The Response of Prices to Shifts in Demand

Robert E. Hall

Hoover Institution and Department of Economics
Stanford University

National Bureau of Economic Research

Hall@Hoover.Stanford.edu

March 11, 2002

Abstract:

An econometric model of industry dynamics implies a slope of the short-run industry supply function of about 0.41. The model uses data on factor inputs and industry output, not data on prices. I estimate a parameter measuring the ratio of the actual response of prices to demand shocks to the slope of the supply function. The estimate is slightly *negative* and the hypothesis that price and output move along the industry supply function in response to demand shocks is resoundingly rejected in favor of no movement of price at all. The finding is repeated in an alternative body of data on industry prices. In addition, I check that productivity shifts are insufficiently correlated with output changes to explain much of the finding. I show that the expected response of prices is not deferred, but rather never occurs at all.

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. Introduction

When the demand function for an industry's output shifts upward, output and prices rise. This proposition has few exceptions in standard economic thinking. For a competitive industry, the only exception is perfectly elastic supply, a borderline case arising only if all factors are freely adjustable and in perfectly elastic supply to the industry. For an industry where price is a markup over marginal cost, it is conceivable that an upward shift in demand could trigger a sufficiently large reduction in the markup ratio to keep prices stable, but models with this property have never penetrated the outer limits of plausibility.

Empirical research on prices has found demand effects elusive. A constant markup over average cost has always been a hard model to beat—see Rotemberg and Woodford [1999]. However, research in this area has remained unpersuasive because it is challenging to isolate exogenous demand shifts from other sources of variation. For example, if the forces perturbing output and prices include productivity shifts as well as demand shifts, the correlation of output with prices will have a negative element from the productivity shifts as well as a positive element from the demand shifts and the two may easily cancel each other. The question calls for something more sophisticated than a regression of price on output.

A closely related question of interest is the timing of the response of prices to demand changes. Models contemplated by macroeconomists have incorporated lags in price responses. In the simplest models, sellers freeze their prices for a designated period—say a year—and thus may wait for up to that long to adjust prices in response to demand shifts. In a currently popular model, the waiting period has random rather fixed length.

The research in this paper makes use of a panel of 35 industries covering the entire U.S. economy from 1958 through 1999, compiled by Dale Jorgenson and his co-authors (Jorgenson and Stiroh [2000]). Figure 1 displays the entire body of data as a scatter plot of rates of change of output (horizontal axis) and prices (vertical axis). Notice that price volatility is substantial. Prices are far from rigid. But the correlation of output and price changes is essentially zero. Either productivity shifts offset demand shifts or demand shifts do not affect prices.

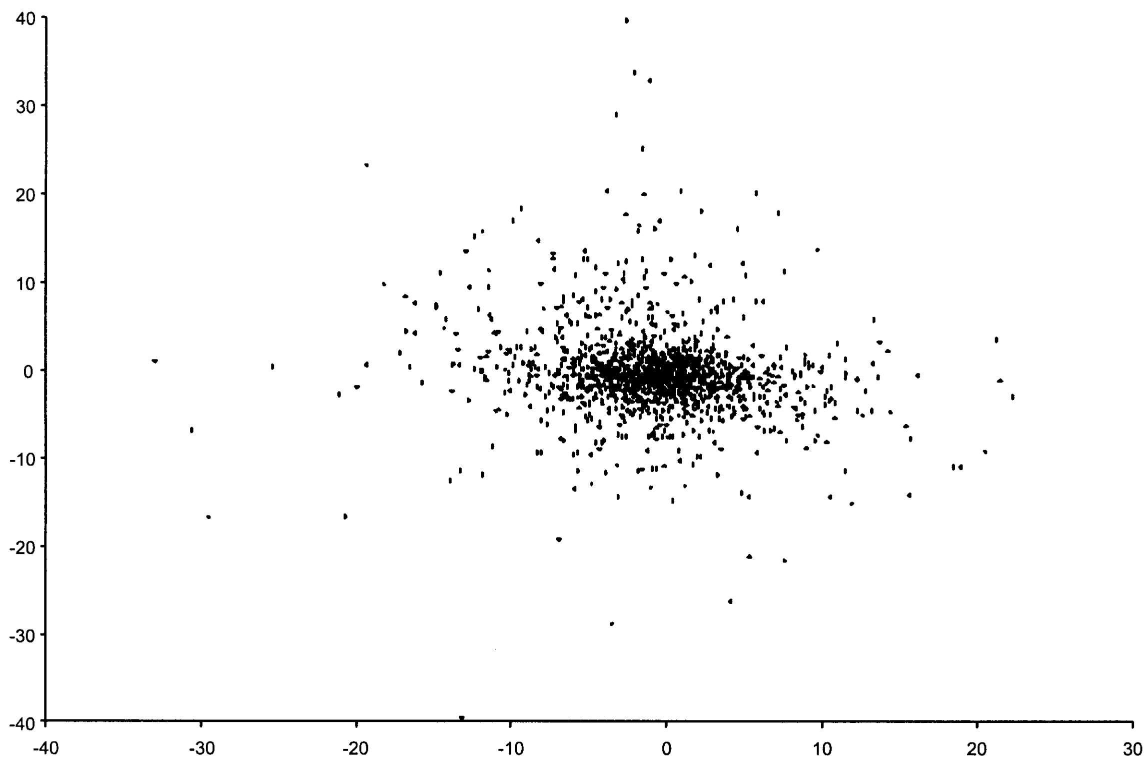

Figure 1. Output changes (horizontal axis) and price changes (vertical axis) in 35 U.S. industries, 1958-1999, in percent.

To measure the response of prices to shifts in demand, I make use of an identifying strategy from Hall [2002]. I focus on the cross-sectional response of the 35 industries to aggregate shocks, measured as time-series innovations in real GDP. These shocks result from forces that have various effects on each industry. Some operate through factor prices and others through product demand. I also consider the possibility that the shocks originate, in part, from shifts in productivity. I hypothesize that, *across industries*, the factor-price and productivity shocks or effects are not correlated with the demand shocks or effects. In other words, industries whose demand is perturbed more strongly by aggregate shocks do not face larger or smaller factor-price or productivity shocks associated with the aggregate shock. The logic is that aggregate factor-price effects—such as higher interest rates—affect industries in proportion to factor intensities, not in proportion to the magnitude of their demand effects. The strategy is successful because of huge cross-industry variation in the elasticity of demand shifts with respect to real GDP innovations. I am able to provide confirming evidence about the assumption that the correlation of industry productivity with real GDP innovations is not related across industries to the correlation of industry output with real GDP innovations.

The basic finding of the paper is striking. Rather than track an industry's short-run supply function, prices actually fall slightly, in relation to prices in general, when a shift in demand raises output. The confidence interval around this finding is tight. The hypothesis that prices move along the supply function is overwhelmingly rejected. I check whether the reason is that the positive effects are delayed, as in sticky-price models, but reject that view definitively as well. I also show that the finding does not stem from favorable shifts in productivity that coincide with the demand shifts.

II. Model

The approach in this paper is to build a standard model of industry dynamics. I do not attempt to incorporate any special factors such as costs of adjusting prices that might provide an explanation for a finding of limited price response to demand shocks. I use an estimation strategy that determines the slope of the short-run supply function without reference to data on output prices. I then compare the actual price-quantity movements induced by demand shifts to those implied by the model's dynamic supply function.

The key issue in the response of price to shifts in demand is the slope of the industry's short-run supply function. The model invokes adjustment costs in capital and labor as determinants of that slope. As I estimate the model, adjustment costs also characterize short-term inelasticities in factor supply. Hall [2002] explains the model in detail. Briefly, it comprises first a production technology, relating labor input, n_t, materials input, m_t, and capital input, k_t, to output, q_t:

$$q_t = A_t n_t^{\alpha} m_t^{\psi} k_t^{1-\alpha-\psi} . \tag{2.1}$$

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

$$q_t = d_t z_t^{\omega} p_t^{-\delta} ; \tag{2.2}$$

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. Adjustment costs are convex in the inputs:

$$\frac{\lambda}{2} w_{t+1} \frac{(n_{t+1} - n_t)^2}{n_t} + \frac{\gamma}{2} w_{t+1} \frac{(k_{t+1} - k_t)^2}{k_t} . \tag{2.3}$$

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 Hall [2002] that discrete costs have little role in industry dynamics despite their essential role in understanding plant-level dynamics.

Finally, I model the aggregate shift, z_t, as a first-order autoregression, with innovation ε_t.

I solve for industry equilibrium as a function of the aggregate shift in the product demand function. Specifically, I find the derivatives of the equilibrium with respect to the innovation in the aggregate shift (see the computational appendix to Hall [2002]). I concentrate on the immediate effect on output and price. In the case of price, the immediate effect is the largest, because the supply function is least elastic in the short run. I calculate the immediate elasticities of output, employment, capital, and price with respect to the innovation in the aggregate shift. I call these the *innovation loadings*. Let $f_n(\lambda,\gamma)$, $f_k(\lambda,\gamma)$, and $f_p(\lambda,\gamma)$ be the contemporaneous responses predicted by the model of the ratio of the changes in the logs of employment, capital, and price to the change in output induced by an innovation. Further, let π be the ratio of the actual price response ratio to the theoretical ratio. I express the model in the form,

$$\tilde{L}(\theta)=\begin{bmatrix} f_n(\lambda,\gamma) \\ f_k(\lambda,\gamma) \\ \pi f_p(\lambda,\gamma) \end{bmatrix}. \tag{2.4}$$

Here θ is the vector of the three parameters: the two adjustment costs, λ and γ and the price-effect ratio, π. Notice that this setup achieves the goal set out at the beginning of this section. Factor adjustment costs are estimated purely from the behavior of output and the factors and are not affected by price behavior. The parameter π simply records the relation between the model's prediction about price movements and actual movements.

III. Econometrics

I use the strategy of indirect inference for estimation—see Smith [1993] and Gouriéroux, Montfort, and Renault [1993] for general discussion and Hall [2002] for details about the application to the panel data used here. The strategy is to estimate the vector L of the contemporaneous innovation loadings—the immediate responses of the endogenous variables to the shocks in the exogenous driving forces. I obtain these through a simple procedure described below. The analytical model yields theoretical values of innovation loadings, $\tilde{L}(\theta)$ as described in the preceding section The estimates of θ are those that equate the theoretical loadings $\tilde{L}(\theta)$ to the observed loadings, L—to find estimates of the three parameters, I solve $\tilde{L}(\theta) = L$. 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, $\hat{\theta}$.

$$V(\hat{\theta}) = \left(\frac{\partial \tilde{L}}{\partial \hat{\theta}}\right)^{-1} \Sigma \left(\frac{\partial \tilde{L}}{\partial \hat{\theta}}\right)^{\prime -1}. \tag{3.1}$$

A. Application in Panel Data

The estimation strategy in this paper exploits the cross-sectional variation in the response of output and factor inputs to aggregate shocks. My earlier paper discusses the framework more fully. The descriptive model is

$$q_{i,t} = a_i \varepsilon_t + B_q y_{i,t-1} + v_{q,i,t} \tag{3.2}$$

$$n_{i,t} = c_n \varepsilon_t + b_n a_i \varepsilon_t + B_n y_{i,t-1} + v_{n,i,t} \tag{3.3}$$

$$k_{i,t} = c_k \varepsilon_t + b_k a_i \varepsilon_t + B_k y_{i,t-1} + v_{k,i,t} \quad (3.4)$$

$$p_{i,t} = c_p \varepsilon_t + b_p a_i \varepsilon_t + B_p y_{i,t-1} + v_{p,i,t} \quad (3.5)$$

Here y is the vector of the four variables, q, n, k, and p. The row vectors $B_\bullet$ make up a vector autoregression for the four variables. Each industry, i, has its own output innovation loading, a_i, that describes the response of its output to an aggregate shock, measured as the innovation in real GDP, ε. The model describes the effect of the shock on price (or labor or capital) in two ways. The component $c_p \varepsilon_t$ 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 $b_p a_i \varepsilon_t$ captures effects of the shock that are in proportion to the output effects. The coefficient b_p is the price/output response ratio. It describes the ratio of the effects of the innovation ε_t on price and output, and, similarly, the labor/output and capital/output response ratio b_n and b_k describes the ratio of the effects on labor and capital input relative to output.

IV. Evidence

Hall [2002] describes the data and econometric method used here. Briefly, the Jorgenson data form a panel on 35 industries covering the entire U.S. economy (including government enterprises but not government itself) for the years 1958 through 1999. The procedure for extracting responses to innovations in driving forces is shown to be reasonably efficient and highly robust in my earlier paper. I estimate the parameters of equations (3.2) through (3.5) by nonlinear multivariate regression, with consideration of correlation and heteroskedasticity across industries and heteroskedasticity across output, labor input, capital input, and price.

Figures 2 and 3 show the basic idea of the estimation method. Figure 2 is a scatter plot of estimated elasticities of output with respect to GDP, on the horizontal axis, and labor, on the vertical axis. Quite reliably, industries with higher output/GDP elasticities have higher labor/GDP elasticities. But the slope is less than one. Labor adjustment is held back by adjustment costs. My earlier paper showed how the slope reveals the adjustment cost, though the full procedure considering both capital and labor adjustment costs requires simultaneous consideration of the labor and capital responses because they interact with each other. This paper adopts the same approach to adjustment costs for labor and capital.

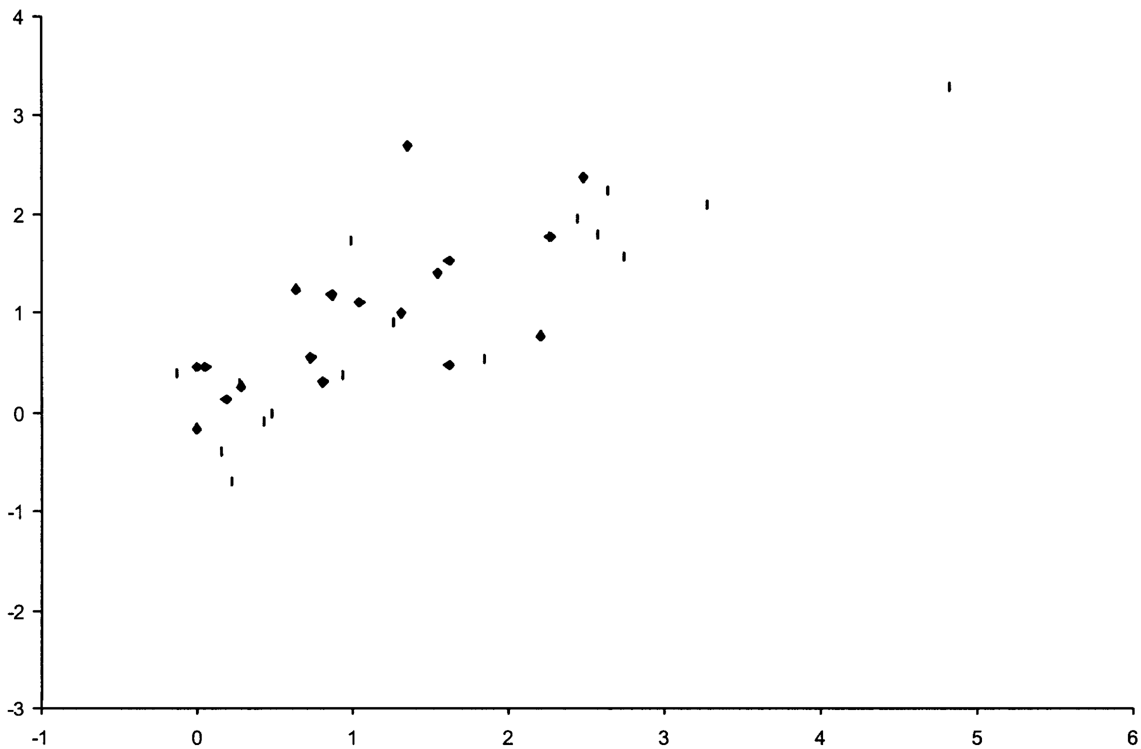

Figure 2. Cross-Sectional Output-GDP (Horizontal Axis) Elasticities and Labor-GDP Elasticities (Vertical Axis)

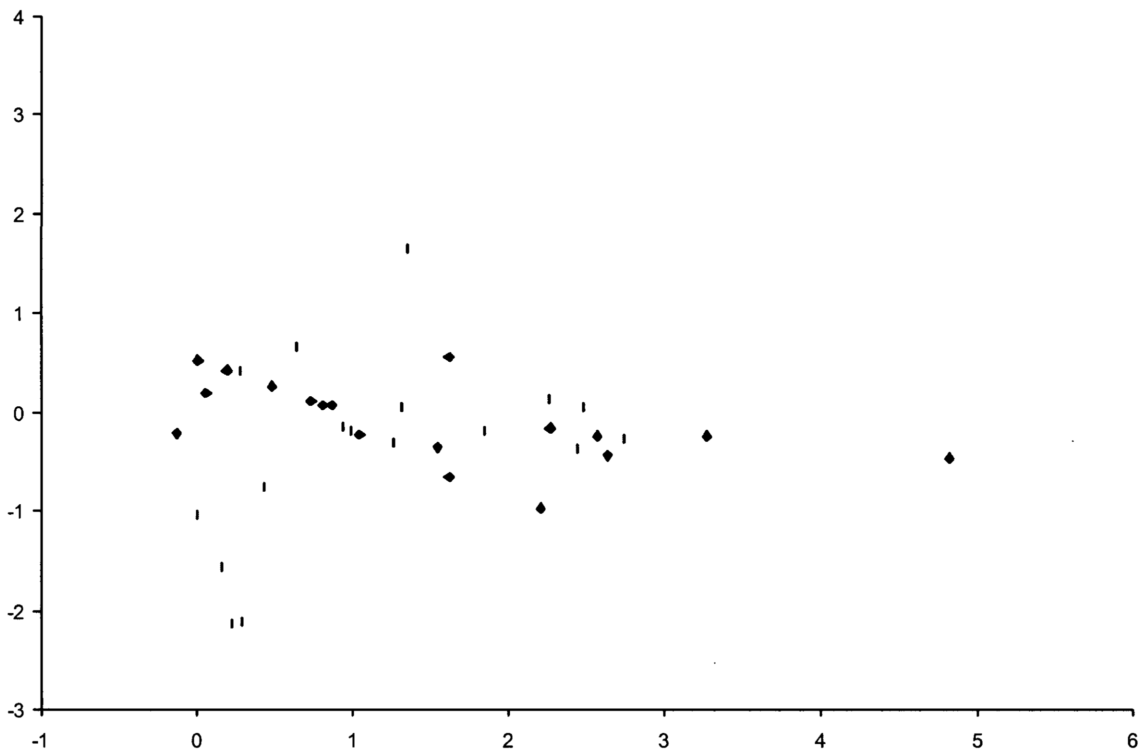

Figure 3. Cross-Sectional Output-GDP (Horizontal Axs) Elasticities and Price-GDP Elasticities (Vertical Axis)

Figure 3 carries out the same exercise for prices. The plot is essentially flat. If I attributed an adjustment cost to prices, its value would be infinity. I infer that prices do not respond to changes in demand.

Table 2 shows the coefficients and standard errors of the descriptive model. For labor input, the GDP innovation has a slight negative uniform effect across industries, interpreted as the effect on employment in an industry where GDP has no effect on output. The slope coefficient of 0.796 is the ratio of the effect of GDP on employment to its effect on output. I interpret this as a pure demand effect. This coefficient is almost the same as reported in my earlier paper. The similar coefficient for capital, 0.241, is a little above the value in that paper, indicating somewhat higher capital adjustment costs than inferred there. The results differ because of the inclusion of the price in equation (3.5), which is not present in the earlier paper.

	Intercept: Factor-price effect, c	*Slope: Response ratio, b*
Labor input	-0.012 (0.082)	0.796 (0.055)
Capital input	0.095 (0.040)	0.242 (0.027)
Price	0.023 (0.065)	-0.060 (0.043)

Table 2. Estimates for the Descriptive Model

The key finding in Table 2 is in the row for the price. The model predicts a negative intercept and a positive slope of about 0.4. With a positive slope, the intercept must be negative because the average effect on price must be zero—the industry prices are stated as ratios to the GDP deflator, which is essentially the average of the prices. In fact, both the intercept and slope are close to zero. Prices do not move the way the model predicts they should move in response to demand shocks. Of course, one possible explanation is that my estimation strategy has not isolated demand shocks. I will say more about this shortly.

A. 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 $\psi = 0.169$. The factor/output response ratios are 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—the response ratios are nearly invariant to its value, as shown in my earlier paper. I take the quarterly survival rate for capital to be 0.975 and the discount rate to be 0.025. I interpret the sum of the first four quarterly effects as the annual effect captured in the descriptive model.

With these other parameters, the resulting estimates of the adjustment-cost parameters are shown in Table 3. The adjustment costs in the table are about the same as found in my earlier work.

Parameter	*Value and standard error*
λ, labor adjustment cost	1.20
	(0.66)
γ, capital adjustment cost	3.20
	(0.91)
π, ratio of actual to theoretical price effect	-0.18
	(0.11)

Table 3. Estimates for the Adjustment-Cost Parameters and Price-Effect Ratio

According to the model, the price elasticity of the short-run industry supply function is 1/0.41. For each percent of output increase caused by a shift in demand, price should rise by 0.41 percent. In fact, the price *falls* slightly, by -0.06 percent, according to the results in Table 2. The estimated value of π is the ratio, -0.18. The standard error of this estimate, 0.11, is sufficiently small to rule out any but small positive values for the price response. The hypothesis that price-output movements track the estimated short-run supply function ($\pi = 1$) is overwhelmingly rejected, with a *t*-statistic of 1.18/0.11 = 11.2.

B. Estimation with Data from the Producer Price Index

I repeated the estimation of equation (2.9) of the descriptive model using data from the Bureau of Labor Statistics' Producer Price Index. I spliced data from the earlier commodity series to the more recent industry series to obtain the longest possible series for the available industries. The appendix shows the sources and years available. I was able to obtain at least some data for 27 of Jorgenson's 35 industries. In all, I had 1186 observations, after losses for lagged observations and to span the transition from commodity to industry prices, covering the years 1947 through 2001. I used the estimated elasticities of output with respect to the aggregate shock (the a_is from equations (2.6) through (2.9)) from the estimates underlying Table 2. Table 4 shows the estimates of the intercept and slope of the cross-sectional relation between the price response to the aggregate shock and the output response. Again, the slope is slightly negative and the hypothesis that it equals the model's prediction of 0.45 is strongly rejected.

	Intercept: Factor-price effect, c	*Slope: Price/output response ratio, b*
Price	0.133	-0.036
	(0.121)	(0.063)

Table 4. Estimates of the Descriptive Model Using Data from the Producers Price Index

V. Productivity

One explanation of the findings of this paper rooted in neoclassical economics is the following: Aggregate shocks captured by the innovation in real GDP are partly productivity shocks. Since higher productivity implies a lower price, the finding of unresponsive prices could reflect a sufficient role of productivity shocks to keep prices approximately constant in the fact of the typical aggregate shock.

Table 5 checks this possibility by examining the behavior of productivity in the same framework as Table 4. The left-hand variable is the standard Solow residual, cumulated into an index of productivity. I estimate in a simplified univariate framework where I take the estimates of α_i as given from the earlier estimation of the entire system, equations (3.2) through (3.5). I estimate

$$A_{i,t} = c_A \varepsilon_t + b_A a_i \varepsilon_t + B_{A,A} A_{i,t-1} + v_{A,i,t} \tag{5.1}$$

The intercept effect, as in the earlier results, is the effect of an aggregate shock on productivity in an industry where the aggregate shock has no effect on output. The slope measures the effect of a shock in proportion to its effect on the industry's output. It measures the effect relevant for the hypothesis that productivity effects are responsible for the lack of response of prices to shocks that change output.

	Intercept: Factor-price effect, c	*Slope: Productivity/output response ratio, b*
Productivity	0.147	0.091
	(0.064)	(0.038)

Table 5. Estimates of Descriptive Model for Productivity

There is statistically unambiguous evidence of a small positive productivity effect. The elasticity of industry productivity with respect to an aggregate shock is 0.091 of the elasticity of the industry's output with respect to the shock. This small effect is enough to offset the *negative* coefficient found in Table 2 and to push the negative coefficient in Table 3 into positive territory, but nowhere near enough to place the coefficient where it belongs in terms of the theoretical model. The simple neoclassical explanation does not work.

VI. Lagged Responses of Prices

Sticky-price models may imply that prices respond with a lag to changes in demand. It is straightforward to check this hypothesis by including lagged innovations in real GDP in equation (3.5). In the same univariate framework as in the previous section, I estimate

$$p_{i,t} = c_{p,0}\varepsilon_t + b_{p,0}a_i\varepsilon_t + c_{p,1}\varepsilon_{t-1} + b_{p,0}a_i\varepsilon_{t-1} + c_{p,2}\varepsilon_{t-2} + b_{p,2}a_i\varepsilon_{t-2} + B_{p,p}p_{i,t-1} + v_{p,i,t}$$

(6.1)

Lag	*Intercept,* $c_{p,\cdot}$	*Slope,* $b_{p,\cdot}$
0	-0.401	0.100
	(0.128)	(0.075)
1	0.170	-0.017
	(0.125)	(0.074)
2	0.013	0.019
	(0.124)	(0.073)

Table 6. Estimates of Lagged Responses of Prices to Aggregate Innovations

The contemporaneous response is now slightly positive (in the simplified univariate setup, it is also slightly positive without the two lagged innovations). The two lagged slope coefficients are both essentially zero. There is no support for the view that the response has the expected magnitude but is delayed. Rather, the small response occurs within the same year as the aggregate shock.

These findings do not resemble the impulse response function for prices predicted by the model, shown in Figure 4.

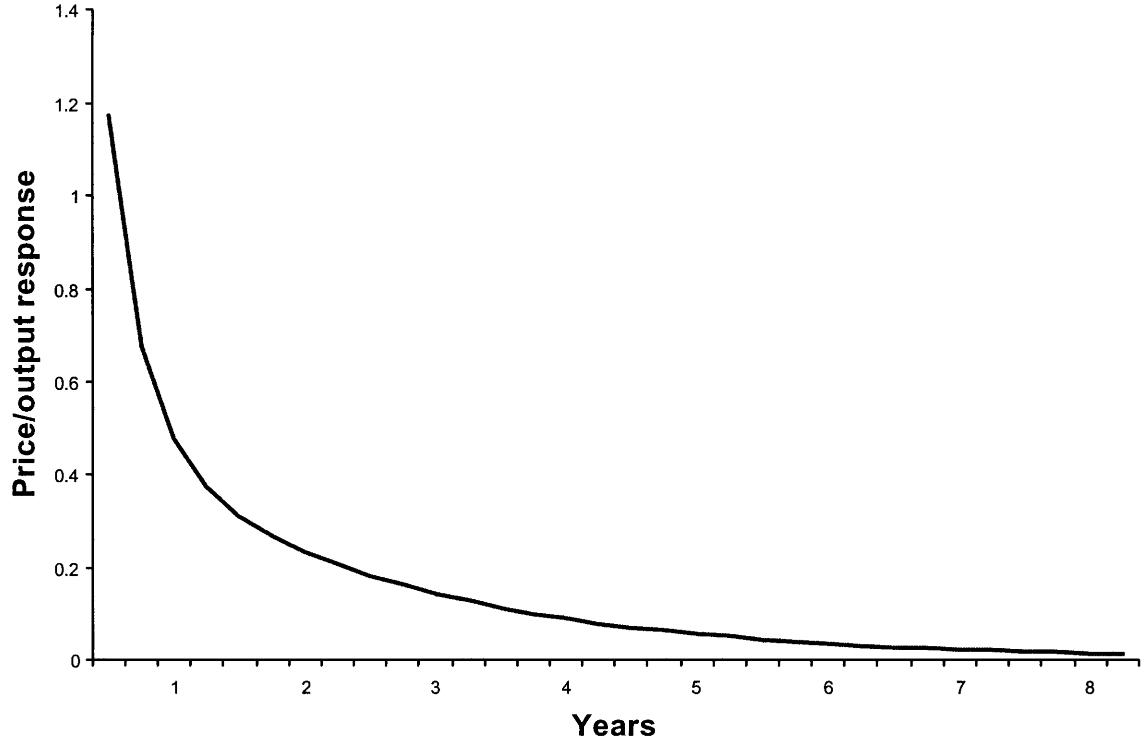

Figure 4. Price/output response ratio implied by estimated adjustment costs

VII. Conclusions

At the industry level, prices are far from rigid. But the evidence is fairly strong that the average price of an industry's products does not move along the industry's short-run supply function as demand shifts. This conclusion follows from an identification scheme that isolates industry-level demand shifts induced by aggregate shocks and uses these shifts to find relatively precise and reasonable estimates of adjustment costs. The same demand shifts have no effects on the industry's price relative to prices in general. A parameter measuring the size of the actual effect of demand on price as a fraction of the theoretical effect has a value slightly below zero, with a relatively small standard error.

I have avoided developing any theory to explain this striking finding. The literature on this type of theory is so extensive and rich that it would be impossible to begin to do it justice in this paper. I do not delve into theories to explain this finding. Rotemberg and Woodford [1999] survey models that invoke declines in price/cost markups that offset changes in cost. In these models, the same forces that increase demand also make markets more competitive. Another line of thought holds that published prices are not true allocational prices, but are installment payments within long-term relationships—see Hall [1980]. Quantity is determined efficiently by some method other than an open market. A finding that price does not track the short-run supply curve confirms the absence of the open market but does not have implications about output and employment.

In the other class of models, sellers quote prices to their customers, who choose quantity unilaterally. The failure of prices to track marginal cost has important implications for efficiency and allocations in this class of models, as Robert Barro [1977] pointed out. Output is inefficiently responsive to demand shifts. Aggregate shocks can result in a general shortfall of output. This class of models underlies almost all modern work on the effects of monetary policy on employment and output.

References

Barro, Robert J. 1977. "Long-Term Contracting, Sticky Prices, and Monetary Policy," *Journal of Monetary Economics* 3(3): 305-16. July.

Gouriéroux, Christian, Alan Montfort, and E. Renault. 1993. "Indirect Inference," *Journal of Applied Econometrics* 8: S85-118. Supplement. December.

Hall, Robert E. 1980. "Employment Fluctuations and Wage Rigidity," *Brookings Papers on Economic Activity* No. 1: 91-123

______________. 2001. "The Stock Market and Capital Accumulation," *American Economic Review* 91: 1185-1202. December.

______________. 2002. "Industry Dynamics with Adjustment Costs," available at http://www.Stanford.edu/~rehall/page2.html.

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: 125-211.

Rotemberg, Julio J., and Michael Woodford. 1999. "The Cyclical Behavior of Prices and Costs," Chapter 16 in John Taylor and Michael Woodford (eds.) *Handbook of Macroeconomics*, vol. 1B: 1051-1135.

Smith, Anthony A., Jr. 1993. "Estimating Nonlinear Time Series Models Using Simulated Vector Autoregressions," *Journal of Applied Econometrics* 8:63-84.

Appendix: Data from the Producer Price Index

		Commodity data			Industry data	
	Industry description	*PPI Commodity Code*	*Start*	*End*	*PPI Industry Code*	*Start*
1	Agriculture	WPU01	47	01		
2	Metal mining	WPU1011	47	84	PCU10	86
3	Coal mining	WPU051	47	84	PCU12	86
4	Oil and gas extraction	WPU056	47	84	PCU13	86
5	Non-metallic mining	WPU1399	47	83	PCU14	85
6	Construction				PCUBNEW	87
7	Food and kindred products	WPU02	47	83	PCU20	85
8	Tobacco	WPU152	47	83	PCU21	85
9	Textile mill products	WPU03	47	83	PCU22	85
10	Apparel	WPU035	47	77	PCU23	85
11	Lumber and wood	WPU08	47	83	PCU24	85
12	Furniture and fixtures	WPU121	47	83	PCU25	85
13	Paper and allied	WPU09	47	83	PCU26	85
14	Printing, publishing and allied	WPU093	81	83	PCU27	85
15	Chemicals	WPU06	47	83	PCU28	85
16	Petroleum and coal products	WPU057	47	84	OCU29	86
17	Rubber and misc plastics	WPU07	47	83	PCU30	85
18	Leather	WPU04	47	83	PCU31	85
19	Stone, clay, glass	WPU131	47	83	PCU32	85
20	Primary metal	WPU1013	47	81	PCU33	85
21	Fabricated metal	WPU103	47	83	PCU34	85
22	Machinery, non-electrical	WPU11	47	83	PCU35	85
23	Electrical machinery	WPU117	47	84	PCU36	86
24	Motor vehicles	WPU1411	47	83	PCU371	85
25	Transportation equipment & ordnance	WPU14	69	84	PCU37	86
26	Instruments	WPU1172	67	83	PCU38	85
27	Misc. manufacturing	WPU15	47	84	PCU39	86
28	Transportation					
29	Communications					
30	Electric utilities					
31	Gas utilities					
32	Trade					
33	Finance Insurance and Real Estate					
34	Services					
35	Government enterprises				PCU43	80