FIXED COSTS OF ADJUSTMENT, COORDINATION, AND INDUSTRY INVESTMENT

Joanne M. Doyle and Toni M. Whited*

Abstract—We test whether smooth industry-level investment dynamics result from explicit aggregation of asynchronous and possibly lumpy firm-level investment. We compare the deviations of optimal from actual firm behavior across industries categorized by their ratios of idiosyncratic uncertainty to the sum of idiosyncratic and aggregate uncertainty. The deviations are represented by the residuals of a cointegrating regression that is derived from the firm's first-order condition under no adjustment costs. In support of models with asynchronous firm decisions, we find a significant negative relationship across industries between idiosyncratic uncertainty and the persistence of these residuals.

I. Introduction

FINDING a model to explain aggregate investment dynamics has baffled macroeconomists for years. The slow capital stock adjustment evident at various levels of aggregation has suggested models in which representative agents smooth investment over time because of increasing, convex costs of adjusting the capital stock. These models have the advantage of providing direct connections between theory and data, in the form of either a decision rule or an Euler equation. However, as noted in the survey by Chirinko (1993), most of their numerous confrontations with both microeconomic and macroeconomic data have been disappointing. Furthermore, the models are at odds with the lumpy adjustment observed in plant-level data (Doms & Dunne, 1998; Cooper, Haltiwanger, & Power, 1999). For example, Doms and Dunne (1998) find that approximately 25% of an average plant's cumulative investment over seventeen years is concentrated in a single year. More recent theoretical research has proposed an alternative source for smooth aggregate investment dynamics, suggesting that they result from aggregation of infrequent, asynchronous, and sometimes lumpy firm-level investment,[1] which is in turn due to fixed costs of adjustment or irreversibility. For example, Blinder (1981) and Caplin (1985) study fixed-costs models in the context of inventories, and Bertola and Caballero (1994) develop and implement a model of investment irreversibility.

The intuition behind this latter class of models is straightforward. First, consider an individual firm. If it faces fixed adjustment costs, it invests only when its capital stock is sufficiently far from the optimal level, otherwise preferring to remain inactive to avoid any lump-sum costs. As in Caballero and Leahy (1996), for example, this sort of lumpy investment behavior often takes the simple form of an inaction or "(S,s)" interval for the difference between the firm's actual and optimal capital stocks. As long as the difference resides in the (S,s) interval, the firm does nothing. When a combination of capital stock depreciation and cumulated shocks to the net profitability of capital bring the deviation to an interval endpoint, the firm adjusts. The deviation usually differs from zero, although the firm's policy of preventing the deviation from leaving the (S,s) interval induces it to be stationary.

Although intuitively appealing, these models have received limited empirical attention because, unlike their representative-agent counterparts, they do not provide direct connections between theory and data. The empirical work accomplished so far has, nonetheless, been encouraging. Caballero, Engel, and Haltiwanger (1995) find support for a fixed-costs model in the behavior of U.S. manufacturing plants: most plants invest only when the difference between their actual and desired capital stocks is large. Similarly, Cooper et al. (1999) estimate that the probability that a plant will experience an episode of intense investment is increasing in the time since the last episode. Cooper and Haltiwanger (1998) fit a model that encompasses both fixed and convex adjustment costs to microeconomic plant-level data and find that both types of costs are important. Using aggregate U.S. data, Bertola and Caballero (1994) find that a model containing microeconomic investment irreversibility is broadly consistent with aggregate investment fluctuations. Caballero and Pindyck (1996) implement a model with irreversibility and free industry entry, finding, as predicted by their model, that an increase in aggregate uncertainty raises the industry entry trigger. Finally, Caballero and Engel (1999) examine annual industry-level data, observing evidence of the types of nonlinearities in dynamic investment behavior consistent with a model in which firms face random fixed costs. They also find that their fixed-costs model outperforms a simple partial-adjustment model.

We approach the problem of testing these models from a new angle, by looking at one of their implications for aggregate investment behavior. We exploit the idea that, when many firms follow (S,s) types of investment policies, the behavior of their aggregated investment depends on the extent to which they synchronize their actions. For example, if individual firms experience no idiosyncratic uncertainty and are subject only to aggregate shocks, then the behavior of the aggregate should resemble the behavior of the individual, and aggregate investment should occur episodically. Further, the aggregate deviation between the actual and

Received for publication February 12, 1998. Revision accepted for publication November 29, 2000.

* James Madison University and University of Iowa, respectively.

We are grateful to Mike Arnold, Tim Erickson, Jane Ihrig, John Leahy, and two anonymous referees for helpful comments. Jennifer Westberg provided excellent research assistance. We also thank Alan Auerbach and Kevin Hassett for providing us with their tax adjustments for the user cost of capital.

[1] For convenience, we use the term *firm* to refer to the relevant microeconomic decision-making unit, whether it be a firm, subsidiary, or plant.

The Review of Economics and Statistics, November 2001, 83(4): 628–637

desired capital stocks should display substantial persistence. To understand this implication, note that most theoretical (S,s) models either specify or derive a nonstationary process for the deviation, which makes sense because there would be no need to regulate the deviation to an interval if the deviation were sufficiently mean reverting. It is only these bands that induce stationarity, and the deviation spends a great deal of time away from zero. On the other hand, if firms experience only idiosyncratic shocks, then, as demonstrated in a number of papers such as Bertola and Caballero (1990),[2] the cross-sectional distribution of the difference between the actual and optimal capital stocks will have a near-zero mean. Idiosyncratic shocks continually reshuffle the cross-sectional positions of the individual firms, leaving the mean roughly constant. Therefore, the aggregate capital stock should track the frictionless optimum, and we should observe little persistence in the aggregate deviation of the optimal and actual capital stocks. Ceteris paribus, as the ratio of aggregate to idiosyncratic uncertainty increases, we should observe an increase in the persistence of the deviation.[3] It is this proposition that we test.

Although we have couched our argument within the confines of lumpy adjustment models, models that predict smooth investment at the microeconomic level have similar implications. For example, the aggregate capital stock will tend to deviate persistently from its frictionless optimum if there is little idiosyncratic uncertainty because convex costs imply that individual firms will adjust slowly to an optimal capital stock after a shock to productivity. Further, if firms face idiosyncratic shocks, the deviations between their actual and optimal capital stocks will tend to average out cross-sectionally. Nonetheless, we prefer to emphasize lumpy adjustment models for two reasons. First, as noted by Foote, Hurst, and Leahy (2000), the intermittent adjustment observed at the microeconomic level in the studies previously mentioned can be thought of as prima facie evidence in favor of these models. Second, we present some evidence that contradicts one of the main predictions of smooth adjustment models.

We use two different data sources. First, our information on quantities and prices comes from quarterly, two-digit, industry-level data. Quarterly data have an important advantage over the readily available, annual, industry-level data in the NBER productivity database or in COMPUSTAT because low-frequency data may obscure the lack of coordination among individual firms by aggregating their actions over time. To avoid this problem, we have constructed our own quarterly series.

Second, to measure idiosyncratic and aggregate (in this case, industry-wide) uncertainty, we use stock returns. These data have the obvious disadvantage that not all of the investment in our industry data comes from firms listed on major stock exchanges. On the other hand, stock prices capture many sources of uncertainty relevant to investment decisions, such as output prices, costs, interest rates, and inflation.[4] We use these data in two ways. First, we use ex post measures of uncertainty calculated over the entire sample period. We also try an ex ante measure of uncertainty that is calculated pre-sample. We report results obtained using both of these procedures; as will be seen, the evidence is similar and supports the same inference.

Using these data, we compare the deviations of optimal from actual firm behavior across industries categorized by their ratios of aggregate to idiosyncratic uncertainty. The deviations are represented by the residuals of a cointegrating regression that is derived from the firm's first-order condition under no adjustment costs. Using two sources of data distinguishes our method from previous work using aggregated data. The work with aggregated data mentioned previously has used one data set to compare observed investment behavior with that implied by a model that dictates the relationship between the uncertainty ratio and the time series properties of the deviation.[5] In contrast, we use two data sets to compare persistence in the deviation and uncertainty ratio across industries. In support of models with asynchronous firm decisions, we find that industries with high idiosyncratic uncertainty have residuals with low persistence and that industries with low idiosyncratic uncertainty have residuals with high persistence. We also provide evidence that this result is not an artifact of our ceteris paribus assumption and that our results are inconsistent with simple convex adjustment cost models.

The paper is organized as follows. We outline our empirical model in section II and describe our data in section III. Section IV presents our results and discusses their robustness to alternative interpretations. We provide concluding remarks in section V.

II. Methodology

Under the fiction of no costs of adjustment, the firm's desired stock of capital satisfies the well-known condition

[2] Although Bertola and Caballero (1990) examine the case of consumer durables, their model is general enough to encompass the case of firm investment.

[3] An important exception is that of one-sided (S,s) models. As demonstrated by Caplin and Spulber (1987), microeconomic adjustment costs may have no effect on the behavior of the aggregate in this class of models. Nonetheless, we believe that this result is not a problem for our application, primarily because one-sided (S,s) models derive from situations in which a decision-making unit would want to adjust in only one direction. Because, in the absence of adjustment costs, both downward and upward adjustment may be optimal, a one-sided (S,s) model is unlikely to characterize investment.

[4] Leahy and Whited (1996) also use stock prices to examine the impact of uncertainty on investment.

[5] The one exception here is Caballero and Pindyck (1996), whose model implies no role for idiosyncratic uncertainty. They assume a linear relationship between the marginal revenue product of capital and idiosyncratic shocks, which allows an expectation operator to average away the effects of these shocks.

that the marginal product of capital equals the user cost of capital: the relative price of capital goods times the sum of the interest rate and the depreciation rate. Given the assumptions of a constant elasticity of substitution production function and isoelastic demand, this first-order condition can be written as

$$k_{it}^* - y_{it} = \alpha_0 + \alpha_1 r_{it}, \tag{1}$$

where all variables are in natural logarithms, k_{it}^* is the desired capital stock of firm i at time t, y_{it} is output, and r_{it} is the user cost. Note that the parameter α_1 is the elasticity of substitution between capital and labor.

If the firm faces costs of adjustment, the actual observed capital stock, k_{it}, will deviate from k_{it}^* by an amount z_{it}. Under a variety of conditions this deviation, z_{it}, will be regulated to lie in a nonstochastic (S,s) type of interval and will therefore be stationary. For example, an (S,s) interval will arise as the firm's optimal investment policy in a partial equilibrium model in which the firm takes prices and productivity shocks as exogenous, if prices and productivity shocks follow geometric Gaussian random walks and if the firm faces nonstochastic adjustment costs proportional to the size of the firm. If capital goods are indivisible, then the second condition can be replaced by a price differential between new and used capital goods.

This observation leads to the standard cointegrating relationship

$$k_{it} - y_{it} = \alpha_0 + \alpha_1 r_{it} + z_{it}. \tag{2}$$

Aggregating this log-linear equation over firms, we obtain

$$k_t - y_t = \alpha_0 + \alpha_1 r_t + z_t, \tag{3}$$

where we have dropped the i subscript to indicate an aggregate variable.[6]

As discussed by Caballero (1994), estimating this cointegrating regression by OLS can produce significant small-sample biases, because the error term z_t results from the presence of adjustment costs. We therefore estimate equation (3) using the dynamic OLS estimator in Stock and Watson (1993), in which leads and lags of the first differences of the regressors are included to reduce the small-sample biases. We choose the lag and lead length by starting with a large number and then decreasing it until the excluded regressors are jointly significant.

[6] Note that our data represent logs of aggregate variables, whereas our theoretical model holds for the sum of the logs of individual variables. Using arithmetic averages to approximate the theoretically correct geometric averages introduces an approximation bias. Lewbel (1992) shows that this bias will be small to the extent that changes in the mean (across firms) of the user cost do not affect the cross-sectional distribution of the user costs (relative to this mean). This condition will hold, for example, if tax reforms or changes in interest rates affect all firms in an industry uniformly.

III. Data

We have been able to gather quarterly data on the output, capital stock, and tax-adjusted user cost of capital for ten, two-digit, U.S. manufacturing industries covering the period 1967:2–1992:4. Of the twelve two-digit industries for which we have investment data, we exclude SIC 36 (electric machinery) due to significant industry redefinition. Because the investment data come from a survey of companies and the output data from a survey of establishments, we exclude SIC 29 (petroleum) following Abel and Blanchard (1988), who report that discrepancies between the two surveys are problematic only for this industry.

We construct seasonally unadjusted output as the sum of real shipments and the change in real finished goods inventories. Because our shipments data from the Bureau of Economic Analysis are seasonally adjusted, we unadjust them as in Reagan and Sheehan (1985), using adjustment factors extracted from the seasonally adjusted and unadjusted nominal values from the Census Bureau's Monthly Manufacturers' Shipments and Inventories. The output price deflator is constructed from the real and nominal series for shipments. We convert the inventories data from a cost to a market basis following West (1983). The monthly values for output are aggregated to quarterly values by summing over the months, and the quarterly output deflator is defined as the geometric mean of the monthly deflator. We define value added as real output less materials input. For materials input, we use the annual fraction of materials usage to real output from the Annual Survey of Manufacturers and interpolate quarterly fractions, which we then apply to the quarterly output series.

Because published capital stock data are available only annually, we combine quarterly investment data from the Commerce Department's New Plant and Equipment Survey with annual capital stock data from the BEA to construct a quarterly capital stock series. We deflate investment expenditures using a weighted average of the price deflators for producers' durable equipment and nonresidential structures, both taken from NIPA, where the weights vary over industries and time and are interpolated from the industry-specific annual share of equipment and structures in new investment. We apply the perpetual inventory method to the investment data using the 1959 annual capital stock from the BEA as a benchmark for 1959:4.[7] Our depreciation rates are a weighted average of $\delta_E = 0.03$ per quarter for equipment and $\delta_S = 0.01$ per quarter for structures.

As an informal check on our data construction methods, we compare the ratio of value added to the capital stock with comparable figures obtained from the annual data in the 1997 annual, full coverage, and research COMPUSTAT files. For this latter data set, we first delete any firm-year

[7] Although separate capital stock series on equipment and structures would be preferable for our purposes, these data are not available.

TABLE 1.—A COMPARISON OF OUTPUT TO CAPITAL RATIOS IN COMPUSTAT AND OUR INDUSTRY-LEVEL DATA

SIC	Industry	COMPUSTAT	Industry
20	Food	0.230	0.263
22	Textiles	0.314	0.332
26	Paper	0.274	0.261
28	Chemicals	0.226	0.252
30	Rubber	0.298	0.322
32	Stone, clay, glass	0.248	0.292
33	Primary metals	0.202	0.251
34	Fabricated metals	0.482	0.486
35	Industrial machinery	0.239	0.342
37	Transportation equipment	0.293	0.291

TABLE 2.—UNCERTAINTY RATIO SUMMARY STATISTICS: 1967:2–1992:4

SIC	Industry	Firms	σ_I	σ_A	$\sigma_I/(\sigma_I + \sigma_A)$
20	Food	104–200	0.233	0.111	0.678
22	Textiles	39–93	0.236	0.158	0.598
26	Paper	41–83	0.206	0.115	0.642
28	Chemicals	124–361	0.272	0.136	0.667
30	Rubber	40–81	0.219	0.146	0.601
32	Stone, clay, glass	37–83	0.189	0.135	0.583
33	Primary metals	73–113	0.200	0.121	0.624
34	Fabricated metals	79–178	0.285	0.137	0.676
35	Industrial machinery	143–454	0.312	0.154	0.669
37	Transportation equipment	90–156	0.212	0.144	0.595

observations that have missing data on sales, the cost of goods sold, and net property, plant, and equipment. This screening device eliminates 11% of our observations. We then construct time series that run from 1977 to 1996 as follows. First, we measure value added as sales less the cost of goods sold. Then, for each year and each two-digit industry, we sum value added across firms and then divide by the sum of the reported net capital stocks. In table 1, we report the means of these series for both our industry data and our COMPUSTAT data. The two sets of figures appear to be quite comparable. Six of the ten means from the COMPUSTAT data are within 10% of the means from our data, and another three are within 20%. The one sizeable deviation occurs for SIC 35, which exhibits a deviation of 40%.

The user cost of capital, r_t, is a weighted average of the user costs for equipment and structures, r_t^E and r_t^S. Specifically,

$$r_t^E = \frac{p_t^E}{p_t^Y} \frac{(\delta_E + \rho_t)(1 - v_t - \tau_t z_t^E)}{(1 - \tau_t)}$$

$$r_t^S = \frac{p_t^S}{p_t^Y} \frac{(\delta_S + \rho_t)(1 - \tau_t z_t^S)}{(1 - \tau_t)},$$

where τ_t is the marginal corporate income tax rate;

v_t is the investment tax credit, applicable only to equipment expenditures;

p_t^E and p_t^S are price deflators for producers' durable goods and nonresidential structures;

p_t^Y is the industry-specific output deflator;

the real discount rate, ρ_t, comes from the Federal Reserve Board's quarterly FRB/US model and is a weighted average of the cost of debt and equity;

z_t^E and z_t^S come from Auerbach and Hassett (1992) and measure the present discounted value of future depreciation allowances for a \$1 investment made at time t; and

both r_t^E and r_t^S are annual rates at quarterly frequency.

Quarterly stock returns are from the 1996 CRSP tapes. Before describing our method for using returns to decompose uncertainty, we'd like to discuss the relationship between the z_{it} shocks and stock returns. As a first step, consider the example of a firm whose dividends are a constant proportion of its cash flows, whose cash flows follow a Markov process, and whose discount rate is constant. Then, the firm's stock price will be proportional to its current cash flow. Under perfect competition and constant-returns technology, the firm's cash flows will equal its marginal product, and the stock return will therefore equal the percentage change in the marginal product. Because, under our technological assumptions in section II, the z_{it}'s are proportional to the increment to the log of the marginal product,[8] the z_{it}'s will also be approximately proportional to returns. Clearly, the assumptions stated above will never hold exactly, but to the extent that they are approximately correct, our use of stock returns to measure uncertainty remains a useful methodology.

To measure idiosyncratic and industry-wide uncertainty, we use the following framework. Let x_{it} denote the return on firm i at time t, and suppose x_{it} can be decomposed as

$$x_{it} = \mu + w_{it} \equiv \mu + u_{it} + e_t, \tag{4}$$

where μ is the unconditional mean of x_{it}, u_{it} is an idiosyncratic component that may be serially correlated, and e_t is a stationary industry-wide component. To decompose the variance of x_{it}, we turn to the panel data literature and interpret equation (4) as a random-effects regression, in which the only regressor is a constant term. We estimate the variances of u_{it} and e_t following the method outlined by Greene (1997, pp. 626–627). Given this model of returns, it is natural to measure idiosyncratic uncertainty as the standard deviation of u_{it} and industry-level uncertainty as the standard deviation of e_t. We use an unbalanced panel of returns because we allow the number of companies to vary from quarter to quarter, depending on data availability. We include a company if its beginning-of-quarter and end-of-quarter daily prices are available for a given quarter.

Summary statistics for our uncertainty measures for the full sample period are in table 2. The column, labeled σ_I, gives the standard deviation of the idiosyncratic component of returns, the column labeled σ_A gives the standard devi-

[8] If the technology is CES, as in section II, the constant of proportionality will be the absolute value of the inverse of the elasticity of substitution.

ation of the industry-wide component of returns, and the column labeled $\sigma_I/(\sigma_A + \sigma_I)$ gives the percentage of the standard deviation due to idiosyncratic sources. We conjecture that the variation in this last set of figures stems from two sources. The first is the definition of the industry. For example, SIC 35, industrial machinery, is a very broad category and has a large $\sigma_I/(\sigma_A + \sigma_I)$, whereas SIC 30, rubber, is much narrower category and has a small $\sigma_I/(\sigma_A + \sigma_I)$. We are essentially exploiting outdated definitions of an industry to generate part of the variation in the uncertainty ratio. Second, in a related manner, $\sigma_I/(\sigma_A + \sigma_I)$ tends to rise with the average number of firms in an industry, the correlation between these two variables being 0.66.

At this point, it is useful to note that differences in the breadth of our industry classifications violate our implicit assumption that all types of cross-sectional heterogeneity can be captured by σ_I. Such an assumption violation can alter the interpretation our results. For example, some sub-industries in SIC 35, industrial machinery, are growing rapidly, yet others are stagnating. In this case, the z_{it}'s will be drawn from more than one distribution, and the differences in these distributions will alter the adjustment policies across the different sub-industries (even if adjustment costs are identical). Just as a high σ_I will lead to a lack of coordination, this heterogeneity in adjustment policies will do so as well. (See Caballero and Engel (1991).) Therefore, we cannot distinguish between differences in σ_I or differences in adjustment policies as a source of the results that follow. Nonetheless, we can attribute any results to differences in coordination among individual agents, which is the primary purpose of this study. In what follows, the reader should interpret the phrase "high σ_I" as indicating a lack of coordination.

IV. Results and Discussion

Before examining the behavior of the residuals, z_t, we test for the presence of unit roots in our series. For every industry, augmented Dickey-Fuller tests indicate that $k_t - y_t$ and r_t are individually I (1) with no drift. Further, using the test in Johansen (1988), we reject the null that $k_t - y_t$ and r_t are not cointegrated in each industry. This second result implies that the standard neoclassical model is a good characterization of long-run movements in the capital stock. It is also consistent with the stationarity of z_t implied by (S,s) models.

Table 3 contains our Stock and Watson (1993) cointegrating regressions, where we have divided our industries into three groups corresponding to the ratio $\sigma_I/(\sigma_I + \sigma_A)$. We present the coefficient estimates, their standard errors, the regression R^2's, and the variance of the residuals. We first note that our estimates of the elasticity of substitution, α_1, have small standard errors and range between -0.342 and -1.232, with five of the ten estimates insignificantly different from -1, the elasticity consistent with Cobb-Douglas

TABLE 3.—COINTEGRATING REGRESSIONS: 1967:2–1992:4

	SIC	α_1	R^2	σ_z
High $\sigma_I/(\sigma_I + \sigma_A)$	20	-0.387 (0.095)	0.156	0.007
	34	-0.411 (0.070)	0.256	0.009
	35	-0.178 (0.121)	0.024	0.014
Low $\sigma_I/(\sigma_I + \sigma_A)$	22	-0.813 (0.128)	0.333	0.011
	32	-0.824 (0.170)	0.224	0.020
	37	-0.628 (0.146)	0.178	0.031
Intermediate $\sigma_I/(\sigma_I + \sigma_A)$	26	-1.310 (0.166)	0.419	0.024
	28	-0.546 (0.076)	0.338	0.013
	30	-0.342 (0.159)	0.044	0.023
	33	-1.178 (0.165)	0.334	0.072

Standard errors are in parentheses under the parameter estimates.

technology. These numbers also indicate that policy designed to change the user cost of capital will have important long-run effects on capital accumulation. Next, it is interesting to compare the standard error of the regression residual, σ_z, across groups. For the industries with high idiosyncratic uncertainty (those with a high $\sigma_I/(\sigma_I + \sigma_A)$), σ_z ranges between 0.005 and 0.013. In contrast, for the low $\sigma_I/(\sigma_I + \sigma_A)$ industries, σ_z ranges between 0.018 and 0.071. This difference is consistent with higher persistence in z_t for the low $\sigma_I/(\sigma_I + \sigma_A)$ group, because more persistent series have higher variances in general.

We now turn to other indicators of persistence in z_t. Figure 1 plots the correlograms of z_t for the three groups of industries. For the group with high idiosyncratic uncertainty, the autocorrelation falls off to zero after five to six quarters. This evidence supports the idea that a high degree of idiosyncratic uncertainty causes substantial reshuffling of firms within an inaction interval. Therefore, when we add up the individual deviations, z_{it}, the sum dissipates quickly in response to aggregate shocks. In contrast, for the industries with low idiosyncratic uncertainty, it takes at least nine quarters for the correlograms to die out, which supports the idea that, if firms coordinate their actions, the aggregate deviation, z_t, should display substantial persistence. The correlograms for the industries with intermediate levels of $\sigma_I/(\sigma_I + \sigma_A)$ display varying degrees of persistence, which is at odds with the result that $\sigma_I/(\sigma_I + \sigma_A)$ is approximately equal for these industries. However, it is worth noting that our stock market-based measures of σ_A and σ_I are only proxies for the components of the standard deviations of the z_{it} shocks. Therefore, because there is noise in these measures, we expect some incorrect sorting of industries. This problem should be more severe when we move from either a low or high $\sigma_I/(\sigma_I + \sigma_A)$ group to the middle group than when we move from the low to the high group. In other

FIGURE 1.—ERROR CORRELOGRAMS: FULL SAMPLE

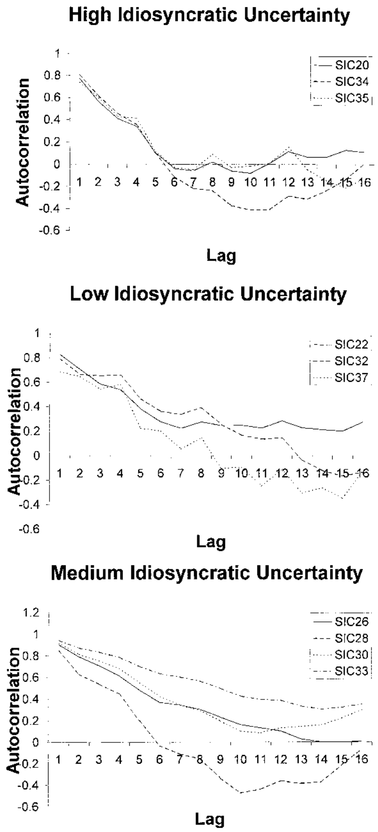

words, the middle group may contain some industries with high or low idiosyncratic uncertainty, which may explain the differences in the correlograms within this group. Finally, because it is less likely that noise would cause an industry to jump from the high- to the low-uncertainty group, we can still maintain that the bulk of our results support models of asynchronous lumpy firm-level investment.

We next examine the statistical significance of our correlograms. Let $\hat{\rho}_i$ be the i^{th}-order autocorrelation of z_t, and let T be the sample size. Then, as shown by Hamilton (1994, p. 111), under the null hypothesis that z_t is Gaussian white noise, the absolute value of $\hat{\rho}_i$ should be less than $2/\sqrt{T}$ approximately 95% of the time. Because we have a sample size of about 100, $2/\sqrt{T} \approx 0.2$. A quick glance at figure 1 reveals that the autocorrelations for the group with high idiosyncratic uncertainty reach this upper confidence bound after four lags. In contrast, in the group with low idiosyncratic uncertainty, the autocorrelations for SICs 22 and 32 are still above this bound after ten lags, and the autocorrelation for SIC 37 falls below this bound at lag seven.

The question remains as to whether the correlograms for the group with high idiosyncratic uncertainty are statistically different from those in the group with low idiosyncratic uncertainty. To address this issue, we use the expression for the approximate variance of an autocovariance given by Hamilton (1994, p. 111). We then perform one-sided t-tests of the null hypothesis that the j^{th} autocovariance for a given industry in the high $\sigma_I/(\sigma_I + \sigma_A)$ group is the same as the j^{th} autocovariance for a given industry in the low $\sigma_I/(\sigma_I + \sigma_A)$ group. Table 4 gives the p-values for these tests, where we examine the first ten autocovariances and where the j^{th} row of the table represents the j^{th} autocovariance. We present the pairwise tests for all possible combinations of industries with high and low idiosyncratic uncertainty. The table is divided into three groups of columns: each group corresponds to one of the industries with high idiosyncratic uncertainty, and each column within the group corresponds to one of the industries with low idiosyncratic uncertainty. We do observe some significant differences between the autocovariances across the high and low groups. For example, starting with the fourth through sixth autocovariances, the autocovariances for SICs 22 and 32 are significantly different at the 5% level from those for the three industries in the high group. However, the autocovariances for SIC 37 are insignificantly different from those for SICs 20 and 35; and only five of the autocovariances for SIC 37 are significantly different from those for SIC 34—four at the 10% level and one at the 5% level. Although not uniform, this evidence provides support for the idea that idiosyncratic uncertainty plays a role in aggregate investment behavior.

FIGURE 2.—ERROR CORRELOGRAMS: TRUNCATED SAMPLE

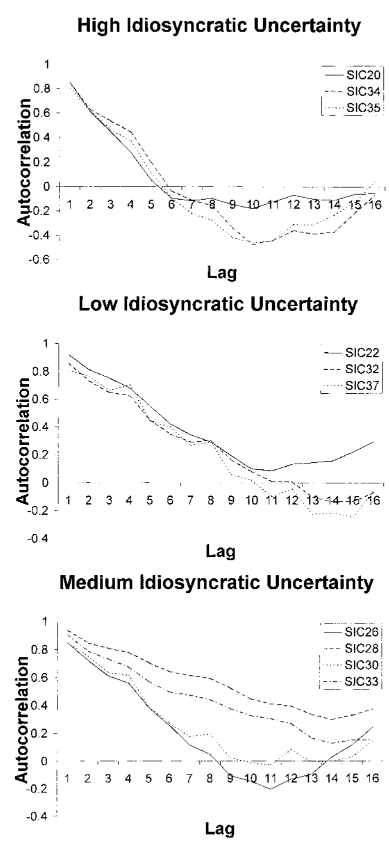

TABLE 4.—*P*-VALUES FOR TESTS OF EQUALITY BETWEEN AUTOCOVARIANCES

	SIC 20			SIC 34			SIC 35		
Order	SIC 22	SIC 32	SIC 37	SIC 22	SIC 32	SIC 37	SIC 22	SIC 32	SIC 37
1	0.354	0.446	0.261	0.440	0.468	0.196	0.279	0.363	0.337
2	0.196	0.276	0.306	0.300	0.395	0.433	0.260	0.350	0.386
3	0.164	0.090	0.224	0.241	0.144	0.316	0.195	0.111	0.261
4	0.143	0.045	0.093	0.167	0.056	0.112	0.244	0.093	0.174
5	0.074	0.034	0.267	0.073	0.034	0.258	0.088	0.041	0.294
6	0.053	0.023	0.102	0.025	0.010	0.051	0.059	0.027	0.113
7	0.074	0.025	0.279	0.013	0.004	0.078	0.083	0.029	0.297
8	0.096	0.034	0.256	0.005	0.001	0.024	0.174	0.070	0.389
9	0.060	0.063	0.400	0.001	0.002	0.083	0.084	0.086	0.341
10	0.047	0.115	0.467	0.001	0.003	0.051	0.091	0.189	0.338

We now examine the robustness of our results to our various data definitions and assumptions. First, we explore whether our use of an ex post measure of uncertainty has led us to any erroneous conclusions. For example, because our uncertainty measures are calculated over the full sample, it is possible that a series of high-variance realizations of the cross-sectional distribution of shocks could drive our results. First, a high-variance realization would cause substantial cross-sectional reshuffling and lower the persistence of z_i. However, it might also lead to a high observed in-sample σ_I. In this example, therefore, it is not our maintained hypothesis that is driving the relationship between persistence in z_i and the uncertainty ratio, but the incidental realization of shocks.

To rule out such a possibility, we calculate the uncertainty ratios using the first twenty observations and then run our cointegrating regressions using the subsequent 83 observations. The uncertainty ratio statistics for this experiment are in table 5, which replicates the sorting of industries observed in table 2. Given this identical sorting of industries into high and low $\sigma_I/(\sigma_I + \sigma_A)$ groups, it is not surprising that the results from our new set of cointegrating regressions in table 6 and the correlograms in figure 2 display the same pattern of results as those observed when we used our in-sample uncertainty ratios. Table 7 gives the *p*-values for the tests of the differences between the autocovariance for the high and low groups. Here, the evidence for differences between the autocovariances is stronger than the evidence in table 4. In particular, the autocovariances for SIC 37 are significantly different from those for all of the industries in the low group.

Next, we turn to the question of whether our use of quarterly data makes a difference. We do so by estimating the cointegrating regressions and calculating the correlograms with the annual two-digit SIC data used by Doyle (1993). Not surprisingly, because of the problem of time aggregation of individual-firm investment decisions, we have found almost no differences across industry groups. We conclude that our efforts to construct quarterly data have in fact been fruitful.

In a similar vein, we have also examined whether our use of seasonally unadjusted data is warranted. This issue is of particular interest because the seasonally unadjusted data can be much more volatile than the adjusted data. Because we are looking at residuals from a regression containing a fairly smooth capital stock series and a much more volatile output series, our results may be an artifact of seasonal patterns in the residuals. To address this difficulty, we have also run our regressions using seasonally adjusted series, with little change in any of the results. We conjecture that

TABLE 5.—UNCERTAINTY RATIO SUMMARY STATISTICS: 1967:2–1972:1

SIC	Industry	Firms	σ_I	σ_A	$\sigma_I/(\sigma_I + \sigma_A)$
20	Food	104–114	0.189	0.117	0.619
22	Textiles	43–60	0.170	0.156	0.521
26	Paper	41–47	0.146	0.123	0.542
28	Chemicals	124–133	0.153	0.111	0.580
30	Rubber	40–47	0.181	0.151	0.547
32	Stone, clay, glass	43–51	0.145	0.135	0.518
33	Primary metals	84–102	0.157	0.124	0.559
34	Fabricated metals	79–90	0.215	0.149	0.591
35	Industrial machinery	143–158	0.197	0.148	0.571
37	Transportation equipment	90–110	0.179	0.154	0.537

TABLE 6.—COINTEGRATING REGRESSIONS: 1972:2–1992:4

	SIC	α_1	R^2	σ_z
High $\sigma_I/(\sigma_I + \sigma_A)$	20	−0.526 (0.079)	0.361	0.006
	34	−0.506 (0.075)	0.360	0.008
	35	−0.446 (0.099)	0.326	0.011
Low $\sigma_I/(\sigma_I + \sigma_A)$	22	−0.940 (0.148)	0.309	0.014
	32	−0.764 (0.103)	0.396	0.017
	37	−1.088 (0.128)	0.493	0.025
Intermediate $\sigma_I/(\sigma_I + \sigma_A)$	26	−1.126 (0.105)	0.585	0.020
	28	−0.546 (0.076)	0.338	0.013
	30	−0.342 (0.159)	0.044	0.023
	33	−1.210 (0.174)	0.338	0.071

Standard errors are in parentheses under the parameter estimates.

TABLE 7.—*P*-VALUES FOR TESTS OF EQUALITY BETWEEN AUTOCOVARIANCES

	SIC 20			SIC 34			SIC 35		
Order	SIC 22	SIC 32	SIC 37	SIC 22	SIC 32	SIC 37	SIC 22	SIC 32	SIC 37
1	0.319	0.479	0.377	0.306	0.470	0.382	0.207	0.356	0.500
2	0.119	0.250	0.190	0.133	0.277	0.211	0.115	0.249	0.187
3	0.052	0.151	0.131	0.115	0.275	0.247	0.056	0.165	0.143
4	0.021	0.044	0.017	0.108	0.181	0.092	0.050	0.095	0.042
5	0.008	0.030	0.031	0.040	0.111	0.114	0.014	0.049	0.050
6	0.007	0.019	0.010	0.014	0.035	0.020	0.006	0.018	0.009
7	0.015	0.031	0.038	0.015	0.030	0.037	0.003	0.008	0.010
8	0.034	0.034	0.036	0.017	0.017	0.018	0.003	0.004	0.004
9	0.054	0.078	0.172	0.006	0.011	0.033	0.002	0.004	0.013
10	0.096	0.119	0.180	0.004	0.006	0.011	0.004	0.007	0.013

this similarity is due to the small observed multiplicative adjustment factors on shipments data.

At this point, it is important to ask whether our results depend on our ceteris paribus assumption. As explained by Bertola and Caballero (1990), not only does the variance ratio matter for the dynamics of the aggregate, but so does the drift in the z_t process and adjustment costs. First, if the drift in z_t is high (for example, because of a high depreciation rate), firms will hit investment triggers frequently, and persistence in z_t will be low. To examine whether differences in drift are driving our results, we use two approaches. The first stems from the observation by Bertola and Caballero (1994) that high drift and frequent adjustment imply high investment rates. When we compare investment rates across industries, we reassuringly find no systematic relationship between these figures and $\sigma_I/(\sigma_I + \sigma_A)$. For example, the investment rates for the low $\sigma_I/(\sigma_I + \sigma_A)$ group are 0.35, 0.32, and 0.26, whereas for the high $\sigma_I/(\sigma_I + \sigma_A)$ group they are 0.31, 0.27, and 0.33.

Second, we use an alternative measure of drift that is a function of stock returns adjusted for depreciation. Recall from our discussion in section III that, under certain assumptions, returns are proportional to the percentage change in the marginal product. Under our assumption of CES technology, by multiplying the average return on an equally weighted portfolio of stocks by our estimates of $-1/\alpha_1$ for a particular industry, we can obtain a measure of drift. We add depreciation rates to these measures, because firms buy new capital goods not only for expansion but for maintenance. Once again, when we compare these measures of drift across industries, we find little relationship between drift and the uncertainty ratio. The drifts for the low $\sigma_I/(\sigma_I + \sigma_A)$ group are 0.11, 0.11, and 0.09, whereas for the high $\sigma_I/(\sigma_I + \sigma_A)$ group they are 0.10, 0.12, and 0.12.

A more serious concern is whether our results are an artifact of an incidental correlation between adjustment costs and $\sigma_I/(\sigma_I + \sigma_A)$. This possibility arises because high adjustment costs imply wide inaction intervals, infrequent adjustment, and high persistence in z_t. However, we find this interpretation unlikely for three reasons. First, Doms and Dunne (1998) report no significant differences in the bunching of plant-level investment across four-digit subindustries contained within our high and low $\sigma_I/(\sigma_I + \sigma_A)$ groups. If adjustment costs were driving our results, we would expect to see differences, but we don't. Second, on an intuitive level, we would expect differences in adjustment costs to explain our results only if the types of investment differed across industries. However, industry investment consists of the aggregation of heterogeneous projects, where many of these projects share important characteristic across industries.[9] Finally, Veracierto (1996) shows that, in a general equilibrium context, changing the fixed cost of disinvesting (irreversibility) does not matter greatly for aggregate investment. The intuition is that, as irreversibility increases, the number of plants investing decreases, but the amount each plant invests increases.

Finally, as noted in the introduction, it is possible that a model with convex costs of adjustment could produce our results. We therefore examine the question of whether it is, in fact, an (S,s) type of micro model with fixed adjustment costs or a micro model with convex costs that is driving our results. We can use an insight from the fixed-costs literature to shed light on this question. Specifically, not only does productivity affect investment, but investment affects productivity, because a firm that is investing will tend to suffer lost output. In a convex costs model, such as those underlying q-theory, the former effect dominates because firms smooth investment over time and do not incur large output losses in the period in which investment occurs. In a fixed-costs model, on the other hand, firms will tend to invest in spikes and incur large output losses during these episodes. In industries characterized by high synchronization of actions, therefore, if fixed costs are important, we should tend to see a low or even negative correlation between investment and productivity. In contrast, in the case of convex costs, we should see a large positive correlation.

To distinguish between these two cases, we first calculate the industry-wide shocks to the marginal product of capital using the fitted values from the regression (3). We then regress the log first difference of the capital stock on this

[9] Caballero and Engel (1991) analyze the case of heterogeneous agents, showing that, in this case as well, the smoothness of an aggregate variable depends on the coordination of individual decisions.

TABLE 8.—REGRESSION OF INVESTMENT ON PRODUCTIVITY

SIC	Industry	Current Productivity	Lagged Productivity
20	Food	0.0043	−0.0070
		(0.0049)	(0.0049)
22	Textiles	−0.0059	0.0095
		(0.0046)	(0.0047)
26	Paper	−0.0431	0.0404
		(0.0240)	(0.0233)
28	Chemicals	−0.0275	0.0251
		(0.0049)	(0.0046)
30	Rubber	0.0007	0.0037
		(0.0028)	(0.0030)
32	Stone, clay, glass	−0.0016	0.0096
		(0.0057)	(0.0057)
33	Primary metals	−0.0151	0.0146
		(0.0067)	(0.0066)
34	Fabricated metals	−0.0048	0.0178
		(0.0033)	(0.0034)
35	Industrial machinery	−0.0184	0.0213
		(0.0036)	(0.0036)
37	Transportation equipment	−0.0115	0.0182
		(0.0039)	(0.0040)

Standard errors are in parentheses under the parameter estimates.

measure of productivity. This simple regression produces nonstationary residuals—a phenomenon known as the "spurious regression" problem. As a remedy, we therefore include lagged productivity in the regression, with the result that a Dickey-Fuller test indicates that the residuals are stationary. The results from this regression can be found in table 8, in which standard errors are in parentheses under the parameter estimates. First, it is worth noting that eight of the ten coefficients on lagged productivity are positive and significant—two at the 10% level and six at the 5% level. This result is consistent with delivery lags for capital goods or with a time to build of more than one quarter. More importantly, however, the first column of the table shows that eight of the ten coefficients on current productivity are negative, and five of these are significantly different from zero at the 5% level. None of the coefficients is significantly positive. Given our preceding argument, this evidence points to a lumpy rather than a smooth adjustment model.

This interpretation hinges on adjustment costs taking the form of foregone output: either resources will be diverted away from productive activities when capital goods are installed, or the firm will have to learn about the new production technology and therefore be less productive during this process. However, adjustment costs could quite possibly be pecuniary, as in the case of a firm that outsources the installation of new computer equipment, and our measure of productivity does not account for pecuniary adjustment costs. Nonetheless, we believe our results are robust to this alternative. For example, if all adjustment costs were pecuniary, a firm would not suffer output losses when investing, and we would therefore not expect to see a negative correlation between investment and our measure of contemporaneous productivity, yet we do.

A different problem with our interpretation of table 8 arises because of the possibility of decreasing returns to scale, which could induce a negative correlation between investment and the marginal product of capital in the absence of any adjustment costs. However, investment goods must be delivered and installed within the period for this problem to arise. Because our observed positive coefficients on lagged productivity are consistent with delivery and installation lags, we view as small the likelihood that decreasing returns are responsible for our negative coefficients. In sum, because we do observe output losses associated with investment and because we suspect strongly that these output losses are not a product of decreasing returns, we view the evidence in table 8 as consistent with the presence of fixed costs.

V. Conclusion

We provide new evidence against representative-agent models, which have been the workhorse of much of modern macroeconomics. We find that the larger the dispersion of financial shocks to firms in an industry, the less sluggish is the adjustment of their aggregated capital stocks. Given the infrequent capital stock adjustment observed at the plant level, this result is clearly at odds with the idea that the individual and the aggregate can be modeled identically. Instead, it is consistent with models that explain smooth aggregate fluctuations as the result of the aggregation of microeconomic decisions. Our approach to this issue has the advantage that our results do not hinge on the assumptions of a particular model. On the other hand, because we do not explicitly test a specific economic model, we have taken care to ensure that our results are not the product of alternative explanations. Further research needs to find new methods for direct tests of models of microeconomic adjustment.

REFERENCES

Abel, Andrew, and Olivier Blanchard, "Investment and Sales: Some Empirical Evidence," (pp. 269–296), in William A. Barnett, Ernst R. Berndt, and Halbert White, (Eds.), *Dynamic Econometric Modeling: Proceedings of the Third International Symposium in Economic Theory and Econometrics* (Cambridge: Cambridge University Press, 1988).

Auerbach, Alan, and Kevin Hassett, "Tax Policy and Business Fixed Investment in the United States," *Journal of Public Economics* 47:2 (1992), 141–170.

Bertola, Giuseppe, and Ricardo J. Caballero, "Kinked Adjustment Costs and Aggregate Dynamics," *N.B.E.R. Macroeconomics Annual* (1990), 237–288.

——— "Irreversibility and Aggregate Investment," *Review of Economic Studies* 61:2 (1994), 223–246.

Blinder, Alan S., "Retail Inventory Behavior and Business Fluctuations," *Brookings Papers on Economic Activity,* 2 (1981), 443–505.

Caballero, Ricardo J., "Small Sample Bias and Adjustment Costs," this REVIEW 76:1 (1994), 52–58.

Caballero, Ricardo J., and Eduardo M. R. A. Engel, "Dynamic (S,s) Economies," *Econometrica* 59:6 (1991), 1659–1686.

——— "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach," *Econometrica* 67:4 (1999), 783–826.

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

Caballero, Ricardo J., and John Leahy, "Fixed Costs: The Demise of Marginal *q*," NBER working paper no. 5508, (1996).

Caballero, Ricardo J., and Robert Pindyck, "Uncertainty, Investment, and Industry Evolution," *International Economic Review* 37:3 (1996), 641–662.

Caplin, Andrew S., "The Variability of Aggregate Demand with (S,s) Inventory Policies," *Econometrica* 53:6 (1985), 1395–1409.

Caplin, Andrew, and Daniel F. Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics* 102:4 (1987), 703–725.

Chirinko, Robert S., "Business Fixed Investment Spending: Modeling Strategies, Empirical Results, and Policy Implications," *Journal of Economic Literature* 31:4 (1993), 1875–1911.

Cooper, Russell, and John Haltiwanger, "On the Nature of Capital Adjustment Costs," Boston University mimeograph (1998).

Cooper, Russell, John Haltiwanger, and Laura Power, "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review* 89:3 (1999), 921–946.

Doms, Mark, and Timothy Dunne, "Capital Adjustment Patterns in Manufacturing Plants," *The Review of Economic Dynamics* 1:2 (1998), 409–429.

Doyle, Joanne M., "An Empirical Investigation into Internal Adjustment Costs in Neoclassical Investment Models," unpublished doctoral dissertation, Boston College (1993).

Foote, Christopher, Erik Hurst, and John Leahy, "Testing the (S,s) Model," *American Economic Review* 90:2 (2000), 116–119.

Greene, William H., *Econometric Analysis* (Upper Saddle River, NJ, Prentice-Hall, 1997).

Hamilton, James D., *Time Series Analysis* (Princeton, Princeton University Press, 1994).

Johansen, Soren, "Statistical Analysis of Cointegration Vectors," *Journal of Economic Dynamics and Control* 12:2/3 (1988), 231–254.

Leahy, John V., and Toni M. Whited, "The Effect of Uncertainty on Investment: Some Stylized Facts," *Journal of Money, Credit, and Banking* 28:1 (1996), 64–83.

Lewbel, Arthur, "Aggregation with Log-Linear Models," *Review of Economic Studies* 59:3 (1992), 635–642.

Reagan, Patricia, and Dennis Sheehan, "The Stylized Facts About the Behavior of Manufacturers' Inventories and Back-Orders Over the Business Cycles: 1959–1980," *Journal of Monetary Economics* 15:2 (1985), 217–246.

Stock, James H., and Mark W. Watson, "A Simple MLE of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica* 61:4 (1993), 783–820.

Veracierto, Marcelo, "Plant Level Irreversible Investment and Equilibrium Business Cycles," Cornell University mimeograph (1996).

West, Kenneth D., "A Note on the Econometric Use of Constant Dollar Inventory Series," *Economic Letters* 13 (1983), 337–341.

Whited, Toni M., "Debt, Liquidity Constraints, and Corporate Investment: Evidence from Panel Data," *Journal of Finance* 47 (1992), 1425–1460.