment function that perfectly fits all aspects of the cross-sectional distribution of investment-to-capital ratios. This section follows an entirely different approach. We start by changing the metric; we are only concerned with fitting the path of aggregate (sectoral) investment-to-capital ratios, and do so with a time-invariant parametric estimate of the adjustment function. For this we approximate the adjustment function in equation 1 by a fourth-degree polynomial with a constant term that is free to vary across two-digit sectors, and higher-degree terms that are common across sectors:³¹

$$(6) \quad A_i(x,t) = a_{0,i} + \sum_{k=1}^4 a_k x^k.$$

Substituting the above expression into equation 1 (and allowing for a free, sector-specific constant) leads to

$$(7) \quad I_{i,t} = c_{0,i} + a_{0,i} M_{i,t}^1 + \sum_{k=1}^4 a_k M_{i,t}^{k+1},$$


where $M_{i,t}^{k+1}$ denotes the $(k + 1)$ -th moment of the cross-section of mandated investment in sector i during period t .

Equation 7 relates sectoral (aggregate) investment data on the left-

30. The investment-to-capital ratio would have been 14.4 percent rather than the actual 12.5 percent. Needless to say, this is just a partial equilibrium experiment.

31. The results do not change significantly for working with a second- or sixth-degree polynomial; nor for eliminating the sectoral variation in the adjustment function; nor for reweighting the sectors by the inverse of the standard deviations of their residuals.

Figure 13. Time-Invariant Adjustment Rate Function Based on Moments



Source: Authors' calculations using data from the LRD.

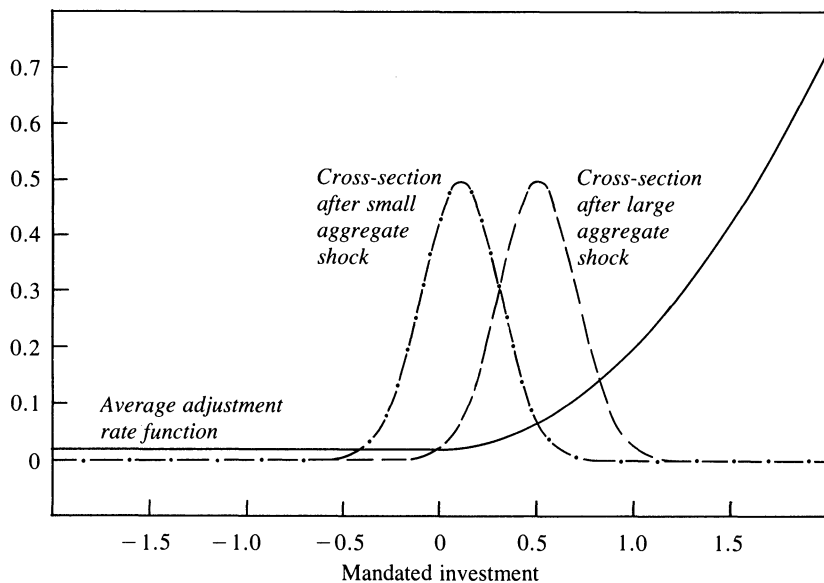
hand side to the first five moments of sectoral mandated investment rates as regressors. We estimate this set of equations using ordinary least squares (OLS). Figure 13 portrays the average (across sectors) adjustment function obtained. Interestingly, it is qualitatively similar to the average adjustment function portrayed in the first panel of figure 8, which was estimated by an entirely different procedure. As before, the adjustment function is clearly increasing for capital shortages and close to zero for retirement decisions.

The Role of Nonlinear Adjustment Functions

Figure 14 presents an example with two hypothetical cross-sections of mandated investment rates immediately before adjustments take place. The cross-section concentrated toward the right reflects a history of larger recent shocks (expansion) than the cross-section concentrated to the left.

Figure 14. Interaction of Nonlinear Adjustment Function with Aggregate Shocks

Actual/mandated investment



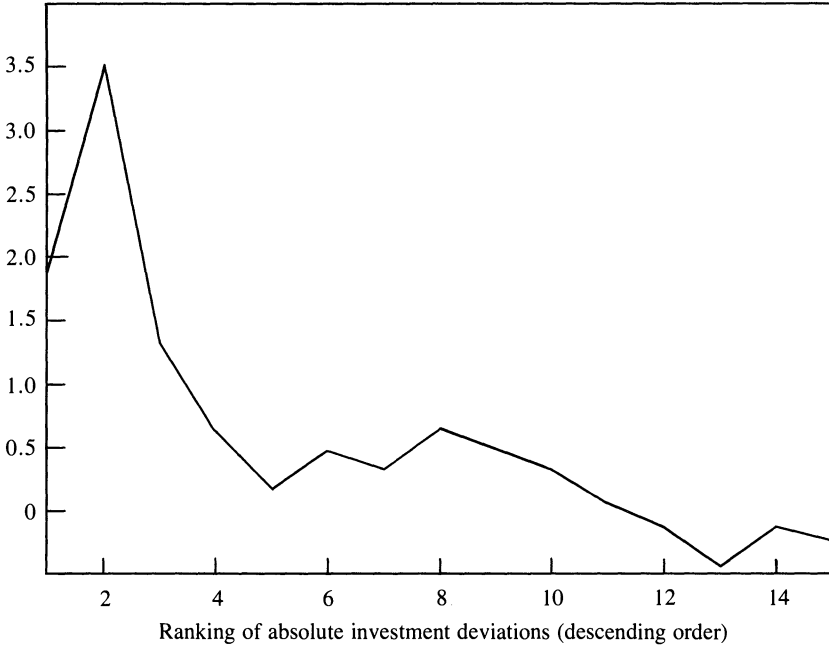
Source: Authors' calculations using data from the LRD.

For expositional purposes, figure 14 assumes that the difference in the histories underlying the two cross-sections is due only to the last aggregate shock; for this reason the cross-sections are shifted versions of each other. The purpose of this figure is to illustrate one instance in which the nonlinear nature of the adjustment function matters. Aggregate investment increases with the size of a (positive) aggregate shock for two reasons. First, as in linear models, a larger shock leads to larger average adjustments by individual plants. Second, after a large shock there is a higher concentration of plants in the region where the adjustment function increases more steeply; thus not only do plants, on average, adjust more, but the number of plants adjusting more is also larger. The latter effect, which is not present in linear models, explains why increasing adjustment functions matter most after large shocks.

To examine how the nonlinear adjustment function behaves when *accumulated* shocks are large, we sort the residuals of each sector by the size of their absolute deviation from median investment for both the semiparametric estimate of the adjustment function (increasing adjust-

Figure 15. Normalized Differences of the Sums of Squared Residuals

Difference of SSR



Source: Authors' calculations using data from the LRD.

ment) and the estimate of the constant adjustment function (partial adjustment) obtained by setting a_k equal to zero in equation 6, for values of k from one through four, and reestimating the remaining parameters by OLS. Figure 15 plots, from left to right, the sum across sectors of the differences between the squared residuals corresponding to the largest investment deviations for the partial adjustment and increasing adjustment functions, the sum of the differences between squared residuals corresponding to the second largest investment deviation, and so on. Each one is normalized by the standard deviation of the difference in squared shocks. It is apparent that the difference between the sum of squared residuals (SSR) is large for the largest investment episodes and decreases with the size of investment deviations, showing that the non-constant adjustment function model matters most during periods of large accumulated shocks.

In contrast to linear models with fixed parameters, adjustment func-

tion models have the flexibility to capture, in a simple and parsimonious manner, changes in the aggregate responsiveness to shocks over the business cycle. This flexibility can be illustrated by calculating the marginal response (*MR*) of the aggregate investment-to-capital ratio to aggregate shocks:³²

$$(8) \quad MR = \int A(x,t)[1 + e(x,t)]f(x,t)dx,$$

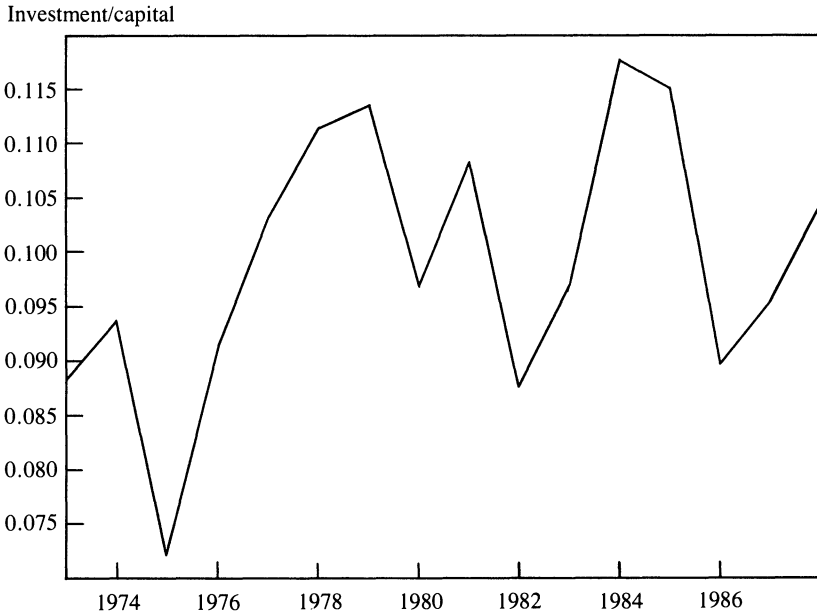
where $e(x,t)$ denotes the elasticity of the adjustment rate at time t with respect to mandated investment x . The first term on the right-hand side is also present in constant adjustment function (linear) models, although without x as an argument: it is equal to the fraction of mandated investment that actually takes place, on average. The second term is a weighted average of mandated investment elasticities, with weights proportional to the contribution to aggregate investment of plants with different mandated investment rates. Figure 16 portrays the path of the average marginal response of investment across sectors after suppressing the exogenous time variation in the adjustment function. It is apparent that these changes are important, even at the most aggregate level.

Figure 17 depicts the relative contribution of the time-varying marginal response. The conceptual experiment underlying this figure is the decomposition of the difference between actual investment and investment in the absence of the current shock into a linear component and a nonlinear component that is equal to zero with a constant adjustment function. The figure illustrates the ratio of the nonlinear to the linear component. A positive value reflects an amplification effect, while a negative value reflects an off-setting effect. The impact of the time-varying marginal response appears to be especially large during 1986; the decline in investment is 20 percent greater than it would have been in the absence of the nonlinear component.³³

32. Recalling that v is linear in the log of the aggregate shock, this amounts to calculating dI/dv evaluated at $v = v_t$. The derivation is provided in appendix C.

33. In comparing figures 16 and 17, the relevant marginal response of investment to shocks for figure 17 lies somewhere between those depicted for 1985 and 1986 in figure 16. That is, the marginal response depicted in figure 16 for 1986 (1985) represents the marginal response *after* the shock in 1986 (1985). It is worth noting that this finding of a large effect from the 1986 tax reform is related to, but quite different than, the findings of Cummins, Hassett, and Hubbard (1994). They find evidence of a large effect on the cross-sectional pattern of investment from the cross-sectional changes in taxes imme-

Figure 16. Marginal Response of Aggregate Investment to Aggregate Shock, 1973–88^a



Source: Authors' calculations using data from the LRD.

a. The y-axis gives the effect of doubling the cost of capital on the investment-to-capital ratio.


Conclusion

This paper started on a pessimistic note and concludes on a more positive one. Despite the simplicity of the “almost accounting” framework used, it has yielded a view on U.S. equipment investment that is not at odds with common sense:

—There are large long-run elasticities of investment with respect to cost of capital. These vary from about -0.01 for transportation to -2.0 for textiles, with an average of around -1.0 , the neoclassical benchmark.

—Over the short run, the responses are substantially smaller and

diately after tax reforms, including the 1986 reform. In contrast, our result shows that there is a large impact on aggregate investment due to the 1986 reform, induced, in part, by the time varying elasticity.

Figure 17. Relative Contribution of Time-Varying Marginal Response, 1974–88

Source: Authors' calculations using data from the LRD.

variable over time; they range from 0.07 to 0.12 of the corresponding long-run elasticities.³⁴

—The variability of these aggregate elasticities comes from the non-linearity in microeconomic policies. Plants invest a larger fraction of their mandated investment rates when these rates are large and positive than when they are small or negative.

—The tax reforms of the 1980s had substantial impact on equipment investment: positive in 1981 and large and negative from 1986.

—The effect of the 1986 reform was exacerbated by a large elasticity with respect to shocks that had been brought about by the preceding expansion.

—Fluctuations in microeconomic policies also played a role in determining aggregate investment during the sample period. This was particularly true during election years. Other things equal, investors

34. *After* suppressing exogenous fluctuations in the adjustment function.

were overly reluctant when Presidents Carter and Bush were elected, and overly optimistic when President Reagan was first elected.

There is plenty of space for improvement on the approach used in this paper. In technical terms, this would primarily involve enriching the set of variables that characterize the state of plants. Specifically, we pay too little consideration to agency problems and additional dynamic factors accruing to time-to-build and other sources of inertia after investment decisions have been made. We leave no margin for shocks of different persistence to affect the responses of plants; and we treated the data as if they were measured without noise. These are important factors and must be largely responsible for the fluctuations in the microeconomic policies observed.

APPENDIX A

Cross-Sectional Dynamics

THIS APPENDIX provides explicit expressions for the evolution of the cross-section of mandated investment rates. The evolution of the density of deviations during period t is affected by three inputs. First, the initial density (the final density of the previous period), $f_1(x, t - 1)$, is convolved with the density of idiosyncratic shocks, giving rise to $f_2(x, t - 1)$. We let the density of idiosyncratic shocks depend on initial deviations and denote it by $g(v, t|x)$. Thus

$$f_2(x, t) = \int f_1(x - v, t - 1)g(v, t|x - v) dv.$$

Second, there is an aggregate shock that shifts all units by v , in state space, yielding $f(x, t)$. Third, denoting by X_t and $X_{1,t}$ the random variables corresponding to $f(x, t)$ and $f_1(x, t)$, it follows that $X_{1,t} = X_t(1 - J_t)$, where J_t denotes the fraction of its mandated investment rate by which a plant adjusts. We denote the density of the J_t by $a(j, t|x)$, which satisfies the constraint $A(x, t) = \int ja(j, t|x)dj$, and summarize this last step by

$$f_1(x, t) = \int \frac{1}{u} a \left(1 - u, t \left| \frac{x}{u} \right. \right) f \left(\frac{x}{u}, t \right) du.$$

APPENDIX B

Frictionless Capital

LET Y REPRESENT the *value* of output (“output”) of an individual plant. Imperfect competition and fixed factors other than equipment yield a decreasing returns output function

$$Y = AK^\gamma F^\phi, \quad \gamma + \phi < 1,$$

where A , K , and F are indexes of productivity and contribution of fixed factors, equipment capital, and flexible factors, respectively.

Optimizing over flexible factors yields a profit function

$$\Pi(K) \equiv \max_F Y - w_F F,$$

where w_F is the price of flexible factors. Frictionless capital is the argument of the maximization of this profit function,

$$K^* = \operatorname{argmax}_K \Pi(K) - cK.$$

Taking the logarithm of this expression, with some manipulation, yields

$$k^* - k = \eta \{y - k - c\},$$

where $\eta \equiv (1 - \phi)/(1 - \gamma - \phi)$. This equation corresponds to equation 5, but in that case the constraint that cost of capital elasticity be the same as the output elasticity was relaxed. This assumption was viewed as responsible for the (to some, artificial) significance of cost of capital in early estimations of neoclassical equations.³⁵

Estimating η is not an easy task, for it requires assumptions, or

35. Michael Woodford has pointed out to us that allowing this extra degree of freedom may be inconsistent with constant values of η . It is therefore comforting that our estimates of the cost of capital elasticity are not very far from one, at least on average.

guesses, about many intermediate steps. In principle, it requires estimates of the elasticity of demand faced by the plant and of the fraction of factors that would be easily adjusted if capital were to change (flexible factors). Instead, we rewrite η as

$$\eta = \frac{1 - \phi}{1 - \gamma - \phi} = \frac{1}{1 - \alpha} \left(\frac{1 - \alpha}{1 - \gamma} \right) \left(1 + \frac{\gamma\phi}{1 - \gamma - \phi} \right),$$

where α is the cost share of equipment capital. Note that the approximation

$$(B1) \quad \eta \approx \frac{1}{1 - \alpha}$$

is fairly good for a wide range of ‘‘reasonable’’ assumptions, for example, if $\alpha = 0.06$, the markup is 30 percent, and the cost share of flexible factors is 0.7, then the ratio of the exact value of η and its approximation is only 1.07. Since equipment cost shares are easy to compute, we use the approximation in equation B1.

APPENDIX C

Marginal Response of Investment to Shocks

THIS APPENDIX derives equation 8. Denoting by $f^*(x,t)$ the cross-section density *before* the aggregate shock in period t , we have that equation 1 implies that aggregate investment, as a function of the aggregate shock, v , is equal to

$$I_t = \int (x + v)A(x + v,t)f^*(x,t)dx.$$

Differentiating with respect to v leads to

$$(C1) \quad I'_t(v) = \int [A(x + v,t) + (x + v)A'(x + v,t)] f^*(x,t)dx.$$

Recalling that $f^*(x,t) = f(x + v_t,t)$, equation C1 leads to

$$(C2) \quad I'_t(v_t) = \int [A(x,t) + xA'(x,t)] f(x,t) dx.$$

From this, simple algebra leads to equation 8. Equation C2 shows that if $A(x,t)$ is approximated by a k -th degree polynomial in x , then the marginal response of investment to aggregate shocks can be calculated from by using the first $(k + 1)$ moments of the cross-section of mandated investment rates. Figures 16 and 17 are constructed using this result.

APPENDIX D

Data Appendix

THE CAPITAL STOCK series for equipment and for structures for each plant are constructed as follows.³⁶ The real capital stock is initialized in 1972 by dividing the book value in 1972 by a two-digit-based capital deflator. The 1972 capital deflator is the ratio of the current dollar book value in 1972 to the constant dollar value in 1972 for the two-digit industry, in 1987 dollars. Starting with the initial capital stock, subsequent values of capital stock are created by using equation 3. The expenditures data on capital are deflated with four-digit, new capital deflators from the Gray-Bartelsman data.

The in-use depreciation rate and the real retirements measure are constructed as follows. Equations 2 and 3 together yield

$$(D1) \quad \delta^a = \delta_t - (R_t/K_{t-1}).$$

We implement equation D1 at the two-digit level using BEA two-digit depreciation rates to measure δ_t . Note, however, that real retirements depend on the in-use depreciation rate. That is, real retirements are given by

$$(D2) \quad R_t = \sum_{j=0}^{t-1} (v_{jt}/p_{jt})(1 - \delta^a)^{t-j},$$

36. We construct the capital stock for structures in order to be able to compute the cost share of equipment.

where v_{jt} is the component of nominal retirements in year t associated with capital purchases made in period j , and pi_j is the new capital deflator for period j . Nominal retirements are explicitly available in the LRD from 1977 to 1988. For the period before 1977, retirements can be measured from the identity relating the change in book value from the beginning to the end of the period, expenditures, and retirements. Equations D1 and D2 are used to solve for the in-use depreciation rate at the two-digit level under the assumption that each plant exhibits a FIFO retirement pattern. The two-digit in-use depreciation rates are then used with the new capital deflators, plant-level data, and equation A5 to measure real retirements at the plant-level.

Additional details of the estimations of the terms and variables in equation 5 are as follows. The value of output at the plant level is measured by adding the nominal change in inventories to the nominal total value of shipments, and then deflating by the four-digit output deflator from the Gray-Bartelsman data. The capital cost shares are calculated using the plant-level cost data on wages and salaries, materials costs, energy costs, cost of contract work, and capital costs. Capital costs are measured by using two-digit rental prices from the Bureau of Labor Statistics and the constructed plant-level capital stocks on equipment and structures.

The sample used in this analysis is drawn from all large, continuously operating plants for the period 1972–88. The sample terminates in this year because retirements were collected in the ASM until 1988. Further, a small number of plants (about 100) are excluded from the analysis because the above procedures yield extremely large investment rates. Specifically, plants with negative investment rates in excess of 100 percent in absolute magnitude (in principle, this is impossible) and those with positive investment rates in excess of 500 percent are excluded. The results are not sensitive to the precise value of the cutoff for positive rates; for example, the results are essentially the same for only excluding plants with positive investment rates in excess of 1,000 percent.

Comments and Discussion

Michael Woodford: This is a very impressive paper. I am impressed, first of all, by the sheer size of the data set that the authors have had to characterize and the amount of computational effort that has obviously been required, both in computing the parameters of their model and in performing the various simulation exercises.

I am also impressed by the fact that two of the authors studied the ability of this type of model to explain aggregate investment dynamics in an earlier paper that looked only at the properties of (sectoral) aggregate series and concluded (on the basis of rather indirect evidence) that the model did well.¹ I find it very courageous of them to take on a coauthor with a lot more data and actually look at the plant-level data to see if the model for which they had already declared victory would survive this much more intensive empirical scrutiny.

The results that the authors find are quite striking. Their framework for organizing the data brings into focus some very suggestive regularities. Indeed, the data appear to be quite consistent with the general theoretical framework from which the authors approach the problem. However, the data provide much more support for some aspects of their analysis than for others.

The authors argue that investment decisions at the microeconomic level have straightforward determinants (including an important role for a Jorgensonian “cost of capital”), but that these relations are masked by aggregation, if one seeks to explain aggregate investment in terms of average values of the shock variables. Their explanation of

1. Caballero and Engel (1994).

why aggregation is problematic involves two elements. First, they argue that investment is “lumpy”; adjustment occurs at discrete times at the plant level, even though one observes relatively continuous variations in aggregate investment. Second, they argue that there is a nonlinear relationship at the plant level between “mandated investment” (the discrepancy between the plant’s desired and actual capital stocks) and the average rate of investment that results. In particular, capital shortfalls result in proportionally greater responses than do capital excesses, as stressed in the literature on the irreversibility of investment.

The results presented offer some support for both of these sources of aggregation problems. However, it is the evidence of nonlinearity (and in particular, asymmetry of the kind associated with irreversibility) that is most striking. The support for discrete adjustment is much less clear-cut. Specifically, it seems that the data do not support a simple (S, s) model of discrete adjustment of the kind hypothesized by Caballero and Engel.²

This is worth commenting on because the great advantage of the analytical framework in this paper is that it allows one to model discrete adjustments of a relatively general sort. The main competing approach in modern work on investment, the “ q -theoretic” approach that estimates Euler equations for intertemporal substitution of investment spending, can deal with irreversibility, as long as investment (although possibly intermittent, in the sense that intervals of time pass with no investment) is not discrete.³ Hence if irreversibility is the main obstacle to the validity of an aggregative model, a disaggregated q -theoretic model might be appropriate, rather than the disaggregated “partial adjustment” model proposed by the authors.

Reviewing the main findings of the paper, the authors construct mandated investment levels for each plant in each year of their sample and compare these with the measured levels of investment spending. This requires construction of a capital stock series for each plant, and then a series for the plant’s desired capital stock for each year. The indirect method used to construct the latter series is clever. It amounts to observing that each plant’s output-to-capital ratio and its cost of capital are arguably cointegrated, although neither series is stationary,

2. See Caballero and Engel (1994).

3. See Abel and Eberly (1994).

and then interpreting the stationary linear combination of the two series as a measure of mandated investment. The idea is that in the absence of any costs of adjusting one's capital stock, the output-to-capital ratio would vary instantaneously with variations in the cost of capital in exactly the way that the two series are observed to comove in the long run. Thus any observed temporary fluctuations in the stationary cointegrating vector are attributed to failure to adjust the plant's capital stock immediately to its desired level, and are taken to be proportional to the discrepancy between actual and desired capital.

The authors describe a special case in which this constructed series would correctly measure mandated investment, although they make no effort to demonstrate that their data are consistent with the implications of that special case. For example, the derivation in Caballero and Engel's earlier paper relies upon the shock variables following random walks but no attempt is made in the present paper to show that they do so, even though series are constructed for the shocks to each plant's desired capital.⁴ Similarly, the derivation in appendix B implies that equation 5 should hold with θ_i equal to -1 , but in the empirical work θ_i is treated as a free parameter; the estimated values used in constructing mandated investment series range between -0.01 and -2.0 . (Another problematic feature of the use of values for θ_i other than -1 is that if the cointegrating relation is really of that sort, then there should be no stationary "cost share of equipment capital." Yet such a cost share is assumed to be measurable for each sector and is used to calibrate the value of η_i in equation 5.)

The primary case to be made for the validity of the proposed measure of mandated investment is that actual investment is clearly related to it. This is shown in the second panel of figure 8 by the nonconstancy of the plot of the average investment rate as a function of the mandated investment rate. Not only is the plot upward-sloping, but many of its gross features are those that the authors' framework would predict. Average investment is near zero when mandated investment is near zero. The response to positive mandated investment is much stronger than the response to negative mandated investment of the same size (indicating an important degree of irreversibility); and the response to large positive values of mandated investment is larger to an extent that

4. See Caballero and Engel (1994).

is more than proportional to the increase in mandated investment (as would follow from a model of discrete adjustment in which the increased profits from the adjustment must be large enough to offset a fixed cost of adjusting). The plausible character of this plot suggests that the authors' method does capture some important features of the data.

Still, it is important to be clear about what is *not* established by this finding. The authors stress that their estimates suggest significant effects of cost of capital variations (and hence of tax incentives) upon investment at the microeconomic level, although previous studies using aggregate data have typically failed to find them. But it is not clear how much evidence there is in their results of cost of capital effects. The second panel of figure 8 shows a strong effect of their measure of mandated investment, and that measure involves the cost of capital. But there is no demonstration that the strong relation between average investment and mandated investment owes anything to the way that cost of capital variations are used in constructing the mandated investment series. It seems likely that most of the variation in mandated investment in the authors' sample is due to variation in the output-to-capital ratio (over time and across plants), and that the relation found between mandated and actual investment largely has to do with the well-known ability of "accelerator" equations to predict investment.

And the second panel of figure 8, in itself, provides little evidence in favor of the authors' hypothesis of discrete adjustment. (Note that the asymmetry of this plot is important evidence in favor of *irreversibility*, and hence of a need to go beyond a simple aggregative framework in order to account for variations in aggregate investment.) They suggest that the convexity of the plot for positive values of mandated investment (or equivalently, the fact that the adjustment rate function is increasing in the first panel of figure 8) indicates "increasing returns in the adjustment technology." But a similar plot could equally well be generated by a model with convex adjustment costs (and hence continuous adjustment), if the marginal profits associated with an increase in the capital stock were sufficiently steep at low levels of capital stock relative to output.

The authors also suggest that the skewness and excess kurtosis of the distribution of plant-level investment rates (shown in the first panel of figure 6) indicate discrete adjustment, since the distribution of plant-

level shocks (shown in the second panel of figure 6) has neither property. But while this finding surely indicates a *nonlinear* response to the shocks, it need not involve discrete adjustment. In fact, a model of continuous adjustment with the kind of nonlinear response required by irreversibility could easily result in a distribution of investment rates with both properties.

A simple example will illustrate this. Suppose that a plant continuously adds to its capital stock, with an investment rate i given by

$$(E1) \quad i = Ae^{\gamma x},$$

where $A > 0$, $\gamma > 0$, and x is mandated investment. Note that equation E1 is an increasing, convex function like the plot in the second panel of figure 8, and investment is never negative, in accordance with the idea of irreversibility. As in the text, mandated investment is given by

$$(E2) \quad x = \bar{k} - k,$$

where \bar{k} evolves in response to exogenous shocks. Let \bar{k} follow a Brownian motion with drift,

$$(E3) \quad d\bar{k} = \alpha dt + \sigma dw,$$

where w is a unit Wiener process. (Note that this implies that the distribution of shocks is normal, as in the second panel of figure 6.) The evolution of the capital stock is finally given by

$$(E4) \quad dk = [i - \delta]dt.$$

Equations E1–E4 imply a law of motion for mandated investment of the form

$$dx = [\alpha + \delta - Ae^{\gamma x}]dt + \sigma dw.$$

The stationary distribution for mandated investment, $f(x)$, therefore satisfies a diffusion equation of the form

$$[\alpha + \delta - Ae^{\gamma x}]f(x) = \frac{\sigma^2}{2}f'(x).$$

It follows that

$$f(x) = B[\exp(vx - \rho Ae^{\gamma x})],$$

where $\nu \equiv 2(\alpha + \delta)/\sigma^2$, $\rho \equiv 2/\gamma\sigma^2$, and $B > 0$ is chosen so as to make f a probability density. Transforming variables using equation E1, the stationary distribution for the investment rate, $g(i)$, is given by

$$(E5) \quad g(i) = C(i^{\nu-\gamma/\gamma})e^{-\rho i}$$

over the range $i > 0$, where $C > 0$ is chosen so as to make g a probability density.

One observes that because of the nonlinearity of relation E1, the normally distributed shocks result, nonetheless, in a decidedly non-normal distribution of investment rates (equation E5). This distribution is easily seen to exhibit skewness and excess kurtosis, as in the first panel of figure 6. Yet there are no large, discrete adjustments in the hypothesized investment dynamics—simply a large *continuous* rate of investment in the case of a large capital shortfall.

The only direct evidence that the authors report relating to the discreteness of adjustment is the distributions of investment rates reported in the third and fourth panels of figure 8. Here they present complete histograms for the distribution of investment rates observed for plants with a particular level of mandated investment, rather than simply the average investment rate as in the second panel of figure 8. These histograms are suggestive of discrete adjustment in that in both cases the largest bin is the one that includes zero investment. This remains true even when mandated investment is large, although in that case the probability of zero investment falls and large positive investment rates become more frequent. Certainly the histograms are not consistent with the simplest sort of continuous adjustment model, under which one would expect the investment rates always to cluster tightly around a modal rate that tracks the average investment rate plotted in the second panel of figure 8.

At the same time, they do not much suggest a simple (S,s) model of the kind hypothesized in Caballero and Engel's earlier paper, for the nonzero investment rates do not cluster around a single value corresponding to "full adjustment."⁵ (In this respect, the results here are rather different than those reported in the authors' companion paper that deals with employment adjustment. In that paper, the corresponding histograms do often have two sharp peaks, at zero and full adjust-

5. See Caballero and Engel (1994).

ment, respectively.⁶) Instead, the most frequent positive investment rate continues to be the smallest one, even when mandated investment is large; and frequency falls off nearly monotonically with the size of the investment rate, in the case of both positive and negative rates.

These histograms do not much suggest any tendency of investment to be lumpy, and certainly not a tendency to occur in lumps of the size of the authors' measure of mandated investment. Rather, they instead suggest a process of *intermittent* continuous investment. Because there are frequent spells of zero investment, the zero bin is the most frequent. But when investment is observed during a year, the total investment over the year is often small (because much of the year may have consisted of spells of zero investment), and so the small-investment bins are the next most frequent, and so on. A large cumulative investment over the year would be observed only when the incentive for investment becomes so large that investment occurs continuously at a significant rate over the entire year. If the measure of mandated investment were only a very noisy proxy for plants' actual incentives to invest each year, then one would expect a distribution of outcomes in the case of a given measured value of mandated investment that would always include a sizable number of plants with zero investment, although fewer when mandated investment is measured to be high.

Intermittent continuous investment of this sort is predicted by a model with convex costs of adjustment combined with irreversibility (for example, a lower price for sales of capital than for purchases), as shown by Avinash Dixit and Robert Pindyck, and Andrew Abel and Janice Eberly.⁷ Thus an important order of business remains the determination of the extent to which a model of that type, rather than one emphasizing discrete adjustment, can account for the patterns in the plant-level data that are identified by the authors.

None of this detracts from the importance of the authors' conclusion that because the response of actual investment to mandated investment is significantly nonlinear, one cannot hope to explain aggregate investment solely in terms of the average values of the various variables involved in the model. Their findings do suggest significant nonlinearities in investment dynamics, and of an intuitively plausible sort. These

6. See Caballero, Engel, and Haltiwanger (1995).

7. See Dixit and Pindyck (1994) and Abel and Eberly (1994).

appear to result, on the one hand, from a greater difficulty in shedding capital than in acquiring it, and on the other, from a greater marginal benefit from acquiring capital when the capital shortfall becomes large. Insofar as this is true, it is unavoidable that the sensitivity of aggregate investment to shocks will vary over time, as the authors argue, and that an adequate model of investment dynamics will have to track higher moments of the distributions of various state variables, as they here propose.

Robert E. Hall: Ricardo Caballero, Eduardo Engel, and John Haltiwanger break a tremendous amount of new ground in the field of empirical research on investment. Although advances in investment theory have been impressive in recent years, until this paper empirical applications have been disappointing.

The accomplishment shown in figure 8 is nothing short of breathtaking. In the first panel, the vertical position of the curve is the ratio of actual investment to the amount of investment mandated by investment theory, averaged over all of the plants in a category of mandated investment that is one percentage point wide. Mandated investment is positive to the right of the vertical line and negative to the left. Plants with high levels of mandated investment actually invest almost 30 percent of the mandated amount. The adjustment rate falls to around 10 percent if mandated investment is small but positive. Remarkably, plants with low levels of mandated disinvestment have actual disinvestment that is also around 10 percent of the mandated amount. But the adjustment rate for high levels of mandated disinvestment is in the range of 3 to 5 percent.

The second panel of figure 8 shows investment itself, in place of the investment rate in the first panel. The nonconstancy of the investment rate is shown as convexity of investment. There is a little mystery here, not explained in the text. Investment is positive, on the average, for plants with small negative levels of mandated investment. The corresponding adjustment rate ought to be large and negative, but is shown in the graph as moderate and positive.

The nonlinearity found in figure 8 gives strong support to modern investment theory. Plants are expanded fairly aggressively when mandated investment is positive, but to shrink a plant when mandated investment is negative is much more difficult. Irreversibility is a first-

order issue for investment. To put it differently, the bid-ask spread for industrial equipment is extremely high.

Two other figures strongly confirm the validity of the specification and the importance of the nonlinearity. First, the model is well specified if the measured adjustment functions are stable over time, and the explanation of the time-series behavior of investment comes from changes in mandated investment and movements along the adjustment function. Figure 10 shows that freezing the adjustment function over time robs the model of relatively little of its explanatory power. Second, the ideas in the paper are important if a good fraction of the volatility of investment comes from movements along the nonlinear adjustment function. Figure 17 shows the importance of these movements relative to the movements caused by changes in mandated investment.

One of the big contributions of this paper is to make use of census data on capital retirements. The authors assume, as does the BEA, that capital depreciates in use but still has positive marginal product when it is scrapped. Retirements include sales of equipment, where that assumption makes sense. For true retirements depreciation is, presumably, already complete. In fact, equipment that is no longer in use may remain in the factory for many years before it is removed to make room for new equipment.

The introduction to the paper suggests that the authors do something quite interesting: study purchases of new equipment and retirements of old equipment separately. In fact, the only variable that they consider is purchases less retirements. Thus in figure 8 the positive adjustment rate on the disinvestment side means that retirements exceed purchases.

Haltiwanger, in his highly productive collaboration with Steven Davis, has become famous for developing an indirect measure of gross changes in employment. The Davis-Haltiwanger measures of job creation and destruction have taught analysts a tremendous amount about the operation of the labor market. This new knowledge is incremental to what we learned from the difference between creation and destruction, net employment change. I would love to see the same philosophy brought to the investment side. Even though it would raise more questions than it would answer, a new version of figure 8 showing separately the adjustment rate for purchases of equipment and that for retirements would be an important addition to the paper. There would likely be nonzero adjustment for retirements in the area of positive mandated

investment. Firms may wait to retire equipment that is not in use until they need to make room for new equipment. As a general matter, the separation of the adjustment functions for new purchases and retirements may help to explain the role of reorganization in the aftermath of recessions.

The concept of mandated investment is central to the measurement exercise of this paper. Mandated investment is a throwback to investment theory of the 1960s. Just as Edmund Phelps and Sidney Winter were writing in 1970 that “a landing on the non-Walrasian continent has been made. Whatever further exploration may reveal, it has been a mind-expanding trip: We need never go back to $\dot{p} = \alpha(D - S)$ and $q = \min(D, S)$,” investment theory moved beyond the idea of defining disequilibrium and then asserting an adjustment process driven by that disequilibrium.¹ The investment theory of the 1970s assumed convex adjustment costs and showed that there was a simple relation between the flow of investment and the value of a costate variable, usually called q . Moreover, it was believed that q might be measurable from data on the market values of firms, although this idea has never worked out, in practice. Even without direct measures of q , the Euler equation approach to investment has flourished.

More recent investment theory has assumed nonconvex costs, and therefore has come to grips with infrequent large adjustments in the capital stock.² The idea of a mandated capital stock reappears in this theory as a means of summarizing the demand facing the firm and the factor prices that it pays. Moreover, the mandated capital stock does play a role in the adjustment process, in that a firm definitely will not adjust today if its capital is close to the mandated value. On the other hand, the adjustment process itself takes a hammer-and-tongs dynamic programming analysis. The firm does not adjust mechanically to the mandated level. It usually pays to overshoot, once adjustment becomes appropriate.

The mandated, or desired, capital stock has proven to be an elusive concept in investment theory. Dale Jorgenson was the first to derive an expression for desired capital, K^* , when the technology admits factor substitution.³ Jorgenson posed the problem of maximizing the value of

1. Phelps and Winter (1970, p. 337).
2. See Caballero and Engel (1994).
3. See Jorgenson (1963).

a firm with a Cobb-Douglas technology. His model did not include adjustment costs or time to build. As a result, his K^* would hold for frictionless investment. He used the resulting formula in the framework mentioned above: the firm adjusts its capital stock toward K^* .

Jorgenson's framework suffers from two shortcomings. First, the value maximization problem is ill posed for a competitive firm with constant returns to scale. The output and desired capital of such a firm, in this context, is either infinite, indeterminate, or zero. Second, Jorgenson states K^* as a function of the firm's output and the rental price of capital. When he evaluates K^* he uses actual output, not the level of output that the firm would have chosen in the frictionless setting.

As a result, Jorgenson's K^* has the following rather contorted definition: K^* is the amount of capital that the firm would hold, absent frictions, if the wage were sufficiently different so that the firm would choose the same output as it actually produces, given the constraint of the amount of capital already installed.

In the subsequent literature on investment these issues are side-stepped, rather than solved. The problem of value maximization is well posed with constant returns for a firm with adjustment costs, and the concepts derived in the theory of investment with adjustment costs are easy to match with their empirical counterparts. In the recent theory with nonconvex adjustment costs and lumpy adjustment, the problem of value maximization has been made well posed by assuming imperfect competition.

However, Caballero, Engel, and Haltiwanger use a concept of mandated capital that is quite close to Jorgenson's K^* ; it is about as contorted. In their paper K^* is the amount of capital that the firm would use if it first optimized its use of variable factors, given the amount of capital currently available, and then optimized its amount of capital, taking as given the previously optimized levels of the variable factors. This definition makes K^* dependent on the last period's actual capital stock, which may seem rather odd. But it does have the desirable property that a firm that would make more profit with more capital will have K^* in excess of its actual capital, and vice versa.

There is no good answer to the K^* conundrum. That is, there does not seem to be a good measure of the firm's need for capital in the

quantity dimension. Apparently there is no good measure in the price dimension, either.⁴

The estimation of the elasticity of demand for capital in the paper duplicates efforts in the vast empirical literature on the estimation of factor demand. The results seem to be consistent, but there ought to be more discussion of the connection with these other studies.

Finally, this paper perpetuates the Hall-Jorgenson mistake of talking about the “effects” of tax policy.⁵ Given what is being held constant, that word should be banished. What is measured is the shift in the firm’s capital demand function.

But these are minor quibbles. This paper amply repays the effort needed to understand it. After decades of complaining about the defects of aggregate equations, someone has finally done something constructive to disaggregate investment.

General Discussion

Anil Kashyap noted an interesting implication of the paper for monetary policy: In setting interest rates, central banks should take account of the underlying distribution of a firm’s conditions. Minor changes in interest rates might lead to large investment responses when many firms are near an investment trigger, and lead to a much smaller response when many firms have small, or even negative, discrepancies between their desired and actual capital stock. For example, the central bank ideally should track not only the volume of new orders, but also the fraction of firms that are placing them.

Kashyap also noted that the cost of capital, as constructed in the paper, assumes that the real interest rate is constant; variations in the cost of capital come from differences in taxes, depreciation, and industry price deflators. He was skeptical about this assumption and surprised that, despite it, the results are as good as they are. He encouraged the authors to explore alternative specifications in which the real rate varied across time, and possibly varied by the size and financial structure of firms.

William Brainard observed that in order to provide the rich distri-

4. See Caballero and Leahy (1995).

5. See Hall and Jorgenson (1967).

butional information on the discrepancy between desired and actual capital stock and on investment conditional on these discrepancies, the authors had sacrificed information about dynamics, at both the firm and the aggregate levels. The simulations of the effect of a change in the cost of capital on aggregate investment, for example, take the distribution of mandated investment as given and do not trace the implications of the induced investment on the distribution of mandated investment in future periods, which itself is a determinant of investment in subsequent periods. Brainard noted that, at the firm level, the model focuses on investment one period in the future. Yet some investments, for example, those following from a decision to build an entire new plant or production facility, will be spread out over several periods. In the same spirit, he noted that some investment is undertaken to meet expected future demand, rather than in response to low capacity relative to current output or to the retirement or obsolescence of the existing capital stock. He wondered whether there was any evidence in the data that investment “caused” output, rather than the other way around.

References

- Abel, Andrew B., and Janice C. Eberly. 1994. "A Unified Model of Investment Under Uncertainty." *American Economic Review* 84(5): 1369–84.
- Bar-Ilan, Avner, and Alan S. Blinder. 1992. "Consumer Durables: Evidence on the Optimality of Usually Doing Nothing." *Journal of Money, Credit, and Banking* 24(2): 258–72.
- Bartelsman, Eric J., and Wayne Gray. 1995. "The NBER Manufacturing Productivity Database." Unpublished paper. Board of Governors of the Federal Reserve System (January).
- Bertola, Giuseppe, and Ricardo J. Caballero. 1994. "Irreversibility and Aggregate Investment." *Review of Economic Studies* 61(2): 223–46.
- Caballero, Ricardo J. 1994. "Small Sample Bias and Adjustment Costs." *Review of Economics and Statistics* 76(1): 52–58.
- Caballero, Ricardo J., and Eduardo M. R. A. Engel. 1994. "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach." Working Paper 4887. Cambridge, Mass.: National Bureau of Economic Research (October).
- Caballero, Ricardo J., Eduardo M. R. A. Engel, and John Haltiwanger. 1995. "Aggregate Employment Dynamics: Building from Microeconomic Evidence." Working Paper 5042. Cambridge, Mass.: National Bureau of Economic Research (February).
- Caballero, Ricardo J., and John V. Leahy. 1995. "Fixed Costs: The Demise of Marginal q ." Unpublished paper. Massachusetts Institute of Technology (December).
- Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12(3): 383–98.
- Chirinko, Robert S. 1993. "Business Fixed Investment Spending: Modeling Strategies, Empirical Results, and Policy Implications." *Journal of Economic Literature* 31(4): 1875–1911.
- Cooper, Russell, John Haltiwanger, and Laura Power. 1995. "Machine Replacement and the Business Cycle: Lumps and Bumps." Working Paper 5260. Cambridge, Mass.: National Bureau of Economic Research (September).
- Cummins, Jason G., Kevin A. Hassett, and R. Glenn Hubbard. 1994. "A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments." *BPEA*, 2:1994, 1–59.
- Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment Under Uncertainty*. Princeton, N.J.: Princeton University Press.
- Doms, Mark, and Tim Dunne. 1993. "Capital Adjustment Patterns in Manufacturing Plants." Unpublished paper. U.S. Bureau of the Census, Center for Economic Studies.

- Hall, Robert E., and Dale W. Jorgenson. 1967. "Tax Policy and Investment Behavior." *American Economic Review* 57(3): 391–414.
- Harrison, J. Michael, Thomas M. Sellke, and Allison J. Taylor. 1983. "Impulse Control of Brownian Motion." *Mathematics of Operations Research* 8(3): 454–66.
- Hulten, Charles R., and Frank C. Wykoff. 1981. "The Measurement of Economic Depreciation." In *Depreciation, Inflation, and the Taxation of Income From Capital*, edited by Charles R. Hulten. Washington: Urban Institute.
- Jorgenson, Dale W. 1963. "Capital Theory and Investment Behavior." *American Economic Review, Papers and Proceedings* 53(2): 247–59.
- Phelps, Edmund S., and Sidney G. Winter, Jr. 1970. "Optimal Price Policy under Atomistic Competition." In *Microeconomic Foundations of Employment and Inflation Theory*, edited by Edmund S. Phelps and others. New York: Norton.
- Rotemberg, Julio J. 1987. "The New Keynesian Microfoundations." In *NBER Macroeconomics Annual 1987*, edited by Stanley Fischer. Cambridge, Mass.: MIT Press.