

NBER WORKING PAPER SERIES

INDUSTRY DYNAMICS WITH ADJUSTMENT COSTS

Robert E. Hall

Working Paper 8849
<http://www.nber.org/papers/w8849>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2002

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. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

© 2002 by Robert E. Hall. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Industry Dynamics with Adjustment Costs
Robert E. Hall
NBER Working Paper No. 8849
March 2002
JEL No. E3, J23, D92

ABSTRACT

Adjustment costs determine the dynamics of the response of an industry's output to a shift in demand. Absent any adjustment costs, an increase in demand not accompanied by any change in factor prices raises output, labor, capital, and materials in the same proportion. In the presence of adjustment costs, the elasticity of the response of factors with higher costs is less than one while the elasticity of those without adjustment costs exceeds one. I develop a model of industry dynamics to capture these properties and a related econometric framework to infer adjustment costs from the observed ratios of factor responses to output responses. I find relatively precise evidence of moderate adjustment costs.

Robert E. Hall
Hoover Institution
Stanford University
Stanford, CA 94305
and NBER
hall@hoover.stanford.edu

I. Introduction

Output and factor inputs are surprisingly volatile in many U.S. industries. Figure 1 shows data for the electrical equipment industry.

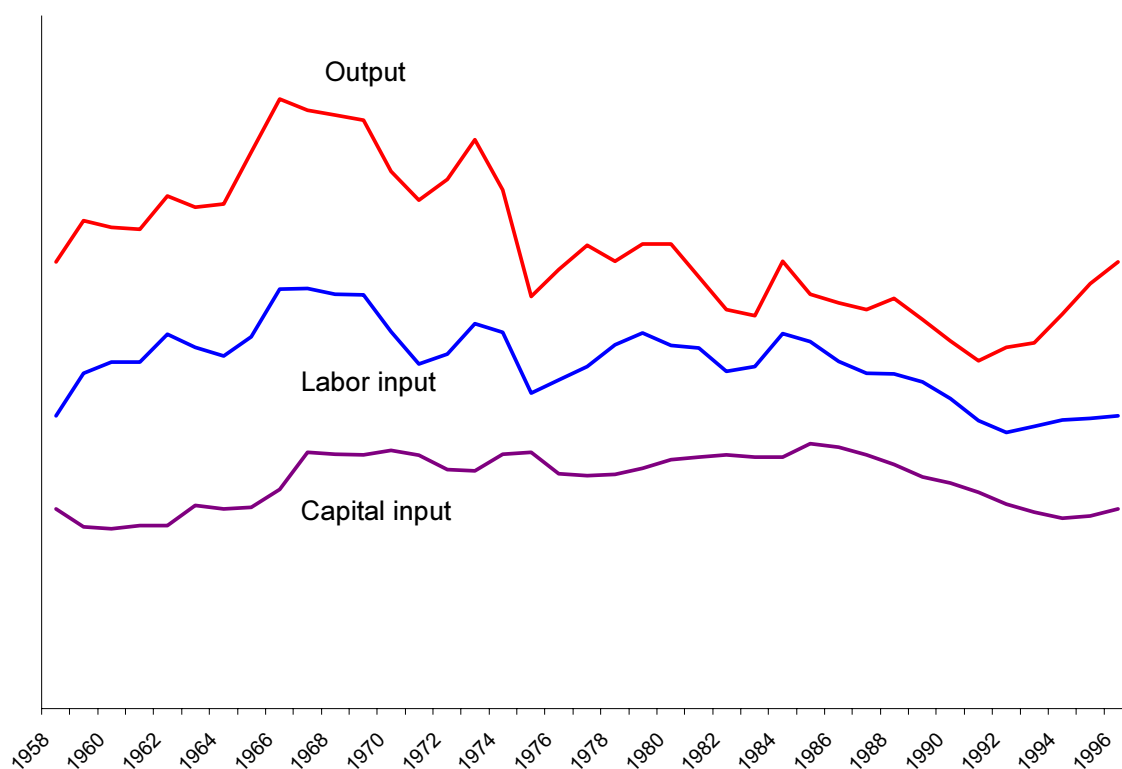


Figure 1. Output, Labor Input, and Capital Input in the Electrical Equipment Industry

Source: Data compiled by Dale Jorgenson; see Jorgenson and Stiroh [2001]. Data are detrended with arbitrary scales. The horizontal axis is zero for all three series.

This paper builds an equilibrium model of industry dynamics. The model resembles—and is inspired by—dynamic stochastic general equilibrium models of the entire economy, but its scope is more modest. Questions that can be answered with the model include those about the dynamics of business earnings, the issue that led me to

develop the model. I take as the baseline model a Cobb-Douglas technology, competitive product and factor markets, and adjustment costs for labor and capital.

The central topic of investigation is adjustment costs. Some authors, such as Hamermesh and Pfann [1996], have concluded that adjustment costs substantially limit changes in labor and capital input in the short run. In such an economy, an increase in demand for an industry's product causes an extended period of rents, as the industry moves up its supply curve and prices rise. I find that adjustment costs are modest for both labor and capital. The supply functions implicit in this work are quite elastic. 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. This paper, however, does not consider that subject any further—it is almost entirely concerned with adjustment costs. I do find that there are adjustment costs for labor as well as capital, a source of transitory rents not considered in my earlier work.

The paper tries to take data, identification, and estimation seriously. I use Dale Jorgenson's 35-industry panel of annual data for output and factor inputs over the period 1959 through 1999. These data embody careful consideration of input quality and price measurement. My identification strategy exploits the panel nature of the data. I study the time-series correlations of output, labor input, and capital input with GDP for each industry. In the typical industry, the elasticity of output with respect to GDP is somewhat higher than the elasticity of labor input with respect to GDP and substantially higher than the elasticity of capital input. The key empirical measures in the paper are the *labor/output response ratio* and its counterpart for capital. The labor/output response ratio is the ratio of the labor-GDP elasticity to the output-GDP elasticity. Absent adjustment costs and responses of factor prices, an increase in demand would cause factor inputs to change in the same proportion as output, and all the factor/output response ratios would be 1.0. If a factor has adjustment costs, its own response ratio is less than one and the response ratios

of other factors are higher than they would be otherwise. If any factor has adjustment costs, factors without any adjustment cost will have output response ratios exceeding one.

I use GDP as an index of aggregate shocks. These shocks include changes in government purchases, shifts in consumption and investment behavior, changes in the terms of trade, aggregate changes in productivity growth, and the like. The aggregate shocks affect individual industries in a number of ways—they shift industry demand, they affect productivity, and they affect factor prices. The aim of my approach to estimation is to isolate the demand effects from the other effects.

The correlation of GDP with industry output shows tremendous variation across industries. The measured elasticities range from -0.2 to 5. My identifying assumption is that this variation results from heterogeneity in demand effects. Effects operating through national markets should be reasonably uniform across industries—for example, the elasticity of the capital stock with respect to an increase in real interest rates caused by an aggregate shock should be about the same across industries. My estimation method measures uniform effects of aggregate shocks—indexed by real GDP—and does not use them to infer adjustment costs. Only *differences* across industries in GDP elasticities enter the measurement of adjustment costs. I measure the labor/output response ratio as the ratio of the slope of the labor-GDP elasticities across industries to the slope of the output-GDP elasticities. The intercept of this relationship captures the factor-price and aggregate productivity effects that are approximately uniform across industries.

In effect I measure the slope of the short-run supply functions and short-run factor demand functions of the industries. My econometric method makes the notion of the short run precise—the short run is the response within a year to a time-series innovation in GDP in that year. As a result, I cannot distinguish internal from external adjustment costs. For example, the limited short-run response of employment could result from training and startup costs within the firm, or it could reflect the short-run inelasticity of labor supply to an industry owing to labor mobility costs. For capital, what I measure as adjustment costs

could include costs that actually arise in the capital-goods-supplying industries even though I measure them in the using industries.

I find that the labor/output response ratio is 0.77 and capital/output response ratio is 0.35. Both estimates have small standard errors—0.06 and 0.05. Materials inputs make up the difference and have a response ratio greater than one, though I do not pursue this feature in my approach to estimation. To interpret these two empirical findings, I use the model and employ the strategy of indirect inference. First, I show that the model's prediction about these factor/output response ratios is close to invariant to parameters other than the adjustment costs. Hence, I can solve the model for the values of the adjustment costs of labor and capital that predict these response ratios without knowing the values of the other parameters. From the derivatives of that solution, I can also find the standard errors of the adjustment-cost parameters. The resulting estimates suggest adjustment costs somewhat below those estimated by other methods.

In my empirical work and parallel analysis with the industry model, I concentrate on the responses to an innovation in the economic environment. In the model, this is an innovation in real GDP, taken as a general measure of the effects of aggregate influences on all industries. I argue that the study of responses to innovations is a robust approach to the estimation problem. I show that it is much less sensitive to specification errors than the existing approach based on the Euler equation, as in Shapiro [1986]. In addition, the method is less sensitive to specification error than estimation based on observed measures of Tobin's q , although a huge specification error is needed to explain the bias in estimates of adjustment costs with data from the U.S. stock market.

II. Industry Model **Equation Section 2**

Consider the following broad class of dynamic models. An exogenous driving force, x_t , evolves according to

$$x_t = f(\varepsilon_t, x_{t-1}). \quad (2.1)$$

Here ε_t is a serially uncorrelated random shock and f is a stationary rule mapping the shock into a new value of the driving force given its previous value. Participants in the industry observe ε_t and the lagged values of the driving force, and they know the rule f . Their interaction results in an equilibrium described by a vector of endogenous variables, y_t , and a function, g :

$$y_t = g(v_t, y_{t-1}, \varepsilon_t, x_{t-1}). \quad (2.2)$$

Here v_t is a vector of random variables associated with the formation of the endogenous variables and with other random aspects of the industry's environment.

In my application of this framework, the endogenous variables, y_t , resulting from industry equilibrium are output, price, employment, capital stock, and materials input. In the model, firms face adjustment costs that require looking into the future to maximize value. I solve for the entire industry equilibrium over time. The autoregression for the demand shift is part of the dynamic model. The model solves for the equilibrium in the industry as a function of the current innovation in demand with the knowledge that future levels of demand will probably be higher as a result of this year's innovation.

I model industry equilibrium as follows. First is the production technology, relating labor input, n_t , materials input, m_t , and capital input, k_t , to output, q_t :

$$q_t = A_t n_t^\alpha m_t^\psi k_t^{1-\alpha-\psi}. \quad (2.3)$$

A_t is an index of productivity, growing over time at a possibly variable rate. Next is product demand,

$$q_t = d_t z_t^\omega p_t^{-\delta}; \quad (2.4)$$

The unobserved disturbance, d_t , shifts the position of the demand function, z_t is an aggregate shift, ω is its elasticity, and δ is the price elasticity of demand, taken to be a constant.

Third is the factor adjustment technology. I assume that labor is the only input required to increase the level of labor input to production and to increase the level of capital input to production. These are recruiting and training costs for labor and planning and installation costs for capital. Adjustment costs are convex in the inputs—they amount to

$$\frac{\lambda}{2} w_{t+1} \frac{(n_{t+1} - n_t)^2}{n_t} + \frac{\gamma}{2} w_{t+1} \frac{(k_{t+1} - k_t)^2}{k_t} . \quad (2.5)$$

These costs have constant returns to scale as discussed in Hall [2001]. Notice that I do not include discrete costs of adjustment, despite their clear importance for understanding factor adjustment at the plant level. I argue in Section III that discrete costs have little role in industry dynamics.

I approximate the stochastic equilibrium of the industry by solving the model with infinitesimal demand innovations. The quality of this approximation is known to be high. The firm maximizes the present value of its future cash flows. The relevant terms of the present value for a decision made in period t are:

$$\begin{aligned} & \frac{1}{1+r_t} \left[p_t A_t n_t^\alpha m_t^\psi k_t^{1-\alpha} - w_t n_t - p_{m,t} - p_{k,t+1} (k_{t+1} - s k_t) - \frac{\lambda}{2} w_{t+1} \frac{(n_{t+1} - n_t)^2}{n_t} - \frac{\gamma}{2} w_{t+1} \frac{(k_{t+1} - k_t)^2}{k_t} \right] \\ & - p_{k,t} (k_t - s k_{t-1}) - \frac{\lambda}{2} w_t \frac{(n_t - n_{t-1})^2}{n_{t-1}} - \frac{\gamma}{2} w_t \frac{(k_t - k_{t-1})^2}{k_{t-1}} \end{aligned} \quad (2.6)$$

Here w_t is the wage, $p_{m,t}$ is the price of materials inputs, $p_{k,t}$ is the price of new capital goods, and s is the survival rate for capital (one minus the deterioration rate).

Let $g_{n,t} = \frac{n_t - n_{t-1}}{n_{t-1}}$ and similarly for $g_{k,t}$. The first-order conditions for value maximization are:

$$\frac{1}{1+r_t} \left[\alpha \frac{p_t q_t}{n_t} - w_t + \lambda w_{t+1} (g_{n,t+1} + g_{n,t+1}^2) \right] - \lambda w_t g_{n,t} = 0 \quad (2.7)$$

$$\psi \frac{p_t q_t}{m_t} = p_{m,t} \quad (2.8)$$

$$\begin{aligned} \frac{1}{1+r_t} \left[(1-\alpha-\psi) \frac{p_t q_t}{k_t} + s p_{k,t+1} + \gamma w_{t+1} (g_{k,t+1} + g_{k,t+1}^2) \right] \\ - p_{k,t} - \gamma w_t g_{k,t} = 0 \end{aligned} \quad (2.9)$$

I model the aggregate demand shift as a first-order autoregression,

$$z_t = \rho z_{t-1} + \varepsilon_t \quad (2.10)$$

My goal is to derive the difference between the equilibrium resulting from an innovation in the driving force and the stationary equilibrium. The stationary equilibrium solves

$$q^* = A n^{*\alpha} m^* \psi k^{*1-\alpha-\psi} \quad (2.11)$$

$$y^* = dz p^{*-\theta} \quad (2.12)$$

$$\alpha \frac{p^* y^*}{n^*} = w \quad (2.13)$$

$$\psi \frac{p^* y^*}{m^*} = p_m \quad (2.14)$$

$$\frac{1-\alpha}{\mu} \frac{p^* y^*}{k^*} = \frac{1-s}{1+r} p^k \quad (2.15)$$

To find the response function to an innovation, I find the derivatives of the model with respect to ε_0 with all later innovations taken to be zero. For convenience I take the base level of all the exogenous variables to be 1.

I then solve for the derivatives of equations (2.3), (2.4), (2.7), (2.8), (2.9), and (2.10) with respect to ε_0 together with those of the boundary conditions,

$$k_{-1} = k^*, n_{-1} = n^*, k_{T+1} = k^*, \text{ and } n_{T+1} = n^* . \quad (2.16)$$

All 5 equations apply over periods 0 to $T-1$. I choose T to be large enough to approximate the response in the case of an infinite horizon. There are 5 endogenous variables, y , p , m , k , and n . The derivatives of the first three are to be determined in periods 0 through $T-1$, and the second two in periods -1 through T . There are $5T$ equations and 4 boundary conditions. There are $3T + 2(T+2)$ values to be solved. I solve for the derivatives of the equilibrium of the model by solving the two-point boundary value problem as a large system of linear equations.

III. Some Properties of the Model **Equation Section (Next)**

A. Differences across Industries and over Time

When an aggregate impulse stimulates demand in one industry, the industry's output rises and its use of all factors of production rises. The factors with lower adjustment costs rise more. The basic strategy of the paper is to infer adjustment costs from the ratios

of the increase in factor inputs to the increase in output when demand rises. I associate aggregate impulses with changes in GDP. There is wide variation in the empirical elasticity of industry output with respect to GDP. An important feature of the model, however, is that the ratio of factor responses to demand shifts and output responses to those demand shifts is essentially invariant to other parameter values. Neither the elasticity of demand nor the elasticities of the production function have any significant influence on this ratio. Thus, if adjustment costs are roughly the same in all industries, they can be estimated in panel data despite variation in the other parameters across the industries in the panel.

The factor/output response ratios of the model are literally invariant to the elasticity ω of the demand function with respect to the demand shift. The evidence suggests wide variation in these elasticities. Invariance with respect to them is the essential property of the model leading to the estimation strategy of the paper. Industries that respond very differently to aggregate shocks nonetheless respond the same when the factor responses are normalized by the output responses.

The factor/output response ratios are close to invariant with respect to the industry elasticity of demand. Both output and factor responses are smaller in industries with more elastic demand, but they are smaller in proportion. Because I have not yet found a workable strategy for estimating demand elasticities, this near-invariance is fortunate. Lowering the demand elasticity from 1 to 0.5 changes the labor/output response ratio from 0.77 to 0.75 and the capital/output response ratio from 0.35 to 0.33.

An impulse that raises GDP does more than shift an industry's demand function. It also affects factor prices. The estimation method used here requires a separation of the demand effects from the factor-price effects. The method exploits the substantial variation across industries in the demand effects. I hypothesize that there is no cross-sectional correlation of the factor-price effects and the demand effects. For example, an increase in government purchases raises real interest rates. I assume that the effect of the higher interest rate on, say, capital employed in an industry is uncorrelated with the industry's

demand shift. In the model, the effect of a change in the interest rate is literally invariant to the demand-shift elasticity, ω . In fact, the only parameter that has much influence on a factor's response to prices is the adjustment cost for that factor and the adjustment costs for other factors. Thus the hypothesis, more than anything else, is an assertion that adjustment costs are not correlated across industries with responses to aggregate impulses operating through factor prices.

B. Discrete Adjustment Costs

The model embodies the traditional specification of convex adjustment costs. A vibrant recent literature has developed an analysis of non-convex costs—generally a discrete cost of undertaking any change—and demonstrated the empirical importance of those costs. In particular, the evidence is overwhelming that discrete costs influence plant-level investment—the great majority of plants have zero investment in a given year, an inexplicable circumstance with convex adjustment. See Caballero [1999] for a survey. Caballero and Engel [1999], Cooper and Haltiwanger [2002], and Thomas [2001] investigate the implications of discrete adjustment costs for industry aggregate investment.

In brief, models with discrete adjustment costs work as follows: Within a zone of inaction, there is no adjustment, as its benefit 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 industry-wide 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 [2000] 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. Thomas [2001] reports that discrete costs have almost no implications for aggregate general equilibrium dynamics.

IV. Econometric Framework **Equation Section (Next)**

To estimate the adjustment costs, I use indirect inference (Smith [1993] and Gouriéroux, Montfort and Renault[1993]). In that framework, a formal model has deep parameters θ . The formal model for this paper appeared in Section II. But estimation of the deep parameters is challenging or intractable. Instead, one estimates the parameters of a simple descriptive model of the data. The descriptive model summarizes the relevant properties of the data. Then one takes as estimates of the deep parameters the values that cause the model to generate the same descriptive model. In most applications of this general idea, the formal model simulates data for the relevant variables, so many estimators in this framework are called simulation estimators. In my application, I calculate the factor/output response ratios from the model rather than from simulations. But the philosophy is the same. Cooper and Haltiwanger [2002] applied indirect inference to capital adjustment costs in an earlier paper with a rather different setup.

The standard rationalization for using indirect inference rather than, say, maximum likelihood, is the intractability of conventional estimators. In the application discussed here, the use of indirect inference with a focus on the contemporaneous effect of the demand innovation on the endogenous variables appears to be robust in the presence of specification errors as well. I call this the *innovation loading estimator*.

An advantage of the innovation loading estimator (ILE) over instrumental variables (IV) applied to the Euler equation arises from imposing the boundary conditions in the model. The Euler equation holds for wildly non-optimal behavior as well as for optimal behavior that satisfies the terminal conditions. Consequently, an estimator that incorporates the terminal conditions pins down parameter values more effectively than one that considers only the Euler equation. On the other hand, the ILE fails to include information from the responses of variables to lagged innovations. But the ILE performs much better in the presence of likely forms of specification error.

Figure 2 illustrates why the focus on the immediate effect of an exogenous shock is appropriate. An innovation in the demand shifter results in a delayed response of employment because of adjustment costs. With lower adjustment costs, the immediate response is correspondingly higher. The same parameter controls both the height of the curve at the left and its rate of decline from that point. Hence one can infer the adjustment cost from the height alone without losing too much. The gain is to avoid entanglement in the lower-frequency aspects of the model.

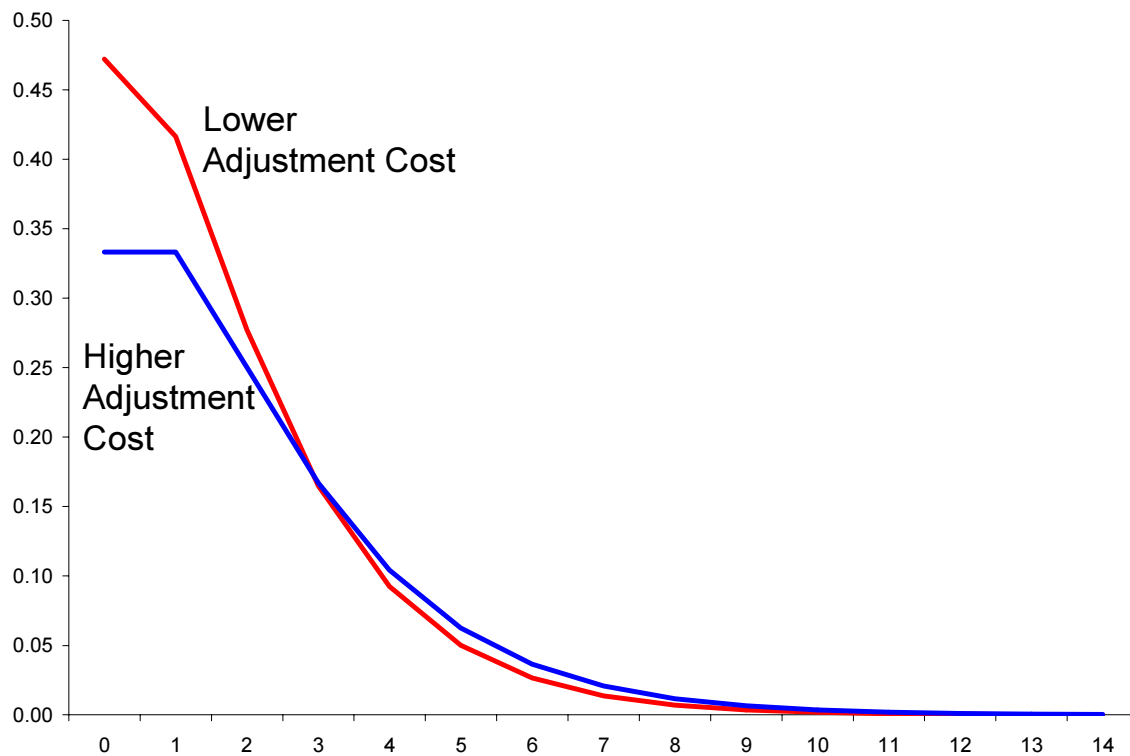


Figure 2. Innovation Response Functions of Employment to Demand Shift with Two Values for the Adjustment Cost

A third estimation strategy exploits data on the value of factors in place. The market value of installed capital, or the market value of an employment relationship with an existing worker, reveals marginal adjustment cost directly. The first-order condition for optimal levels of factor inputs with adjustment costs equates the marginal adjustment cost to the shadow value of installed factors. Tobin [1969] suggested estimating adjustment costs by reading the value of installed capital from the market value of the firm. Tobin's idea will only work with a single factor subject to adjustment costs. Thus it is disqualified from use here by my finding that both capital and labor have adjustment costs and both factors contribute rents that cause the market value of the firm to differ from the acquisition cost of the firm's capital stock.

I show later in this section that the theoretical performance of Tobin's approach could be spectacular. It is the unavailability of the appropriate data that stands in its way. If there were active markets for separate claims on installed capital and on employment relationships, Tobin's method would probably perform best among all simple estimation methods.

A. The Innovation Loading Estimator

First, I assume that the process generating the driving forces, described by f in equation (2.1), is an autoregression:

$$x_t = \rho x_{t-1} + \varepsilon_t. \quad (4.1)$$

Second, I use the following model as my descriptive model:

$$y_t = B y_{t-1} + L \varepsilon_t + v_t. \quad (4.2)$$

The vector L contains the contemporaneous innovation loadings—the immediate responses of the endogenous variables to the shocks in the exogenous driving forces. Solution of the analytical model yields theoretical values of the impulse response functions given the vector of deep parameters θ . Denote the vector of contemporaneous theoretical innovation responses by $\tilde{L}(\theta)$. The estimates of θ are those that equate the theoretical loadings $\tilde{L}(\theta)$ to the observed loadings, L .

The autoregression for the driving forces and the descriptive model form a joint system:

$$\begin{aligned} x_t &= \rho x_{t-1} + \varepsilon_t \\ y_t &= B y_{t-1} + L \varepsilon_t + v_t \end{aligned} \quad (4.3)$$

The identifying condition for this system is the orthogonality of ε_t and v_t : $E(\varepsilon_t v_t') = 0$. This condition is essentially definitional—it says that ε_t is an aggregate effect and v_t

captures effects at the industry level not contained in ε_t . Let $v_t = L\varepsilon_t + v_t$ be the composite disturbance in the descriptive model, so $v_t = v_t - L\varepsilon_t$. The identifying moment condition is $E[\varepsilon_t(v_t - \varepsilon_t L)] = 0$. Thus $L = \frac{E\varepsilon_t v_t}{E\varepsilon_t^2}$. The corresponding expression in the sample is the vector of OLS coefficients for the regressions of the residuals from the endogenous equations on the innovations from the exogenous autoregression.

In principle, efficient estimation of equations (4.3) calls for multivariate regression (seemingly unrelated regressions) with cross-equation restrictions. In my application, this is difficult to accomplish in existing software, because I use panel data for the second equation. Monte Carlo results showed that little is lost from estimating the autoregression first and then including the residuals in the second equation. In particular, the standard errors of the descriptive coefficients are actually slightly overstated by the reported standard errors. So, to estimate the parameters of the descriptive model, L , I regress the driving variable on its own lagged value to obtain the residuals ε_t . Then I regress the endogenous variables on their own lagged values and on the residuals from the equation for the driving force. The resulting regression coefficients for the residuals, L , form the descriptive model.

Let Σ be the estimated covariance matrix of L . The final step is to calculate the implied covariance matrix of the estimates of the deep parameters, $\hat{\theta}$.

$$V(\hat{\theta}) = \left(\frac{\partial \tilde{L}}{\partial \hat{\theta}} \right)^{-1} \Sigma \left(\frac{\partial \tilde{L}}{\partial \hat{\theta}} \right)'^{-1}. \quad (4.4)$$

B. Comparison to Other Estimators

To compare the innovation loading estimator to other estimators, consider the following simplified version of the adjustment cost problem. Demand is linear with unit slope, $d_t - p_t$. One unit of labor produces one unit of output, and the wage is one.

Adjustment costs are $\frac{\lambda}{2}(n_{t+1} - n_t)^2$. With zero discounting, the first-order condition for competitive profit maximization is

$$d_t - n_t - 1 = \lambda [n_t - n_{t-1} - (E_t n_{t+1} - n_t)]. \quad (4.5)$$

The left side is the marginal contribution of one worker to profit absent adjustment cost (price minus wage cost) and the right side is current less future marginal adjustment cost.

The demand shifter, d_t , is the exogenous driving force, which I take to be a random walk. I take $\lambda = 1.2$. Optimal employment then follows the process,

$$n_t = (1 - L)n_{t-1} + Ld_t. \quad (4.6)$$

Here L is the innovation loading and is one minus the smaller root of $\lambda x^2 - (1 + 2\lambda)x + \lambda = 0$, or 0.587.

Let $q_{n,t}$ be the shadow value of a worker in place—the analog for labor of Tobin's q for capital. The standard first-order condition equates marginal adjustment cost to the shadow value:

$$q_{n,t} = \lambda(n_t - n_{t-1}). \quad (4.7)$$

Thus, given employment from the solution to the optimization, the shadow value is just the adjustment cost times the rate of change of employment.

I compare three estimators for λ :

(1) *Instrumental variables* (IV) applied to

$$d_t - n_t - 1 = \lambda [n_t - n_{t-1} - (n_{t+1} - \eta_{t+1} - n_t)]. \quad (4.8)$$

where η is the innovation in n , using d_t, d_{t-1} , and n_{t-1} as instruments

(2) The *q-based estimator* (QE) using the regression

$$n_t - n_{t-1} = \frac{1}{\lambda} q_{n,t} \quad (4.9)$$

(3) The *impulse loading estimator* (ILE) using the regression

$$n_t = (2 - L)n_{t-1} + Ln_{t-2} + L\varepsilon_t \quad (4.10)$$

and solving for the λ corresponding to the estimated value of the loading coefficient, L , from $\lambda = \frac{L-2}{L^3}$.

I consider various cases involving sources of other random variation in the data apart from the demand innovation, ε_t . Equation (2.2) reveals a key difference between the IV estimator and the other two. If the demand innovation ε_t is the only source of random variation in the endogenous variables—there is no other source of randomness v_t —then there is a non-stochastic relation between the observed ε_t and the endogenous variables, and also a non-stochastic relation among the endogenous variables. Either one could be exploited to obtain an exact measure of the adjustment cost λ from a single observation. For the QE, this is equation (4.6) for any period when employment changes (a non-stochastic relation among the endogenous variables), and for the ILE, it is equation (4.10), a non-stochastic relation between employment and the observed innovation in demand. The QE and ILE obviously dominate the IV estimator in this case. Further, because the IV estimator always contains a stochastic element from the demand innovation that is absent from the other estimators even when there are other sources of variation, the IV estimator suffers in comparison to the others in those cases as well.

Table 1 reports the results of Monte Carlo experiments with 10,000 repetitions apiece. In each repetition, I draw 50 random normals and use them to calculate values for n_t from equation (4.6). The top line of the table considers estimation without any other source of random variation. In this case, the ILE and QE both give exact estimates and are not shown in the table. The average IV estimate of λ is 0.976 with a standard deviation of 0.075. The average reported standard error is 0.078, which is slightly conservative. The IV estimator suffers from a considerable downward bias even when the equation is properly specified.

<i>Additional random component</i>	<i>Estimator</i>	<i>Mean estimate of λ (true value is 1.20)</i>	<i>Standard deviation</i>	<i>Mean reported standard error</i>
None	IV	0.976	0.194	0.206
White noise in employment	IV	0.255	0.152	0.119
	ILE	1.237	0.378	0.538
Random walk in employment	QE	1.203	0.065	0.081
	IV	0.572	0.191	0.227
White noise in q	ILE	1.218	0.237	0.350
	QE	1.203	0.057	0.057
Random walk in q	QE	1.286	0.052	0.050
Random walk in q	QE	1.881	3.630	1.375

Table 1. Sampling Properties of Alternative Estimators of Adjustment Cost

Next I add white noise with a standard deviation of 0.2 to the employment numbers. The average IV estimate of λ becomes 0.255 in the face of this specification error, with a standard deviation of 0.152. The average ILE of λ is 1.237 with a standard deviation of 0.378. The IV estimator is fatally biased in the presence of the specification error, while the ILE estimator performs well. The QE performs even better, with an average estimate of λ almost exactly unbiased and a small and correctly reported standard error.

I also add a random walk to employment with an innovation standard deviation of 0.2. The average IV estimate of λ becomes 0.572 in the face of this specification error, with a standard deviation of 0.191. The average ILE for λ is 1.218 with a standard deviation of 0.237. Again, the IV estimator is seriously biased in the presence of the specification error, while the ILE continues to perform well. The QE continues to perform even better.

The last two lines of Table 1 add a random element to the shadow value of labor, q_n . The addition of white noise (again with a standard deviation of 0.2) biases the regression coefficient of q_n downward and thus the estimated value of its reciprocal, the adjustment cost λ , upward. The bias is much more severe in the case of a random walk noise element in q . Nonetheless, the bias is much too small to explain the findings of the q literature, where values of the adjustment-cost parameter at the absurd level of as high as 20 have been reported. The noise component of empirical measures of q must be huge to explain those findings.

These results confirm the finding of Garber and King's [1983] famous unpublished paper: Euler-equation estimation using lagged endogenous variables as instruments is fatally sensitive to specification error of a type likely to occur. The ILE avoids this problem completely, without becoming entangled in the complexities of maximum likelihood estimation. The QE works extremely well if good measures of q are available, though in practice this seems not to occur.

Another conclusion from this experiment is that the reported standard errors are slightly conservative estimates of the actual sampling variation of the parameter estimates, despite the two-step estimator that might appear to bias the standard errors downward.

C. Application in Panel Data

As I discussed earlier, the estimation strategy in this paper exploits the cross-sectional variation in the response of output and factor inputs to aggregate shocks. The descriptive model is

$$q_{i,t} = a_i \varepsilon_t + B_q y_{t-1} + v_{q,i,t} \quad (4.11)$$

$$n_{i,t} = c_n \varepsilon_t + b_n a_i \varepsilon_t + B_n y_{t-1} + v_{n,i,t} \quad (4.12)$$

$$k_{i,t} = c_k \varepsilon_t + b_k a_i \varepsilon_t + B_k y_{t-1} + v_{k,i,t} \quad (4.13)$$

Each industry, i , has its own innovation loading, a_i , that describes the response of its output to an aggregate shock, measured as an innovation in real GDP. The model describes the effect of the shock on employment in two ways. The component $c_n \varepsilon_t$ is a common effect across all industries. It captures effects of the shock that operate through factor prices and the cross-industry element of productivity shocks. There is no reason to expect these effects to be in proportion to the output effects. The term $b_n a_i \varepsilon_t$ captures effects of the shock that are in proportion to the output effects. The coefficient b_n is the labor/output response ratio. It describes the ratio of the effects of the innovation ε_t on labor input and output, and, similarly, the capital/output response ratio b_k describes the ratio of the effects on capital input and output.

V. Data and Estimates **Equation Section (Next)**

A. Data

I use the data compiled by Dale Jorgenson and his co-authors (Jorgenson and Stiroh [2000] and earlier papers cited there; see <http://post.economics.harvard.edu/faculty/jorgenson/data.html>). For each of 35 industries, the data report the value and quantity of output, labor, capital, materials, and energy adjusted for quality, annually from 1959 to 1999. The measure of capital is conceptually the services of capital used during the year. It

is measured as the average of the beginning and end of year stocks. I used the measure for the following year in my empirical work as the year- t measure corresponding most closely to the end-of-year measure of the model. In effect, I am allowing for a 6-month lag of investment behind its determinants, which could correspond to a 6-month period before any adjustment can occur (a time to build of 6 months).

B. Descriptive Model

The coefficient of the autoregression for real GDP is 0.991 with a standard error of 0.009. The residuals from this autoregression are the innovations taken to characterize aggregate shocks. In the rest of this section, when I refer to elasticities with respect to GDP, I mean the coefficient on the GDP innovations in an equation like (4.7).

Figures 3 and 4 show the basics of the descriptive model. The figures plot the industry GDP elasticities for output on the horizontal axis and for labor or capital on the vertical axis. That is, the horizontal position of each dot is the elasticity of output in an industry with respect to GDP and the vertical position is the elasticity of labor or capital with respect to GDP. The slopes are the response ratios. Although it would be an econometric abuse, it would not be terribly different from the procedure actually used in this work to fit regressions to these two scatter plots and call the coefficients the estimated factor/output response ratios. The plots make it clear that both ratios are unambiguously positive and that the labor response ratio is higher than the capital response ratio.




Figure 3. Cross-Sectional Output-GDP elasticities and Labor-GDP Elasticities




Figure 4. Cross-Sectional Output-GDP elasticities and Capital-GDP Elasticities

One might challenge the identification assumption underlying this approach on the grounds that there is feedback from investment shocks common across industries to GDP. Obviously GDP is not an exogenous variable—it is the sum of the values added by the industries. Figures 5 and 6 check this by repeating the elasticity calculations using federal government purchases as the aggregate shock. The movement of this variable is dominated by changes in military spending, a component unlikely to respond to other economic developments, especially investment. Although there is less variation in the output elasticities shown on the horizontal axes of these plots, the upward slopes are still quite visible.

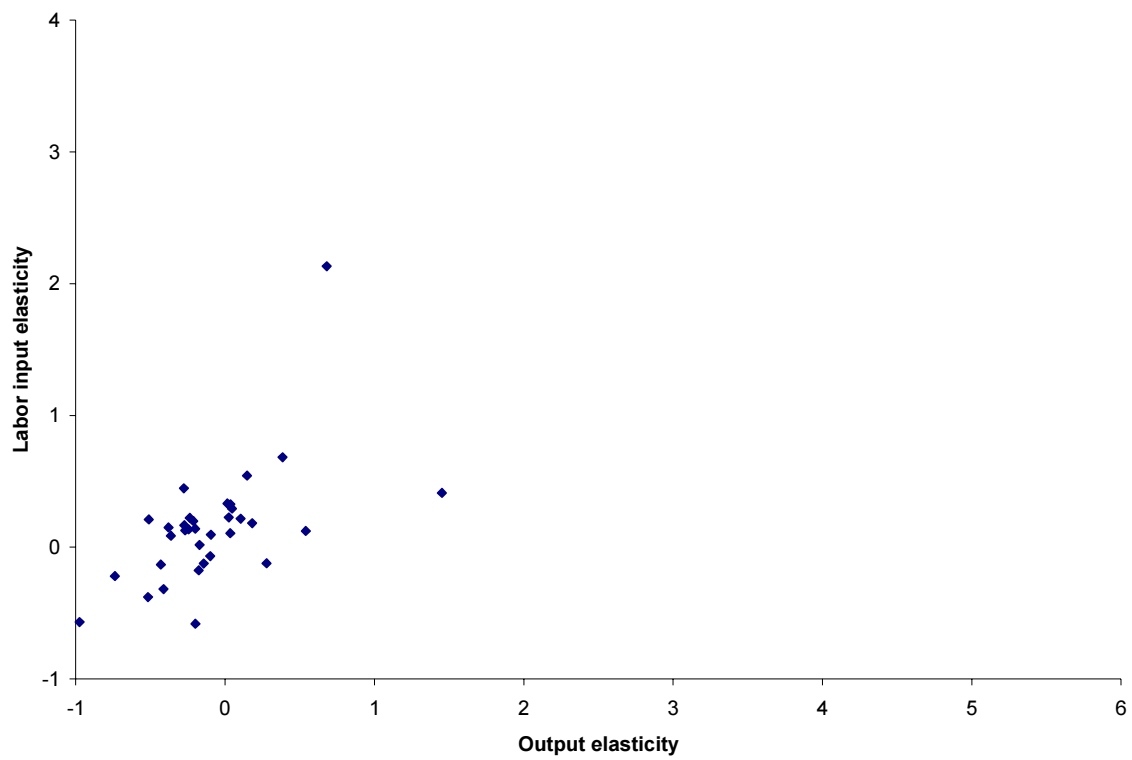


Figure 5. Cross-Sectional Output-Government Purchases Elasticities and Labor Elasticities




Figure 6. Cross-Sectional Output-Government Purchases Elasticities and Capital Elasticities

Table 2 shows the coefficients and standard errors of the descriptive model. The estimation method takes account of correlation and heteroskedasticity across industries and heteroskedasticity across output, labor input, and capital input. It imposes the nonlinear restrictions of equations (4.11), (4.12), and (4.13).

	<i>Factor-price effect, c</i>	<i>Response ratio, b</i>
Labor input	-0.01 (0.10)	0.77 (0.06)
Capital input	0.17 (0.09)	0.35 (0.05)

Table 2. Estimates for the Descriptive Model

C. Estimation of Structural Parameters from the Theoretical Model

In addition to the two adjustment-cost parameters to be estimated from the descriptive model, the industry model has 6 other parameters: two production elasticities, α and ψ , the elasticity of demand with respect to the aggregate demand shock, ω , the price elasticity of demand, δ , the survival rate of capital, s , and the real discount rate, r . For the production elasticities, I use the overall averages of the corresponding factor shares across all industries and years: $\alpha = .345$ and $\psi = 0.169$. I noted earlier that the factor/output response ratios were literally invariant to ω , so I need not specify a value for it. I take the price elasticity of demand, δ , to have the reasonable value 1—as I noted earlier, the response ratios are nearly invariant to its value. I take the survival rate for capital to be 0.9 and the discount rate to be 0.1.

With these other parameters, the resulting estimates of the adjustment-cost parameters are shown in Table 3. For labor, Shapiro [1986] finds zero adjustment cost for production workers and an adjustment cost for non-production workers somewhat above my estimate. For capital, my estimate is a little below those reported by Shapiro [1986]. 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. My estimates are well below the level suggested by Hamermesh and Pfann [1996] and well above the tiny adjustment cost for capital reported by Cooper and Haltiwanger [2002].

<i>Parameter</i>	<i>Value and standard error</i>
λ , labor adjustment cost	1.22 (.41)
γ , capital adjustment cost	1.54 (.48)

Table 3. Estimates for the Adjustment-Cost Parameters

In addition to the Euler-equation approach, as I noted earlier, many authors have pursued James 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.

VI. Implications for Industry Dynamics

Figure 7 shows some of the properties of the model with the estimated adjustment costs. It displays the responses over an 8-year period to a permanent increase in demand (not the slightly decaying increase typical of a shock to real GDP). The figure shows industry equilibrium with factor prices that do not respond to the demand shock. The elasticity of the immediate response of labor exceeds that of the capital stock because labor's adjustment cost is lower. The elasticity of the response of output is greater than the elasticity for labor or capital because materials respond with an elasticity greater than one.

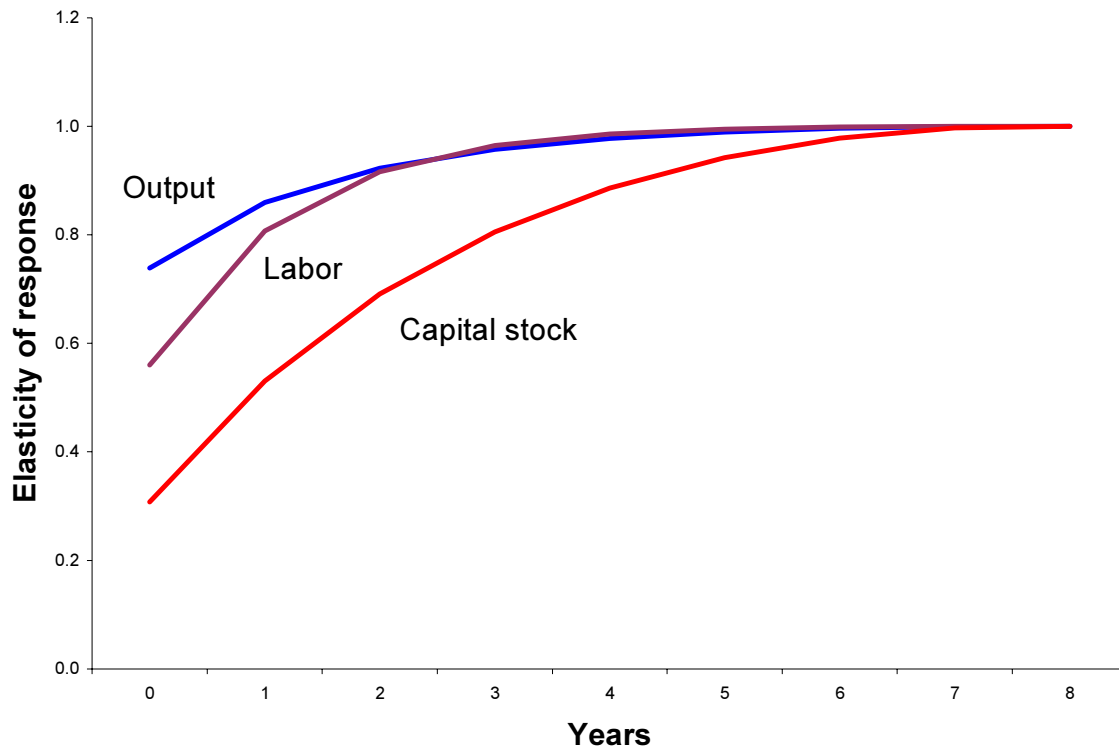


Figure 7. Responses of Output, Labor, and Capital to an Innovation in Demand When Demand is a Random Walk

As discussed in Hall [2001], part of the fluctuations in the value of firms recorded in securities markets comes from capitalized rents associated with adjustment costs. Because these costs make the short-run supply curve of the firm slope upward, a force that moves industry equilibrium up the supply curve creates rents. The dynamic model describes the transition from the short run, where rents are greatest, to the long run, where the rents disappear. Figure 8 shows the total capitalized rent and its labor and capital components resulting from a permanent change in demand. The vertical position of the curves is the percent change in the value of the firm in response to a one-percent innovation in demand.

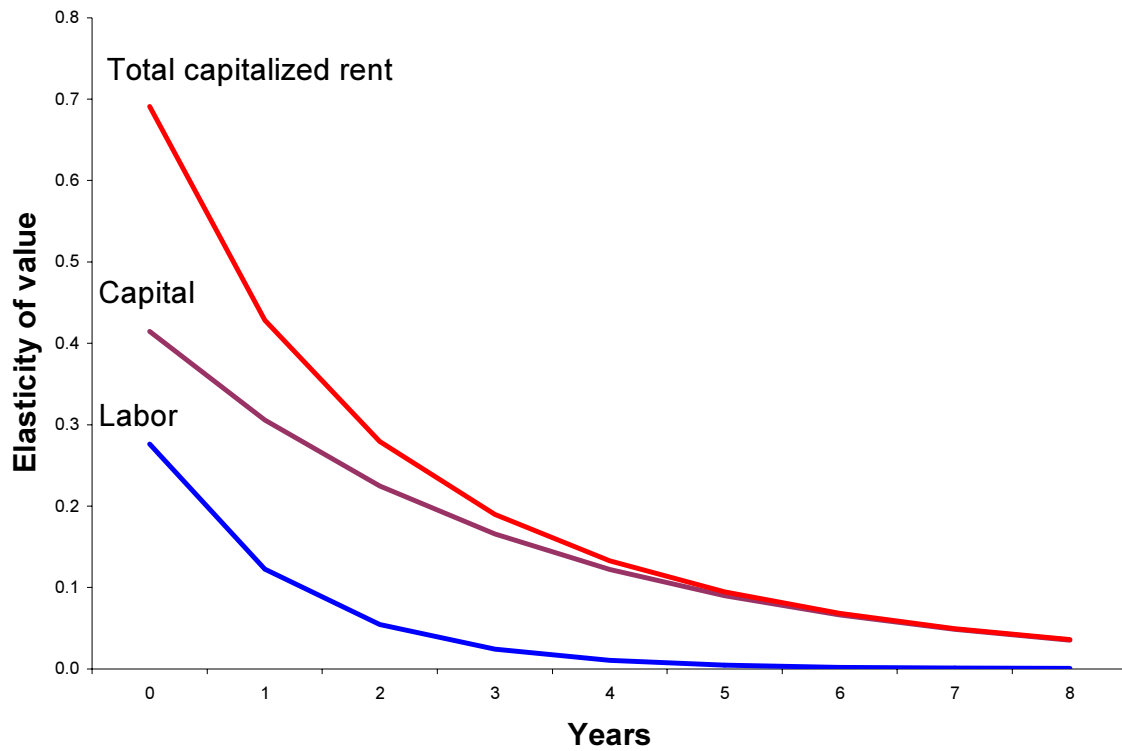


Figure 8. Value of Rents Earned as a Result of a Permanent Increase in Demand: Elasticity of the Value of the Firm with Respect to the Demand Shock

The value of capitalized rent from labor is smaller at all times, and decays faster, because labor’s adjustment cost is smaller. Although rents on the labor side are too big to be ignored, it seems unlikely that they could overturn the conclusion in Hall [2001] that rents from temporary fixity of factors are a small part of the story of the movements of the stock market.

VII. Comparison of Actual and Predicted Movements of the Endogenous Variables

This section explores the performance of the model under the hypothesis that four observed driving forces—the industry demand shift, industry productivity, wages, and materials prices—are the only driving forces. I solve the model for its predicted innovations in the endogenous variables and compare them to the actual innovations. A finding that the two sets of innovations are essentially identical would confirm that I had selected a full set of driving forces and that there was no randomness arising within the model. Of course, there actually will be residuals suggesting other sources of randomness. A weaker finding that the residuals were uncorrelated with the predicted innovations would be interesting. Such a finding would support a GMM approach to estimation, treating the driving forces as econometrically exogenous.

I calculated time series for the driving forces in each industry as follows: I calculated time series for the industry demand shift as the residual from the demand function with price elasticity of unity. I calculated a time series for productivity as the standard cumulated Solow residual. I took the wage data directly from Jorgenson's data. I aggregated energy and materials input into a single input using Divisia aggregation.

For this exercise, I need to extend the model to include a stochastic model of the driving forces. Following King and Watson [1996], I embed a VAR for the driving forces in the model. Decision makers in the model form expectations of future values of the driving forces from the VAR. I estimate the VAR in the panel data, with the same coefficients on lagged variables for all industries, but with industry-specific intercepts. I use a single lag in the VAR based on preliminary findings that two-year lags added nothing. The VAR is close to a random walk separately for each of the four variables—its off-diagonal elements are tiny. I then solve the model containing the VAR for its innovation loading matrix that shows how the innovations in the exogenous variables drive the innovations in the endogenous variables. The matrix appears in Table 4.

<i>Endogenous variable</i>	<i>Driving force</i>			
	<i>Demand</i>	<i>Productivity</i>	<i>Wage</i>	<i>Materials price</i>
Output	0.71	1.01	-0.17	-0.49
Price	0.29	-1.01	0.17	0.49
Labor input	0.53	0.02	-0.50	0.00
Capital	0.24	0.02	0.02	-0.02
Materials	1.00	0.00	0.00	-1.00

Table 4. Innovation Loading Matrix from the Model

The effects of the industry demand shock in the first column are similar to those underlying the earlier results (there, I divided the labor and capital coefficients by the output coefficient). The productivity shock raises output almost in exact proportion and lowers price in the same proportion (the coefficient would be exactly 1, given unit-elastic demand, except for small cross effects in the VAR). The wage shock lowers output and raises price in the same proportion, and lowers employment with an elasticity of a half. The employment effect is constrained by the employment adjustment cost and the output effect by both adjustment costs. Finally, the materials price shock lowers materials inputs with an elasticity of one and lowers output and raises price with an elasticity of about a half. Again, the output and price effects are held back by adjustment costs for other factors.

To get the predicted innovations in the endogenous variables, I compute the product of the innovation loading matrix and the innovations from the VAR for the exogenous variables. I compare these to the innovations from a VAR fitted to the endogenous variables, again with a single lag. Table 5 describes in relation between the two in terms of the regression coefficients of the actual innovations on the predicted innovations and the correlations of the two series.

<i>Endogenous variable</i>	<i>Regression coefficient and standard error</i>	<i>Correlation</i>
Output	0.500 (0.010)	0.806
Price	0.662 (0.031)	0.501
Labor	0.477 (0.023)	0.493
Capital	0.211 (0.033)	0.167
Materials	0.854 (0.022)	0.727

Table 5. Coefficients of Regressions of Predicted and Actual Innovations in Endogenous Variables, and Correlations

These coefficients are not invariant to the elasticity of demand, but the basic story holds for values of the elasticity of 0.5 and 2 as well as the value of 1 used in deriving Tables 4 and 5. The correlations in the right-hand column speak to the issue of excluded driving forces or endogenous randomness, including errors in measuring the endogenous variables. For output and materials input, the correlations are quite high. For capital, on the other hand, the correlation is low. The model omits variations in the financial determinants of capital, though, because the investment literature has generally been skeptical of their quantitative importance, it seems unlikely that adding them would raise the correlation much. More likely is randomness in investment not captured by the model at all.

If the orthogonality property mentioned earlier held, the regression coefficients in Table 5 would all be 1. In fact, all of the coefficients are below 1—orthogonality is rejected in all cases. GMM estimation treating the driving forces as econometrically exogenous would probably not be a good idea, and I found that some experiments in this area gave perverse results. The shortfall in the coefficients arises from two potential sources—errors in measuring the driving forces and correlation of the driving forces with

the errors in measuring the endogenous variables and in the unexplained randomness of the endogenous variables.

Figure 9 shows the scatter plot of the predicted and actual innovations in output for the entire sample. The plot seems entirely consistent with the view that both measure the same underlying variable, but both contain some errors so that the points do not lie perfectly on the 45-degree line. Figure 10 shows the same plot for capital. The same property holds, but it is clear that most of the variation comes from noise and it is difficult to see much tendency to group on the 45-degree line.




Figure 9. Predicted Output Innovation (Horizontal Axis) and Actual (Vertical Axis)




Figure 10. Predicted Capital Innovation (Horizontal Axis) and Actual (Vertical Axis)

VIII. Concluding Remarks

Discrete adjustment costs appear to have little role in industry dynamics despite their large role in plant dynamics, because the plants that do adjust by large amounts to a shock make up for the plants that do not adjust at all. Industry dynamics are controlled by convex adjustment costs, which lead all plants to spread their adjustment over time to exploit the convexity. Firms earn rents from the upward slope that adjustment costs impart to their supply curves. A model of industry equilibrium traces the response from the short run to the long run, as each supply curve becomes more elastic. 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 flows of new hires 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.

Using the intuitively appealing proposition that surprises in demand should stimulate the use of factors with low adjustment costs by more than it stimulates those with high adjustment costs, I am able to measure adjustment costs separately for labor and capital. I use the model to infer the parameters of adjustment cost from the observed responses, using the econometric strategy of indirect inference. The resulting parameter estimates suggest moderate adjustment costs, at the lower end of those estimated in most previous work.

References

- Caballero, Ricardo J. 1999. "Aggregate Investment" Chapter 12 in John B. Taylor and Michael Woodford (eds.) *Handbook of Macroeconomics* (New York: Elsevier). Pp. 816-862.
- Caballero, Ricardo J., and Eduardo M.R.A. Engel. 1999. "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach" *Econometrica* 67: 783-826. No. 4. July.
- Caplin, Andrew S., and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money" *Quarterly Journal of Economics* 102: 703-726. No. 4. November.
- Cooper, Russell W., and John C. Haltiwanger. 2002. "On the Nature of Capital Adjustment Costs" unpublished. January.
- Garber, Peter M., and King, Robert G. 1983. "Deep Structural Excavation? A Critique of Euler Equation Methods" NBER Technical Working Paper 31. November.
- Gilchrist, Simon, and Charles P. Himmelberg. 1995. "Evidence on the Role of Cash Flow for Investment" *Journal of Monetary Economics* 36: 541-572. No. 3. December.
- Gouriéroux, Christian, Alan Montfort, and E. Renault. 1993. "Indirect Inference" *Journal of Applied Econometrics* 8: S85-118. Supplement. December
- Hall, Robert E. 2001. "The Stock Market and Capital Accumulation" *American Economic Review* 91: 1185-1202. December
- Hamermesh, Daniel S., and Pfann, Gerard A. "Adjustment Costs in Factor Demand." *Journal of Economic Literature*, September 1996, 34(3), pp. 1264-1292.
- Jorgenson, Dale W., and Kevin J. Stiroh. 2000. "Raising the Speed Limit: U.S. Economic Growth in the Information Age." *Brookings Papers on Economic Activity* No. 1, pp. 125-211.
- King, Robert G., and Mark W. Watson. 1996. "Money, Prices, Interest Rates, and the Business Cycle" *Review of Economics and Statistics* 78(1): 35-53, February.
- Shapiro, Matthew D. "The Dynamic Demand for Capital and Labor," *Quarterly Journal of Economics*, August 1986, 101(3), pp. 513-542.
- Smith, Anthony A., Jr. 1993. "Estimating Nonlinear Time Series Models Using Simulated Vector Autoregressions" *Journal of Applied Econometrics* 8:63-84.

Tobin, James. 1969. "A General Equilibrium Approach to Monetary Theory" *Journal of Money, Credit, and Banking*. 1: 15-29. February.

Thomas, Julia K. 2001. "Is Lumpy Investment Relevant for the Business Cycle" University of Minnesota, unpublished. May