MEASURING FACTOR ADJUSTMENT COSTS*

ROBERT E. HALL

I estimate adjustment costs for labor and capital from the Euler equations for factor demand. For both factors, I find relatively strong evidence against substantial adjustment costs. My estimates use annual data from two-digit industries. My results support the view that rents arising from adjustment costs are relatively small and are not an important part of the explanation of the large movements of the values of corporations in relation to the reproduction costs of their capital. I investigate the potential effects of three types of specification error: (1) aggregation over time, (2) aggregation over firms with heterogeneous demand shocks, and (3) estimation of a convex adjustment-cost technology in the presence of nonconvex discrete adjustment costs. I find that the likely biases from these specification errors are relatively small.

I. INTRODUCTION

The notion permeates economics that input factors are costly to adjust. Students learn the distinction between a firm's short-run supply function, with capital held constant, and its long-run supply function, applicable after capital adjusts slowly. Labor may be costly to adjust as well. Costs of adjustment imply that firms earn rents when demand rises unexpectedly, until they and their rivals can increase capital and labor. Some economists believe that these rents are large enough to account for a substantial part of the movements of stock prices relative to the underlying cost of the capital that firms own. I find, on the contrary, that adjustment costs are small for both factors. The supply functions implicit in this work are quite elastic over periods of a year or more. Rents are small and transitory. The results support the position I advocated in Hall [2001], that large movements in the stock market are not the capitalization of rents, but reflect other forces, such as the accumulation of large stocks of intangibles.

The focus of my research is on rents arising from adjustment costs that might persist for a year or more, in line with the persistence of movements of securities values in relation to the

* This research was supported by the National Science Foundation under grant SOC SBR-9730341 and is part of the research program on Economic Fluctuations and Growth of the NBER. I am grateful to numerous seminar participants for comments, to Anthony Fai Chung and Felix Reichling for pointing out a number of errors, and to the editor and referees for many suggestions for improvement. A file containing all data, programs, and estimation results is available from Stanford.edu/~rehall.

The Quarterly Journal of Economics, August 2004

reproduction cost of the capital stock. For this reason, I use annual data at the two-digit level of aggregation. I show that both time aggregation and aggregation across firms is probably not an important source of bias in estimation. I fit a model of convex adjustment costs, which neglects the role of nonconvex discrete costs. I show that discrete costs have relatively little role in rents at the industry level even though discrete costs are central to understanding the pattern of investment at the plant level. In particular, an inference of rents from an annual industry model without discrete costs provides a reasonable approximation to the aggregated rent from a fully disaggregated model with discrete adjustment cost.

Many economists believe that adjustment costs for labor are unimportant from one year to the next, so my finding of low annual adjustment costs is no surprise. I do not dwell on the labor adjustment issue in the paper, but I view the finding of low annual adjustment cost for labor as a validation of the method I use.

The paper uses industry panels of annual data for factor inputs over the period from 1948 through 2001. I study the movements of the ratio of labor cost to materials cost and the movements of the ratio of capital cost to labor cost. Absent adjustment costs and absent changes in the ratios of factor prices, an increase in demand or in another determinant of industry equilibrium would cause factor inputs to change in the same proportion as output, and the cost ratios would remain constant. If labor has adjustment costs, it adjusts less than materials do, and the cost ratio moves correspondingly. Similarly, if capital has adjustment costs and labor does not, the movements of their cost ratio provide information about those adjustment costs. A structural equation derived from value maximization captures this property. To identify two key adjustment-cost parameters, I use standard time-series instruments—the timing of oil-price shock and the level of military spending—to identify exogenous movements in each industry's use of factors.

I use data on actual factor prices to incorporate the effects of changes in factor prices on factor intensities. In principle, I distinguish adjustment costs internal to the firm from those that arise outside the firm. In the case of labor, external adjustment costs include costs of moving and costs of acquiring non-firm-specific skills needed for new jobs. In the case of capital, external adjustment costs include those incurred by capital-goods suppliers. These external adjustment costs are transmitted to the firms in my sample by factor prices. My success in distinguishing internal from external adjustment costs depends on the accuracy of the factor prices. The data measure an industry-specific wage, so that I should be able to distinguish adjustment costs from the short-run inelasticity of labor supply to an industry owing to labor mobility costs. For capital, I use an industry-specific measure of the rental price of capital goods.

The adjustment-cost parameter that I estimate is the reciprocal of the adjustment response considered in Tobin's q-theory of investment. Tobin's q is the ratio of the market or shadow value of installed capital to its acquisition cost from capital goods producers. The adjustment response is the change in the flow of investment stated as a proportion to the capital stock induced by a change in q. A typical value for the adjustment response in the q literature is 0.05 at annual rates, corresponding to an annual adjustment-cost parameter of 20. My estimates of the adjustment-cost parameters are not much above zero for either labor or capital in most industries. I find lower adjustment costs for both labor and capital than does Shapiro [1986], an earlier study with much the same framework as this one, but based on more aggregate data.

II. TECHNOLOGY

The firm uses labor input n_t, materials input m_t, and capital input k_t to produce output y_t, according to the technology:

$$y_t = A_t n_t^{\alpha} m_t^{\psi} k_t^{1-\alpha-\psi} - \frac{\lambda_t}{2} n_t g_{n,t+1}^2 - \frac{\gamma_t}{2} k_t g_{k,t+1}^2. \tag{1}$$

Here $g_{n,t} = (n_{t+1} - n_t)/n_t$, and $g_{k,t} = [k_{t+1} - (1 - \omega\delta)k_t]/k_t$. A_t is an index of productivity, growing over time at a possibly variable rate. The second and third terms embody the factor adjustment technology. The parameter δ is the rate of deterioration of capital. The parameter ω determines whether adjustment costs apply to replacement investment. If $\omega = 0$, adjustment costs arise only from net investment—the change in the capital stock. If $\omega = 1$, replacement investment incurs adjustment costs as well. The adjustment costs are convex (quadratic) in the inputs and have constant returns to scale as discussed in Hall [2001]. Changing the level of labor or capital results in a loss of output: these are recruiting and training costs for labor and planning and installation costs for capital. Notice that I do not include discrete costs of adjustment. I investigate discrete costs later in the paper.

The firm maximizes the present value of its expected cash flows. The first-order conditions are

$$p_t\left[\lambda_t E_t\left(g_{n,t+1} + \frac{1}{2}g^2_{n,t+1}\right) + \alpha A_t n_t^{\alpha-1} m_t^{\psi} k_t^{1-\alpha-\psi}\right] - \lambda_{t-1}(1 + r_{t-1})p_{t-1}g_{n,t} - w_t = 0, \tag{2}$$

$$p_t\left[\gamma_t E_t\left(g_{k,t+1} + \frac{1}{2}g^2_{k,t+1}\right) + (1 - \alpha - \psi)A_t n_t^{\alpha} m_t^{\psi} k_t^{-\alpha-\psi}\right] - \gamma_{t-1}(1 + r_{t-1})p_{t-1}g_{k,t} - \rho_t = 0, \tag{3}$$

$$\psi p_t A_t n_t^{\alpha} m_t^{\psi-1} k_t^{1-\alpha-\psi} - p_{m,t} = 0. \tag{4}$$

Here p is the product price, w is the wage, ρ is the rental price of capital, and p_m is the price of materials inputs.

I estimate the adjustment-cost parameters by applying instrumental variables to equations derived from equations (2), (3), and (4). At the end of the paper, I will argue that this approach to estimation is robust to a number of issues: (i) I omit discrete adjustment costs; (ii) I estimate with annual data even though decisions about factor inputs are likely to be made more frequently; (iii) the theory is ambiguous about the role of adjustment costs for investment that replaces deteriorated capital; and (iv) the industry data aggregate many firms each with its own firm-specific shock.

III. Econometric Framework

III.A. *Why Estimate the Euler Equation?*

When the technology or preferences are time-separable, the econometrician faces a fundamental choice in estimation strategy. Under weak assumptions, key deep parameters such as the adjustment-cost parameter can be estimated from an Euler equation by instrumental variables. More efficient estimators can be constructed by assigning probability distributions to the random elements and dealing with the solution to the full decision problem. Hansen [2002] has described the distinction nicely:

> GMM approaches based on Euler equations were designed to deliver only part of an economic model. Their virtue and liability is that they are based on partial specification of an econometric model. They allow an applied researcher to study a piece of a full dynamic model, without getting hung up on the details of the remainder of the model. . . . It allows an econometrician to learn about *something* without needing to learn about *everything*.

In the context of adjustment costs, an econometrician might consider building a structural model of investment. The structural approach would require the estimation of a time-series model for the driving forces and the derivation of the optimal response of investment to current values of the driving forces. Econometric identification of the resulting model would be a challenge, especially because it is likely that the decision makers observe more information than is contained in the variables observed by the econometrician. Strong and implausible assumptions would be required. The most convenient econometric framework for this approach is probably indirect inference, where the point of contact of the data and the model is not the entire body of data (as a maximum likelihood approach would require) but an adroitly chosen summary of the data. Cooper and Haltiwanger [2002] pursue this strategy.

If adequate precision is available from an IV estimator applied to the Euler equation, then the burden of the implausible assumptions can be avoided. A sensible procedure is to try the IV approach and accept its results if the precision is reasonable and to estimate the parameters from an investment equation otherwise. I believe that the second step is not required for the data I use.

III.B. Specification

To estimate the adjustment-cost parameter for labor, I use a version of the Euler equation obtained by dividing the first-order condition for labor by the one for materials and rearranging as follows:

$$\frac{p_{m,t}m_t}{w_t n_t} = \frac{\psi}{\alpha}\left\{1 + \lambda_t\left[(1 + r_{t-1})\frac{p_{t-1}}{w_t}g_{n,t} - E_t\frac{p_t}{w_t}\left(g_{n,t+1} + \frac{1}{2}g^2_{n,t+1}\right)\right]\right\}. \tag{5}$$

The data reveal slow movements in the factor-cost ratio on the left side of this equation that cannot plausibly be the result of adjustment costs. I hypothesize a geometric random walk for the parameter ratio ψ/α. Let its innovation be $\tilde{\eta}_t$. I also assume that the cost of adjustment rises along with the wage/product price ratio: $\lambda_t = \lambda w_t/p_t$. Then

$$(6)\quad \frac{\alpha}{\psi}\frac{p_{m,t}m_t}{w_t n_t} = 1 + \lambda\left[(1 + r_{t-1})\frac{w_{t-1}}{w_t}g_{n,t} - E_t\left(g_{n,t+1} + \frac{1}{2}g_{n,t+1}^2\right)\right]$$

and

$$(7)\quad \Delta \log \frac{p_{m,t}m_t}{w_t n_t} = \tilde{\eta}_{n,t} + \Delta \log \left\{1 + \lambda\left[(1 + r_{t-1})\frac{w_{t-1}}{w_t}g_{n,t} - E_t\left(g_{n,t+1} + \frac{1}{2}g_{n,t+1}^2\right)\right]\right\}.$$

I approximate the right side as

$$(8)\quad \Delta \log \frac{p_{m,t}m_t}{w_t n_t} = \tilde{\eta}_{n,t} + \lambda\Delta\left[(1 + r_{t-1})\frac{p_{t-1}}{w_t}g_{n,t} - E_t\frac{p_t}{w_t}\left(g_{n,t+1} + \frac{1}{2}g_{n,t+1}^2\right)\right].$$

Finally, I substitute the realization less the disturbance for the expectation:

$$(9)\quad \Delta \log \frac{p_{m,t}m_t}{w_t n_t} = \lambda\Delta\left[(1 + r_{t-1})\frac{p_{t-1}}{w_t}g_{n,t} - \frac{p_t}{w_t}\left(g_{n,t+1} + \frac{1}{2}g_{n,t+1}^2\right)\right] + \eta_{n,t}.$$

The disturbance $\eta_{n,t}$ incorporates both the innovation in the parameter ratio and the expectation error.

My estimate of the labor adjustment-cost parameter λ is quite precisely zero—labor is a freely adjustable factor like materials. For estimating the capital adjustment-cost parameter, I use labor cost in place of materials cost. I assume that the cost of adjustment changes in proportion to ratio of the capital goods price $p_{k,t}$ and the output price ($\gamma_t = \gamma p_{k,t}/p_t$) and estimate

$$(10)\quad \Delta \log \frac{w_t n_t}{\rho_t k_t} = \gamma\Delta\left[(1 + r_{t-1})\frac{p_{k,t-1}}{\rho_t}g_{k,t} - \frac{p_{k,t}}{\rho_t}\left(g_{k,t+1} + \frac{1}{2}g_{k,t+1}^2\right)\right] + \eta_{k,t}.$$

III.C. Instrumental Variables and Identification

I estimate the adjustment-cost parameters under the identifying assumption that two aggregate instrumental variables—military spending and the timing of oil price shocks—are strongly exogenous variables. By strongly exogenous, I mean not only uncorrelated with the expectation errors in the Euler equations, but also uncorrelated with the innovation in the ratio of the share parameters. Consequently, I do not rely on pure timing considerations in my choice of instruments—this is a key difference between my approach and Shapiro's [1986]. His use of lagged endogenous variables is subject to the criticism of Garber and King [1983], although he avoids a leading source of the bias they identify by removing the productivity shock from the unobserved disturbance. A second reason for avoiding lagged endogenous variables as instruments is that, in the case of small adjustment costs, the expectation error in the Euler equation is close to zero and the bulk of the noise in the equation must come from another source, for which identification based on timing has no obvious support.[1]

Specifically, I use as instruments (1) the univariate AR(1) innovation in the chain-type quantity index for federal defense spending from the U. S. National Income and Product Accounts and (2) a dummy variable taking the value one in the oil-shock years 1956, 1974, 1979, and 1990, and zero in every other year.

To achieve orthogonality of the instruments with the expectation errors, the instruments must be dated at or before the time the expectations are formed. The decision variables n_t and k_t are actually time aggregates influenced by events occurring in year t. As a result, the expectation errors are not orthogonal to exogenous variables measured during that year. The latest eligible instrument is dated $t - 1$. In the case of capital, a further consideration affects the timing of the instrument. Because of planning and installation lags (time to build), the decision about k_t is most likely to be made in year $t - 1$, in which case the latest eligible instrument is from year $t - 2$.

III.D. Covariances of Disturbances across Industries

Because the disturbances are somewhat correlated across industries, it might be desirable to reestimate with an estimated

1. I am grateful to a referee for pointing this out.

covariance matrix. However, the actual covariance matrix across industries is extremely close to singular, as there are only 49 observations to estimate the 18×18 covariance matrix for the estimation of labor adjustment costs and 51 observations to estimate the 56×56 covariance matrix for the estimation of capital adjustment costs. Some experimentation suggested that better results follow from treating the covariance matrix as diagonal rather than using its actual, near-singular value in the labor case. In particular, the standard errors of the estimates of the parameters were conspicuously too small with the three-stage least squares estimator.

IV. Data

To estimate the labor adjustment-cost parameter, I use the total factor productivity data compiled by the Bureau of Labor Statistics. For each of eighteen manufacturing industries, the data report the value and quantity of output, labor, capital, materials, and energy, annually from 1949 to 2000. Because this body of data includes materials inputs, it suits the estimation strategy described above.[2]

The BLS data combine inventories and fixed investment. To estimate the adjustment-cost parameter for fixed investment alone and to take advantage of more detailed and extensive industry coverage, I use data from the Fixed Asset Tables associated with the National Income and Product Accounts (NIPAs). This source reports the current-dollar capital stock—the quantity of capital multiplied by the price of new capital goods—and an index of the quantity of capital. I multiply the current-dollar capital stock by the rental price of capital to obtain the current-dollar value of capital services, ρk. I measure investment as the rate of change in the quantity of capital, corresponding to the specification for net investment, $\omega = 0$. I use data on total employee compensation from the NIPAs as the numerator in the ratio on the left side of the estimating equation.

The rental price is the cost of holding one unit of plant and equipment for one year. I calculate it from

2. An earlier version of the paper used related data developed by Jorgenson [1995, chapter 1] with essentially identical results. Because the BLS data cover more years and the NIPA data have more industry detail and more years, and because the government data are more completely documented, I have presented results based on the government data here.

TABLE I
ELEMENTS OF THE RENTAL PRICE OF CAPITAL

Symbol	Interpretation	Source
τ	Corporate marginal tax rate	Ratio of corporate tax to taxable income of corporations, NIPA Table 8.25 (all references are to NIPA tables prior to the December 2003 revisions).
z	Present value of depreciation deductions for tax purposes	Nominal depreciation rate inferred from total investment in plant and equipment, NIPA Table 1.1, and depreciation deductions, NIPA Table 8.22. Present value computed assuming exponential time pattern and nominal discount rate equal to the 6-month Treasury bill rate plus a risk premium of 2 percent.
c	Investment tax credit	Ratio of tax credits taken, NIPA Table 8.25, to investment, NIPA Table 1.1.
r	After-tax financial cost of capital	Taken as 5 percent—see Hall [2003].
$\delta_{i,t}$	Rate of deterioration of capital	For each industry, ratio of depreciation in current dollars, NIPA Fixed Asset Table 3.4ES, to current-dollar capital stock, Fixed Asset Table 3.1ES.
$p_{k,i,t}$	Price of capital goods	Implicit in capital stock data, Fixed Asset Table 3.1ES.

$$\rho_{i,t} = \frac{(1 - \tau_t z_t - c_t)(r + \delta_{i,t})}{1 - \tau_t} p_{k,i,t}. \tag{11}$$

The inputs and sources for this calculation are shown in Table I.

The coefficient of the autoregression of the military spending variable is 0.793 with a standard error of 0.054. The residuals from this regression form one of the two instruments.

V. ESTIMATES

V.A. Cost of Adjusting Labor

Table II shows the IV estimates of the labor adjustment cost λ for each industry. Figure I shows the basics of the finding that

TABLE II
ESTIMATES OF THE ADJUSTMENT-COST PARAMETER FOR LABOR INPUT

Industry	Estimate, $\hat{\lambda}$	Standard error
Food & kindred prod. (SIC 20)	0.017	(1.72)
Textile mills prod. (SIC 22)	−0.045	(0.51)
Apparel & related prod. (SIC 23)	−0.158	(0.38)
Paper & allied prod. (SIC 26)	−0.179	(0.54)
Printing & publishing (SIC 27)	−0.965	(4.15)
Chem. & allied prod. (SIC 28)	−0.593	(2.09)
Petroleum refining (SIC 29)	−0.390	(0.60)
Rubber & plastic prod. (SIC 30)	0.373	(1.43)
Lumber & wood prod. (SIC 24)	0.075	(0.22)
Furniture & fixtures (SIC 25)	−0.085	(0.17)
Stone, clay & glass (SIC 32)	−0.894	(2.42)
Primary metal ind. (SIC 33)	−1.663	(7.51)
Fabricated metal prod. (SIC 34)	−0.337	(1.46)
Ind. machinery, comp. eq. (SIC 35)	0.185	(0.34)
Electric & electr. eq. (SIC 36)	1.254	(9.01)
Transportation equip. (SIC 37)	7.695	(115.50)
Instruments (SIC 38)	0.810	(2.77)
Misc. manufacturing (SIC 39)	0.012	(0.53)

Estimates and standard errors of the parameter λ from equation (9), using oil-shock dummies and the innovation in military spending as instruments, with data from the BLS multifactor productivity study, over the period 1949 through 2000.

the adjustment-cost parameter λ is close to zero in those industries where it is measured with precision. Each diamond in the plot describes the IV estimate of λ in an industry. The horizontal position of the diamond shows the estimate of the parameter. The vertical position shows the precision of that estimate (precision is the reciprocal of the standard error of the estimated coefficient). The individual estimates—especially those with high precision—are mostly very close to zero. The fine curving line shows the precision level such that the coefficient would differ from zero by one standard error (that is, the reciprocal of the corresponding coefficient value). None of the estimates differ from zero by a standard error or more. The calculated standard errors overstate the actual dispersion of the estimates, else at least a third of the estimates would be more than one standard error from zero. The reason is that the estimates are moderately correlated with one another—across all industry pairs, the median correlation is 0.20, the twenty-fifth percentile is 0.02, and the seventy-fifth percentile is 0.45.

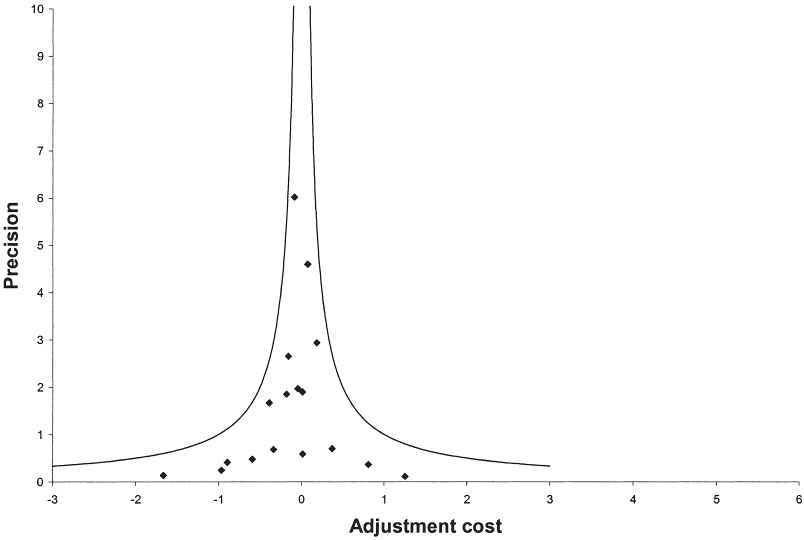

FIGURE I
Instrumental Estimates of Labor Adjustment Cost by Industry

The horizontal position of the marker shows the estimate of the labor adjustment cost λ, for one industry, and the vertical position measures the precision (the reciprocal of the standard error). The fine curving lines mark a precision equal to the horizontal position, corresponding to an estimate equal to its standard error.

To see why the data reveal low adjustment costs for labor, write equation (9) in the following form:

$$\Delta \log \frac{m_t}{n_t} = \Delta \log \frac{w_t}{p_{m,t}} + \lambda \hat{g}_{n,t} + \hat{\eta}_{n,t}. \tag{12}$$

The finding of an IV estimate of λ of zero means that, when an exogenous force drives up the employment growth variable $\hat{g}_{n,t}$, the force moves the materials/employment ratio m_t/n_t only by the amount mandated by the change in the relative factor price, $w_t/p_{m,t}$. By contrast, if adjustment costs held back the movement of labor, then an expansionary force would increase materials more than labor, m_t/n_t would rise by more than the amount mandated by factor prices, and $\lambda \hat{g}_{n,t}$ would account for the extra movement of m_t/n_t. Adjustment costs result in changes in factor intensities not explained by movements of relative factor prices.

Notice that some of the estimates of the adjustment-cost parameter are negative, though none by as much as one standard error. The true value of the parameter cannot be negative. Most of the negative values probably arise from pure sampling errors, though nothing rules out some role for specification errors. The situation is much the same for the capital adjustment-cost estimates presented below.

The instruments are moderately powerful, as judged by the correlation of the fitted value of the right-hand variable with its actual value. The median value of the correlation among the eighteen industries is 0.16, the twenty-fifth percentile is 0.13, and the seventy-fifth percentile is 0.22. There is little evidence against the one overidentifying restriction—the median p-value is 0.32, the twenty-fifth percentile is 0.23, and the seventy-fifth percentile is 0.56, not far from the uniform distribution that would prevail under the null hypothesis. Because most of the correlation of the instruments with the right-hand variable comes from the oil instrument, the power of the test is relatively low, and the acceptance of the overidentifying restriction is not a strong confirmation of the specification.

Serial correlation of the residuals in the Euler equation is generally slightly negative. The median Durbin-Watson statistic across the eighteen equations is 2.24, the twenty-fifth percentile is 1.81, and the seventy-fifth percentile is 2.69.

V.B. Cost of Adjusting Capital

Table III shows the IV estimates of the capital adjustment cost parameter γ for each industry covered by the NIPA data. Figure II shows the corresponding plot of adjustment-cost estimates and precision. Again, the more precise estimates are all close to zero. A few industries with moderate precision have negative estimates. Similarly, a few with moderate precision have estimates as high as one. In addition, there are—as in Figure I—some industries with larger negative or positive estimates of γ with essentially 0 precision. In these industries the instruments are uninformative because their covariances with the right-hand variable are close to zero. The median correlation of the fitted and actual values of the right-hand variable is 0.15, the twenty-fifth percentile is 0.11, and the seventy-fifth percentile is 0.17.

Figure III presents the evidence in a different format. For each industry, represented by a diamond, the horizontal position shows the regression coefficient of the Euler variable (the right side of equation (10)) on the oil instrument (the more powerful of the two instruments), and the vertical position shows the regres-

TABLE III
ESTIMATES OF THE ADJUSTMENT-COST PARAMETER FOR CAPITAL INPUT

Industry	Estimate, $\hat{\gamma}$	Standard error
Farms	0.08	(0.75)
Agricultural services, forestry, and fishing	−0.03	(0.09)
Metal mining	−0.06	(0.43)
Coal mining	0.22	(0.30)
Oil and gas extraction	−0.02	(0.10)
Nonmetallic minerals, except fuels	0.00	(0.28)
Construction	0.20	(0.45)
Lumber and wood products	−0.22	(0.29)
Furniture and fixtures	−0.31	(0.47)
Stone, clay, and glass products	−0.05	(0.18)
Primary metal industries	0.27	(1.28)
Fabricated metal products	0.09	(0.98)
Machinery, except electrical	−0.92	(3.33)
Electric and electronic equipment	−0.09	(0.38)
Motor vehicles and equipment	0.40	(1.30)
Other transportation equipment	3.02	(124.76)
Instruments and related products	1.64	(87.33)
Miscellaneous manufacturing industries	0.05	(0.11)
Food and kindred products	−0.43	(0.88)
Tobacco manufactures	−0.38	(0.82)
Textile mill products	−0.58	(0.86)
Apparel and other textile products	−0.55	(1.62)
Paper and allied products	−0.15	(0.19)
Printing and publishing	0.23	(0.80)
Chemicals and allied products	−0.02	(0.56)
Petroleum and coal products	0.26	(0.95)
Rubber and miscellaneous plastics products	−1.75	(10.17)
Leather and leather products	0.35	(0.33)
Railroad transportation	−0.02	(0.09)
Local and interurban passenger transit	−0.12	(0.33)
Trucking and warehousing	−0.16	(0.34)
Water transportation	0.93	(1.90)
Transportation by air	0.07	(0.13)
Pipelines, except natural gas	−0.07	(0.24)
Transportation services	−0.20	(0.37)
Telephone and telegraph	−0.06	(0.24)
Radio and television	0.38	(2.88)
Electric, gas, and sanitary services	−0.26	(0.30)
Wholesale trade	−0.25	(0.33)
Retail trade	−0.14	(0.62)
Banking	−0.02	(0.31)
Credit agencies other than banks	−0.30	(0.32)
Security and commodity brokers	−0.09	(0.07)
Insurance carriers	−0.13	(0.35)

TABLE III
(CONTINUED)

Industry	Estimate, $\hat{\gamma}$	Standard error
Insurance agents, brokers, and service	0.01	(0.12)
Holding and other investment offices	−0.56	(2.32)
Hotels and other lodging places	−0.08	(0.10)
Personal services	−0.26	(1.63)
Business services	−0.27	(0.35)
Auto repair, services, and parking	−0.12	(0.40)
Miscellaneous repair services	−0.10	(0.18)
Motion pictures	−0.30	(1.09)
Amusement and recreation services	−0.61	(1.03)
Health services	0.51	(2.23)
Legal services	−0.37	(1.29)
Educational services	0.02	(0.23)

Estimates and standard errors of the parameter γ from equation (10), using oil-shock dummies and the innovation in military spending as instruments, with data from the National Income and Product Accounts for expenditure on labor and capital, over the period 1948 through 2001.

sion coefficient of the factor-cost ratio. The coefficient on the Euler variable shows wide dispersion among the industries—investment responses to oil-price shocks vary dramatically by industry. The factor-cost ratio coefficient is very close to zero in every single industry. This finding is inconsistent with even fairly small capital adjustment cost. As equation (10) shows, in the presence of adjustment cost, a force that affects investment must also change the observed factor-cost ratio in proportion to the parameter γ.

Again, there is little evidence against the one overidentifying restriction—the median p-value is 0.52, the twenty-fifth percentile is 0.34, and the seventy-fifth percentile is 0.81, close to uniform. As before, the acceptance of the overidentifying restriction is not a strong confirmation of the specification. The estimates are moderately positively correlated—the median is 0.36, the twenty-fifth percentile is 0.17, and the seventy-fifth percentile is 0.51.

As in the results for labor adjustment, serial correlation of the residuals in the Euler equation is generally slightly negative. The median Durbin-Watson statistic across the 56 equations is 2.23, the twenty-fifth percentile is 2.03, and the seventy-fifth, 2.53.

V.C. Comparison to Other Estimates of Adjustment Costs

For labor, Shapiro [1986] finds zero adjustment cost for production workers and moderate adjustment cost for nonpro-

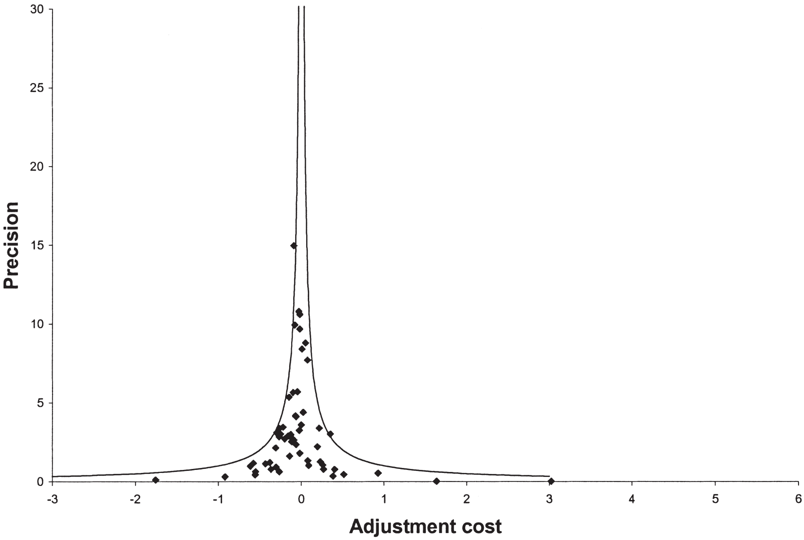

FIGURE II
Instrumental Estimates of Capital Adjustment Cost by Industry
The horizontal position of the marker shows the estimate of the capital adjustment cost γ, for one industry, and the vertical position measures the precision (the reciprocal of the standard error). The fine curving lines mark a precision equal to the horizontal position, corresponding to an estimate equal to its standard error.

duction workers. For capital, my estimate is somewhat below his. See Appendix C of Hall [2001] for a discussion of the interpretation of Shapiro's estimates. His estimates of 8 or 9 for the capital adjustment cost parameter at quarterly frequency correspond to 2 or 2.2 at the annual frequency considered here. Shapiro's specification differs in a number of ways that do not appear to account for the higher adjustment cost he finds—he uses quarterly rather than annual data, he makes adjustment costs proportional to the level of output, and he uses, in effect, the ratio of the value of output to spending on capital where I use the ratio of spending on labor to spending on capital. Rather, all of the difference appears to arise from his specification for the serial correlation of the disturbance. He takes it to be either white noise or a first-order moving average. He uses endogenous variables lagged one or two quarters as instruments, in line with this specification. But the movement of the left-hand variable belies Shapiro's time-series specification, as shown in Figure IV. As I noted earlier, the

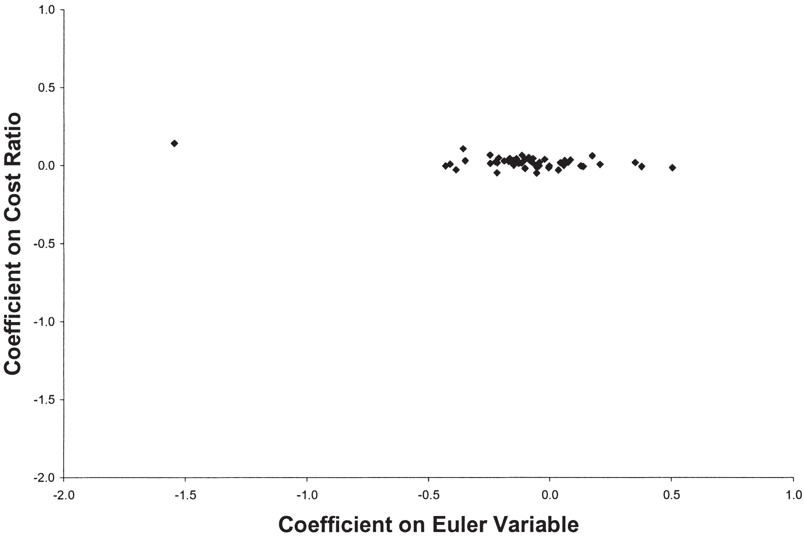

FIGURE III

Reduced-Form Coefficients for Estimates of Capital Adjustment

The horizontal position of a marker shows the estimate of the coefficient of the regression of the right-hand variable in equation (10) (the Euler variable), and the vertical position shows the estimate of the coefficient of the left-hand variable (the cost ratio).

factor expenditure ratio has a random-walk component, and so does Shapiro's inverse of the income share of capital. Since the right-hand variable implied by the Euler equation is a difference of a difference, it cannot have any low-frequency explanatory power—all of the random-walk movement in the left-hand variable must be inherited by the disturbance. Hence his assumption about the exogeneity of lagged endogenous variables must fail. Shapiro tests and rejects the overidentifying restriction in his specification. My finding of only slight negative serial correlation after taking first differences confirms the presence of a low-frequency component of the disturbance that Shapiro did not consider and that appears to be inconsistent with his identifying assumption and choice of instruments.

My estimates are well below the level suggested by Hamermesh and Pfann [1996] and in the range of the small adjustment cost for capital reported by Cooper and Haltiwanger [2002].

In addition to the Euler-equation approach, as I noted

FIGURE IV
Log of the Ratio of Labor and Capital Expenditure, Total Manufacturing
The ratio has a substantial amount of low-frequency drift, a feature of the technology that needs to be incorporated in the adjustment-cost model.

earlier, many authors have pursued Tobin's [1969] insight that the adjustment cost parameter is the reciprocal of the coefficient relating the flow of investment to the ratio of the market value of the capital stock in place to its acquisition cost. That approach yields high—generally absurdly high—estimates of the adjustment cost. For example, in a refined application of the method to excellent data, Gilchrist and Himmelberg [1995] find values for the parameter I call γ of around 20 (see the Tobin's Q columns of their Tables 1 and 2). These findings appear to confirm my conclusion in Hall [2001] that the market values of firms are driven primarily by forces other than the short-term rents earned on capital from adjustment costs. These other factors create an errors-in-variables bias downward in the coefficient of the Tobin regression and thus bias the implied estimate of the adjustment-cost parameter upward substantially.

V.D. Role of the Elasticity of Substitution

My estimates of the convex adjustment-cost parameters are conditional on the hypothesis of Cobb-Douglas technology, with a unit elasticity of substitution. Under a more general technology, the estimates can be interpreted as the product of the adjustment-cost parameter and the elasticity of the corresponding factor demand function conditional on the level of output. In the Cobb-Douglas case, that elasticity is one. In other words, the estimate of the adjustment-cost parameter is the coefficient I estimate divided by the conditional factor-demand elasticity. Over the plausible range of elasticities, the conclusion that the adjustment-cost parameter is generally quite small would continue to hold. For example, Chirinko, Fazzari, and Meyer [2002] estimate an elasticity of substitution between capital and labor of 0.4. Multiplying my estimates of capital adjustment costs by 2.5 would not alter my conclusions substantially.

VI. Discrete Adjustment Costs, Time Aggregation, the Base for Adjustment Costs, and Firm-Specific Disturbances

This section examines a number of issues that arise in estimating adjustment costs:

1. The effect of a discrete adjustment cost—a cost that is incurred to adjust a factor at all;
2. The effect of time aggregation—estimating with annual data when decisions are made more frequently;
3. The choice controlled by the parameter ω between adjustment costs that apply only to net investment ($\omega = 0$) and those that apply to replacement investment as well ($\omega = 1$); and
4. The influence of aggregation across firms with firm-specific shocks.

I carry out this investigation within a laboratory model. To simplify the model, I consider only capital adjustment costs, though I believe the conclusions would carry over to labor adjustment as well. I also omit materials from the list of inputs. A firm in the model inhabits a stationary environment with constant input prices, constant adjustment-cost parameters, and constant discount factor β. I derive the firm's optimal investment strategy as a function of the innovations in its environment. The focus is on the firm's response to changes in product demand, which it sees as changes in the product price.

VI.A. The Laboratory Model

The laboratory technology is

$$y_t = A_t n_t^{\alpha} k_t^{1-\alpha} - \frac{\gamma}{2} g_{t+1}^2 k_t - [g_{t+1}]^0 \phi k_t. \tag{13}$$

The variables and parameters are the same as previously except that I have dropped the k subscript from g. I have added a discrete adjustment cost controlled by the parameter ϕ—the firm incurs a cost $\phi p_t k_t$ to make any change in the capital stock from its base value. $[x]^0 = 0$ if $x = 0$ and 1 otherwise. The following characterization of the optimum follows Abel and Eberly [1994].

Because there is no cost of adjusting labor, the firm chooses its labor/capital ratio to maximize gross margin, $p_t y_t - w n_t$. The maximized value of gross margin per unit of capital is

$$M_t = (1 - \alpha)\alpha^{1/(1-\alpha)} A^{1/(1-\alpha)} w^{1/(1-\alpha)} p_t^{\alpha/(1-\alpha)}. \tag{14}$$

The firm chooses its current investment growth rate g_t to maximize expected discounted value per unit of current capital. I state the first-order condition for capital as a pair of conditions involving the shadow premium on installed capital, $(q_t - 1)p_k$ (the q is Tobin's q):

$$g_t = \frac{\beta p_k}{\gamma p_{t-1}}(q_t - 1) \text{ if } \frac{\beta^2}{2\gamma}\frac{p_k}{p_t}E_t(q_{t+1} - 1)^2 > \phi\frac{p_t}{p_k} \text{ else } 0, \tag{15}$$

$$q_t - 1 = \frac{M_t - \rho}{p_k} + (1 - \omega\delta)\beta E_t(q_{t+1} - 1) + \left[\frac{\beta^2}{2\gamma}\frac{p_k}{p_t}E_t(q_{t+1} - 1)^2 - \phi\frac{p_t}{p_k}\right]^+. \tag{16}$$

Equation (15) is Tobin's [1969] investment equation extended to the case of discrete adjustment cost. Equation (16) states q as a backward recursion or Bellman equation. The last term on the right side measures the benefit of adjusting the capital stock from its base value. Unless it exceeds the threshold of fixed adjustment cost, $\phi p_t/p_k$, the firm will choose $g_t = 0$. In that case, the last term will take the value zero instead.

The shadow value q_t is a random variable depending on all the information relevant to forecasting the future environment of the firm controlled by the random variable M_t. Notice that the

Bellman equation does not involve any endogenous state variables. It is reasonable to assign a discrete distribution to M_t and thereby avoid the problems of interpolation that arise with continuous distributions. Let s be a discrete random variable obeying a first-order Markov process, interpreted as the state of demand. Then the values $q_{t,s}$ can be calculated from the backward recursion of equation (16) by straightforward evaluation over the discrete distribution. In the stationary environment of the laboratory model, the backward recursion starting from $q_{T,s} = 1$ for a distant horizon T converges to a stationary vector q_s.

I model the evolution of the gross margin M_t in the following way: the price is the product of an index $\bar{p}_i$ that affects all firms equally and an index $\bar{z}_j$ that is specific to the firm: $p_t = \bar{p}_{i_t}\bar{z}_{j_t}$. The marketwide disturbance $i_t \in [1,5]$ rises by 1 (if it is less than 5) with probability $(1 - \rho)/2$ and falls by 1 (if it exceeds 1) with the same probability. It remains at its current level with the remaining probability. The firm-specific disturbance $j_t \in [1,3]$ has a similar process with parameter λ. The overall state combines the marketwide and firm-specific elements: $s_t = (i_t, j_t)$. There are fifteen values of $q_{t,s}$ to be calculated from equation (15).

Table IV shows the parameter values used in the calculations. I chose the last four rows from known features of the data. I chose ρ and $\bar{p}$ from data on the serial correlation and variance of labor input per unit of capital, stated as deviation from a trend with constant growth. Cooper and Haltiwanger [2002] report that the standard deviation of the plant-specific shock is 2.7 times as large as the standard deviation of the marketwide shock and that the annual serial correlation of the plant-specific shock is 0.53. I chose the values of $\bar{z}$ and λ in view of these findings. Note that ρ and λ are stated in monthly time units. The annual serial correlations of the price components are 0.77 for the aggregate and 0.30 for the firm-specific. The former is a bit below the level reported by Cooper and Haltiwanger, and the latter is distinctly lower, for the following reason: in the laboratory model, the price is exogenous, to achieve an enormous simplification of the computations, as I noted earlier. Cooper and Haltiwanger measure shifts of the demand schedules facing firms, rather than changes in prices. The prices implied in their work are less persistent than are demand shifts, because of capital adjustment, which makes the longer-run effect of a demand shift on price smaller than the immediate effect.

TABLE IV
PARAMETER VALUES FOR THE LABORATORY MODEL

Parameter	Explanation	Value
α	Elasticity of output with respect to labor input	0.7
β	Discount ratio, at annual rate	0.95
γ	Convex adjustment cost, at annual rate	1.0
ϕ	Discrete adjustment cost	0, 0.00005 (when $\omega = 0$), 0.0002 (when $\omega = 1$)
δ	Rate of deterioration of capital, at annual rate	0.10
$\bar{p}$	Vector of alternative values of marketwide component of price	0.955, 0.9775, 1, 1.0225, 1.045
$\bar{z}$	Vector of alternative values of firm-specific component of price	0.88, 1, 1.12
ρ	Serial correlation parameter for marketwide price component, monthly	0.888
λ	Serial correlation parameter for firm-specific price component, monthly	0.809

The Euler equation in the laboratory model is

$$(17) \qquad M_t - \rho = \gamma\left\{\frac{p_{t-1}g_t}{\beta} - p_t E_t\left[(1-\omega\delta)g_{t+1} + \frac{1}{2}g_{t+1}^2\right]\right\}.$$

Although the presence of p_{t-1} suggests that the right-hand side of the Euler equation depends on last period's state, in fact it does not, as the p_{t-1} cancels the one in equation (15).

I consider several ways to view the behavior of the laboratory model through the lens of estimation. One considers monthly aggregates over firms. Estimating the Euler equation from the aggregate monthly data involves two specification errors. One is the use of aggregates across firms when the true Euler equation applies at the level of the firm. The second is the use of the Euler equation derived from the hypothesis of zero discrete adjustment cost when the actual discrete cost is positive. The laboratory model implies a joint distribution for the monthly aggregates. From the covariances of the joint distribution, I can calculate the

TABLE V
ESTIMATES OF THE CONVEX ADJUSTMENT COST IN THE LABORATORY MODEL

Line	Discrete adjustment cost	Firm-specific demand shock	Net or gross investment	IV estimate of slope of monthly Euler equation (true value = 12)	IV estimate of slope of annual Euler equation (true value = 1.0)
1	No	No	Net	12.0	1.10
2	Yes	No	Net	12.0	1.09
3	No	Yes	Net	12.0	1.09
4	Yes	Yes	Net	13.4	1.18
5	No	No	Gross	12.0	1.11
6	Yes	No	Gross	9.2	0.90
7	No	Yes	Gross	12.0	1.12
8	Yes	Yes	Gross	11.2	1.00

plims of the monthly and annual estimates from the misspecified Euler equation of the convex adjustment-cost parameter γ for all possible alternative specifications—that is, with and without a discrete adjustment cost, with and without a firm-specific demand shock, and with adjustment costs applying to net or gross investment. The calculations for the monthly estimates are exact—I calculate the covariances directly from the transition probabilities.

I also consider aggregation over months to form annual data. I measure the effects of the specification error of fitting an annual Euler equation—derived from the assumption that investment decisions are made once a year—to data that are aggregated from monthly decisions. The calculations for the annual estimates use a Monte Carlo method because the number of possible monthly transitions over a year is way beyond enumeration.

I examine all eight possible combinations of specification error for the net and gross investment specifications. To enforce stationarity on investment, I adjust the mean of the gross margin so that expected net investment is zero in each case. Table V shows the results of these calculations, with the lagged value of the gross margin serving as instrumental variable (the same estimate results from regressing the Euler variable on the margin and taking the reciprocal of the resulting coefficient, as the laboratory model's only disturbance is the expectation error). I take the monthly adjustment cost γ to be 12.

VI.B. Discrete Adjustment Costs

A vibrant recent literature has developed an analysis of discrete adjustment costs and demonstrated the empirical importance of those costs. See Caballero [1999] for a survey. Caballero and Engel [1999], Cooper and Haltiwanger [2002], Thomas [2002], and Khan and Thomas [2003] investigate the implications of discrete adjustment costs for aggregate investment.

For positive discrete adjustment cost, equation (15) defines a zone of inaction, where there is no adjustment, because the benefit of adjustment would not cover the discrete cost. An impulse large enough to push a firm outside its zone of inaction causes a substantial response. Aggregation of firms into industries tends to conceal most of this behavior. An industrywide impulse has no effects on most firms but large effects on some, and the average is not so different from the corresponding setup without the discrete cost. In fact, under certain (stringent) assumptions, there is exact cancellation and no aggregate implications of discrete adjustment costs, a point made by Caplin and Spulber [1987] in the context of discrete costs of price adjustment.

Cooper and Haltiwanger [2002] estimate a model with discrete adjustment costs, using plant-level data, finding, as expected, that the discrete costs add much to the realism of the model for the micro data. They fit a model with convex adjustment to the aggregate of their plants and find that investment predicted by the aggregate model is reasonably close to the aggregate data generated by the underlying model with plant-level discrete costs. Caballero and Engel [1999] study the implications for industry-level time series of a general model that includes discrete costs. They estimate either nonlinear adjustment rates or adjustment costs from panel data by industry and year. They compare the forecasting abilities of these models to a model with a constant adjustment rate, as implied by the quadratic specification I employ. For structures, they find a substantial improvement in forecasting power, but only a small improvement for equipment. Because equipment is about 80 percent of total investment, it appears that their results confirm that, for industry aggregates, the quadratic specification provides a reasonably accurate approximation.

Equation (15) reveals how the presence of a discrete cost of adjustment might bias estimates of the convex adjustment parameter γ if the discrete cost were ignored in the estimation process. Some of the demand states that the estimation model believes should have positive or negative levels of investment, according to the equation with $\phi = 0$, actually have zero levels with $\phi > 0$. But the equation suggests that the bias is small, because the states that are held at zero are ones where investment would be small in the absence of the discrete cost. The more influential observations in the estimation, with larger positive or negative amounts of investment, remain unaffected by the discrete cost. I will note shortly that this conclusion does not hold for the model when adjustment costs apply to replacement investment.

In Table V the column of monthly estimates shows that the bias in estimating the adjustment cost in monthly data arises entirely from the discrete cost. In the presence of the discrete cost, the convex adjustment cost parameter is overstated when there are firm-specific demand shocks (line 4). On line 8 the discrete cost pulls the estimate down while the firm-specific shock pulls it up, with the discrete cost winning by a small margin.

The monthly estimates show that discrete adjustment costs and the firm-specific demand shock have an interactive effect. Absent a discrete cost and a firm-specific shock, the Euler equation holds exactly. There is essentially no bias from firm-specific shocks when there is no discrete adjustment cost. With a discrete cost, the firm-specific shock biases the adjustment coefficient upward. The effect is stronger for gross investment, where it offsets the downward bias from the discrete cost itself.

VI.C. Gross or Net Investment?

Table V shows that the adjustment cost is understated when it applies to gross investment, when there is a discrete adjustment cost, and there are no firm-specific shocks (line 6). Earlier I noted that the discrete cost has relatively little effect on investment when it results in zero levels of investment instead of the small levels that would occur without a discrete cost. This property is special to the case where the average level of the investment variable, $k_{t+1} - (1 - \omega\delta)k_t$ is 0. When replacement investment bears adjustment costs ($\omega = 1$) or when the capital stock is growing fairly rapidly, the average value will be positive, and the situation is different. Khan and Thomas [2003] discuss how this higher sensitivity and nonlinearity biases estimates of convex adjustment costs downward. The effect of the discrete cost is distinctly asymmetric, and the cancellation proposed by Caplin and Spulber [1987] does not occur. The bias is greatest in a setting where most firms constrained by discrete adjustment costs are near the upper ends of their zones of inaction. A positive aggregate shock induces a sharp positive response of investment among the firms who are pushed out of their zones of inaction toward positive investment. But a negative aggregate shock does not have a similar effect of stimulating disinvestment, because few firms are near the lower boundary of the zone of inaction. An estimation method that inferred the convex adjustment-cost parameter from the slope of the relation between the driving force and investment—such as the indirect inference method used by Cooper and Haltiwanger [2002]—could be biased in the presence of the discrete cost. Alternatively, one could use the nonlinearity of the response to infer the magnitude of the discrete cost, the approach Cooper and Haltiwanger actually take.

In this model, the need to replace deteriorating capital accounts for the asymmetry in the investment response in the gross investment model. The asymmetry could also arise in a nonstationary environment with positive growth of the capital stock. Then the typical value of the investment rate would be positive, and the same crowding of firms into the upper range of the zone of inaction would occur. However, because growth rates of capital are generally far below the rate of deterioration of capital (around 10 percent per year), the second effect is much weaker than the first.

VI.C. Time Aggregation

Earlier, I used annual data by industry for estimation. Because decisions about factor inputs are made more frequently than once a year, estimation in annual data results in a bias from time aggregation, as shown in Table V. To explore the bias in the laboratory model, I aggregate the results from a monthly model in the same way that the data are aggregated in the national accounts. Labor input is measured as the total number of hours of work over the year, and the capital stock is measured at the end of the year.

The column for annual estimates shows larger biases, upward in all cases except line 6. Time aggregation is a substantial source of the upward bias. The bias is worst when (1) there is a firm-specific shock, (2) the adjustment cost applies to net investment, and (3) there is a discrete adjustment cost. Even in this case, the bias of 18 percent is not severe.

VI.D. Firm-Specific Shocks

The last of the specification errors is aggregation over firms in the presence of heterogeneous shocks. As Table V shows, aggregation in this dimension does not bias the estimated convex adjustment cost coefficient γ at all unless there are other specification errors. Lines 3 and 7 of the monthly column of the table show the effects of aggregation over firms in the absence of the other specification errors: temporal aggregation and neglect of the discrete adjustment cost. The coefficient is correctly estimated on those lines. But aggregation over firms exacerbates the bias from the other specification errors.

The most conspicuous effect of aggregation over firms is a reduction in the dispersion of the right-hand variable. In that respect, aggregation results in less informative data. But aggregation over firms has remarkably little effect on the slope of the Euler equation, that is, on the estimate of the adjustment-cost parameter γ.

VI.E. Implications for Measuring Tobin's q

The general conclusion from the laboratory model is that the three specification errors—even if simultaneously present—do not seriously bias the estimate of the convex adjustment-cost parameter γ. But that conclusion does not imply that the annual Euler equation that ignores discrete costs of adjustment will give an accurate picture of industry-level investment and related measures in any particular application. The discrete adjustment cost could have important effects even if it does not bias the estimate of the convex cost by much. Because my interest in the measurement of Tobin's q motivated this research, I will examine biases in that measurement that arise from the specification errors.

The natural way to measure Tobin's q is to invert equation (15) under the assumption of zero discrete adjustment cost:

$$\hat{q}_t - 1 = \frac{\gamma p_{t-1}}{\beta p_k} g_t. \tag{18}$$

The issue for investigation is the accuracy of this calculation, when applied to annual aggregate data, as an estimate of the actual shadow value of the capital stock, calculated monthly at the firm level taking proper account of the discrete adjustment cost. Figure V shows the comparison in the laboratory model for

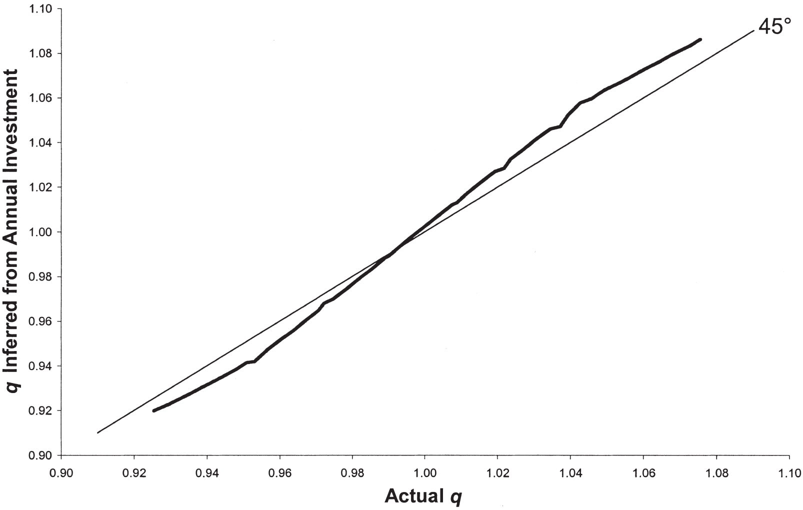

FIGURE V
Relation between Estimated and True Values of Tobin's q

The heavy line shows the relation between the true value of Tobin's q, aggregated from the monthly laboratory model up to annual data (on the horizontal axis), and the value calculated using equation (18) from annual data using the value of the adjustment-cost parameter estimated from annual data. The light line shows an exact relation, and the vertical distance between the lines shows the error made as a result of the specification errors in this approach to inferring Tobin's q.

the case corresponding to line 4 of Table V, with all three specification errors in effect. The horizontal axis of the figure is the correctly measured value of q, and the vertical axis is the expected value of $\hat{q}_t$ from equation (17), conditional on the correct q.

The value of q inferred from annual investment understates the correct value of q when q is less than 1; the error is about 0.01 at the maximum, around a true q of 0.95. Similarly, the inferred q overstates the true value above 1, with a maximum error of a little more than 0.01 at a true q of 1.05. But these errors are tiny in relation to the unexplained movement of securities values in relation to the reproduction cost of capital, as shown in Hall [2001] and elsewhere. Figure V supports the conclusion of my earlier paper that adjustment costs are not a substantial part of the explanation of the movements of securities values by demonstrating that estimates from annual investment data are not badly biased.

VII. Concluding Remarks

Using standard econometric methods on factor-adjustment Euler equations, I am able to measure adjustment costs separately for labor and capital. The resulting parameter estimates suggest low adjustment costs for both factors. Although these estimates are biased by three specification errors—neglect of discrete adjustment costs, aggregation over time, and aggregation across heterogeneous firms—an investigation of the magnitude of the biases in the laboratory model suggests that they are modest. The finding of low adjustment costs in the data cannot be the result of high actual adjustment cost confounded by specification-error bias of the magnitude found in the laboratory model.

The finding of low factor adjustment costs implies that transitory rents from factor adjustment are not an important source of variation in the market values of firms. An unexpected rise in product demand does not push a firm up its product supply schedule and generate rent equal to the area of the triangle under the schedule. Rather, it moves the firm to the right along a flat supply curve. It may also raise the flat supply curve, in which case the factor suppliers—workers or capital goods producers—earn rents. Because capital goods producers also have flat supply curves, the results imply that ultimately labor earns all the rents.

Although theory and intuition agree that one should be able to estimate the parameters of adjustment costs from the relation between the values of the capitalized rents and the changes in labor input and purchases of capital goods, no useful data have been uncovered for that purpose. Data from securities markets appear to be hopelessly contaminated by factors other than capitalized adjustment rents. Rather, the only usable evidence about adjustment costs comes from the joint movements of factor inputs. Labor and capital input move in parallel with fully variable factors—labor and capital do not appear to be held back by significant adjustment costs.

Hoover Institution and Department of Economics, Stanford University, and National Bureau of Economic Research

References

Abel, Andrew B., and Janice C. Eberly, "A Unified Model of Investment under Uncertainty," *American Economic Review,* LXXXIV (1994), 1369–1384.
Caballero, Ricardo J., "Aggregate Investment," in John B. Taylor and Michael

Woodford, eds., *Handbook of Macroeconomics* (New York: Elsevier, 1999), pp. 816–862.
Caballero, Ricardo J., and Eduardo M. R. A. Engel, "Explaining Investment Dynamics in U. S. Manufacturing: A Generalized (S,s) Approach," *Econometrica,* LXVII (1999), 783–826.
Caplin, Andrew S., and Daniel F. Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics,* CII (1987), 703–726.
Chirinko, Robert S., Steven M. Fazzari, and Andrew P. Meyer, "That Elusive Elasticity: A Long-Panel Approach to Estimating the Capital-Labor Substitution Elasticity," Federal Reserve Bank of St. Louis Working Paper, 2002.
Cooper, Russell W., and John C. Haltiwanger, "On the Nature of Capital Adjustment Costs," mimeo, 2002; previously NBER Working Paper No. 7925, 2000.
Garber, Peter M., and Robert G. King, "Deep Structural Excavation? A Critique of Euler Equation Methods," NBER Technical Working Paper No. 31, 1983.
Gilchrist, Simon, and Charles P. Himmelberg, "Evidence on the Role of Cash Flow for Investment," *Journal of Monetary Economics,* XXXVI (1995), 541–572.
Hall, Robert E., "The Stock Market and Capital Accumulation," *American Economic Review,* XCI (2001), 1185–1202.
——, "Corporate Earnings Track the Competitive Benchmark," NBER Working Paper No. 10150, 2003.
Hamermesh, Daniel S., and Gerard A. Pfann, "Adjustment Costs in Factor Demand," *Journal of Economic Literature,* XXXIV (1996), 1264–1292.
Hansen, Lars Peter, "Interview," *Journal of Business and Economic Statistics,* XX (2002), 442–447.
Jorgenson, Dale W., *Productivity: International Comparisons of Economic Growth* (Cambridge, MA: The MIT Press, 1995).
Khan, Aubhik, and Julia K. Thomas, "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?" *Journal of Monetary Economics,* L (2003), 331–360.
Shapiro, Matthew D., "The Dynamic Demand for Capital and Labor," *Quarterly Journal of Economics,* CI (1986), 513–542.
Thomas, Julia K., "Is Lumpy Investment Relevant for the Business Cycle," *Journal of Political Economy,* CX (2002), 508–534.
Tobin, James, "A General Equilibrium Approach to Monetary Theory," *Journal of Money, Credit and Banking,* I (1969), 15–29.