

IS FIRM PRICING STATE OR TIME DEPENDENT? EVIDENCE FROM U.S. MANUFACTURING

Virgiliu Midrigan*

Abstract—If pricing is state dependent, firms are more likely to adjust whenever aggregate and idiosyncratic shocks reinforce each other and trigger desired price changes in the same direction. Using measures of technology shocks derived from production function estimates for four-digit U.S. manufacturing industries, I find that sectoral inflation rates are more sensitive to negative, as opposed to positive, technology disturbances in periods of higher economy-wide inflation, commodity price increases, and expansionary monetary policy shocks. I argue, using a state-dependent sticky price model that matches salient features of the U.S. microprice data, that these results suggest that pricing is state-dependent in U.S. manufacturing.

I. Introduction

MODELS with nominal rigidities play an important role in debates on the sources of business cycle fluctuations, as well as in optimal monetary policy discussions. An important challenge has been to build models with solid microfoundations that are consistent with microeconomic evidence on the price adjustment practices of individual producers. The purpose of this paper is to shed additional light on the price-setting practices of firms in the U.S. manufacturing sector. In particular, the question I ask is one that has been at the center of the debate about the role of nominal rigidities in the monetary transmission mechanism: Is firm price state or time dependent?

Firms price in a state-dependent fashion if the timing of price changes is endogenous and responds to disturbances to the firm's desired price. State-dependent pricing rules are optimal if the sole frictions that prevent price adjustment are physical menu costs of changing prices and communicating the information about price changes to the consumer. In this environment, an optimal policy is for firms to follow (S, s) rules: leave their prices unchanged if these are not sufficiently out of line and reprice only in response to a large disturbance (or cumulative history of disturbances).¹

Firms price in a time-dependent fashion if the date of price changes is a function of calendar time, not of the realization of cost of demand disturbances that affect the firm's desired price. In writing about the first generation of time-dependent models, Taylor (1980) and Calvo (1983), assume that the date at which prices change is exogenous. In these models, firms reprice at exogenously imposed calendar dates even when the potential gains from adjusting in other periods are large. The underlying frictions that justify

time-dependent rules are, in addition to menu costs, institutional restrictions and information-gathering or decision-making costs that render a predetermined schedule of price changes optimal. For example, Zbaracki et al. (2004) report that the pricing season of a large U.S. manufacturing firm occurs once each year and that new list prices are typically distributed every November. Ball, Mankiw, and Romer (1988), as well as Bonomo and Carvalho (2004), explicitly model the frictions that render time-dependent pricing rules optimal. In this second class of time-dependent pricing models, the interval between two consecutive price changes is endogenous and can indeed vary over time in response to changes in the trend growth rate of inflation or the volatility of the environment. Nevertheless, the exact dates at which prices are to be changed are predetermined and independent of contemporaneous disturbances to the firm's price.

The aggregate consequences of the two types of price rigidities can be very different. Firms' ability to respond to idiosyncratic and aggregate disturbances in state-dependent models typically renders monetary policy less potent in these environments. Caplin and Spulber (1987) and Caballero and Engel (1993) show that under special assumptions about the distribution of firm prices, money can indeed be neutral despite nominal rigidities at the firm level. Recent research, grounded in explicit household and firm maximization, and using stochastic forcing processes calibrated from the U.S. data, has overturned this neutrality result, but nevertheless reaches the conclusion that state-dependent pricing models generate smaller real effects from monetary shocks.²

The key mechanism that leads to smaller real effects from money shocks in economies with state-dependent pricing is an endogenous shift in the identity of adjusting firms. The endogenous timing of price changes in these models implies that the distribution of idiosyncratic disturbances to adjusting firms' desired prices varies with the aggregate state of the economy. Firms adjust when they need larger price changes, that is, when idiosyncratic and aggregate shocks reinforce each other and trigger desired price changes in the same direction. As a result, the mix of adjusters varies with the aggregate shock: in times of monetary expansions, adjusters are mostly firms that have received positive idiosyncratic shocks to their desired price. This change in the mix of adjusting firms imparts a greater degree of flexibility to the aggregate price level as firms that do adjust in a particular period are exactly those that need the largest price changes. Golosov and Lucas (2006) call this mechanism the

Received for publication August 30, 2005. Revision accepted for publication July 28, 2008.

* New York University.

I am indebted to George Alessandria, Bill Dupor, Paul Evans, and Mario Miranda for valuable advice and support, as well as Mark Bils, Michael Dotsey, Kevin Huang, Joe Kaboski, Aubhik Khan, and Kei-Mu Yi, an anonymous referee, and Julio Rotemberg, the editor, for useful suggestions.

¹ Barro (1972), Sheshinski and Weiss (1977, 1983), Caplin and Spulber (1987), Caplin and Leahy (1991).

² Dotsey, King, and Wolman (1999); Golosov and Lucas (2006).

selection effect, and Caballero and Engel (2007) refer to it as an *extensive margin effect*.

My goal in this paper is to measure the strength of this effect in the U.S. data. I use sectoral price and input and output data available from the NBER Manufacturing Productivity Database in order to measure idiosyncratic (sectoral) cost shocks. These are measures of technology residuals from production function estimates that allow for increasing returns, imperfect competition, and variable capacity and labor utilization, using the approach of Basu and Kimball (1997). I ask whether economy-wide disturbances alter the responsiveness of sectoral inflation rates to sectoral shocks. The state-dependent model predicts that they do: most adjusting firms are those that have been subject to cost shocks that trigger desired price changes in the same direction as the aggregate disturbance. As a result, the elasticity of sectoral inflation rates to idiosyncratic shocks—primarily determined by the fraction of adjusters in a given sector—should increase for those sectors for which the sectoral cost disturbance has the same sign as the aggregate disturbance.

I find strong evidence that aggregate disturbances affect the responsiveness of sectoral inflation rates to sectoral cost shocks in a manner predicted by the state-dependent model. For example, sectoral inflation rates are much more responsive to negative, as opposed to positive, technology shocks in periods with greater-than-average aggregate inflation, larger changes in commodity prices, and monetary policy shocks. This selection effect is both statistically significant and large: it raises the overall response of the inflation rate in the manufacturing sector to a monetary policy shock by up to 50%.

Several earlier papers provide insights into the price-setting practices of individual producers. Blinder et al. (1998) use survey evidence collected from a survey of CEOs and find that time-dependent rules of price adjustment are twice as common as state-dependent rules. Zbaracki et al. (2004) survey a large manufacturing firm and also find evidence of time-dependent pricing. Cecchetti (1986) and Kashyap (1995) find that the frequency of price changes increases during periods of higher inflation—behavior that is consistent with the implications of state-dependent models, but also of models with endogenous time-dependent pricing. Ball and Mankiw (1994, 1995) illustrate that in menu cost economies with positive trend inflation, an increase in the volatility of idiosyncratic shocks is inflationary, as most adjusters desire price increases. Similarly, changes in the skewness of the distribution of idiosyncratic shocks can also cause movements in the aggregate price level if pricing is state dependent. These authors find support for the state-dependent hypothesis as changes in higher-order moments of the distribution of sectoral relative price changes account for an important fraction of changes in aggregate U.S. inflation. Finally, this paper is complementary to that of Midrigan (2006), who studies the strength of the selection effect implied by mic-

rolevel observations of firm prices in grocery stores. He finds that this effect is muted and money is no longer neutral in a model calibrated to match the distribution of nonzero price changes in the data.

A final comment on terminology is in order. My claim that pricing is state dependent is not also evidence against time dependent in price setting. Distinguishing between pure time- and state-dependent models is not the purpose of this paper. Rather, the goal here is to measure the strength of the selection effect in the U.S. data. I think of the exercise presented here as providing an additional set of moments useful to enrich our understanding of how firms change prices, not an attempt to reject one model at the expense of another.

The rest of this paper is organized as follows. Section II presents a partial equilibrium model with nominal rigidities used to motivate the empirical exercise of this paper. In section III, I discuss the data and my measures of technology shocks. Section IV conducts the empirical analysis. The final section concludes.

II. State- versus Time-Dependent Pricing

In this section I illustrate the selection effect that arises in economies with state-dependent pricing and suggest a set of moments that can be used to evaluate the strength of this effect. To do so, I present two widely used versions of economies with sticky prices in which price stickiness arises endogenously, due to menu costs (I refer to this model as one with state-dependent pricing), and price stickiness is exogenously imposed, in a Calvo (1983) fashion, and in which the timing and frequency of price changes are exogenous (I refer to this model as one with time-dependent pricing). The two models share many features. I thus present the model economy in a unified fashion and then discuss the differences in the two pricing assumptions.

The model is similar to the partial equilibrium problem studied by Sheshinski and Weiss (1977). A firm's profits depend on its relative price, the ratio of the firm's nominal price to that of the aggregate price level, here assumed exogenous. I assume that the aggregate price level \bar{p}_t , evolves according to

$$\bar{p}_t = \bar{p}_{t-1} e^{g_t},$$

where g_t is the growth rate of the price level and evolves according to

$$g_t = \alpha + \delta g_{t-1} + \eta_t,$$

where $\eta_t \sim N(0, \sigma_\eta^2)$. In the discussion that follows, I interpret η_t as monetary policy shocks. Let $z_t = \frac{p_t}{\bar{p}_t}$ denote the firm's relative price, where p_t is the firm's nominal price. I assume constant elasticity demand functions: $q_t = z_t^{-\theta}$.

The firm's real profits in period t are $\Pi(z_t) = z_t^{-\theta}(z_t - \frac{c}{a_t})$, where $\frac{c}{a_t}$ is the (real) marginal cost of production, and a_t is the firm's productivity.³ The firm's productivity is the product of an idiosyncratic and sectoral component, $a_t = \psi_t \phi_t$, which evolves according to $\log(\psi_t) = \log(\psi_{t-1}) + \varepsilon_t$ and $\log(\phi_t) = \log(\phi_{t-1}) + u_t$, where ε_t is a sectoral and u_t a firm-specific productivity shock. The two shocks are drawn from a gaussian distribution with mean 0 and variance σ_ε^2 and σ_u^2 , respectively. Here, a sector is defined as a group of firms that share the same sectoral technology. I make the distinction between firms and sectors as the empirical work is conducted using sectoral data. I thus aggregate firm-level decision rules into sectors, as described below, by computing an average price for firms that share the same sectoral technology ψ_t .

A. State-Dependent Pricing

In this setup, firms face costs of adjusting nominal prices. Specifically, a firm incurs cost ξ every period in which $p_t \neq p_{t-1}$. The firm's problem is to

$$\max_{z_t} E_0 \sum_{t=0}^{\infty} \beta^t \left[\Pi(z_t, a_t) - \xi \mathcal{I}\left(z_t \neq \frac{z_{t-1}}{e^{g_t}}\right) \right], \quad (1)$$

where $\mathcal{I}(\cdot)$ is an indicator function that takes a value of 1 if the firm adjusts its nominal price.

To write this problem recursively, I bound the state-space so as to ensure that the period reward function is bounded. To this end, I define $\hat{z}_t = z_t a_t$. Although a_t is unbounded, \hat{z}_t is not, as optimality requires z_t be proportional to $\frac{1}{a_t}$. Given this normalization, I can write the firm's profits as $\Pi(\hat{z}_t, a_t) = a_t^{\theta-1} [\hat{z}_t^{1-\theta} - c \hat{z}_t^{-\theta}]$. I can thus rewrite the firm's problem as

$$V(\hat{z}, g) = \max [V^a(\hat{z}, g), V^n(\hat{z}, g)],$$

where V^a and V^n denote the firm's value of adjusting and not adjusting its nominal price, respectively, that satisfy

$$\begin{aligned} V^a(\hat{z}_{-1}, g) &= \max_{\hat{z}} \left[\hat{z}^{1-\theta} - c \hat{z}^{-\theta} - \xi \right. \\ &\quad \left. + \beta \int_{\varepsilon \times u \times \eta} e^{(\theta-1)(\varepsilon+u)} V(\hat{z}'_{-1}, g') dF(\varepsilon, u, \eta) \right], \\ V^n(\hat{z}_{-1}, g) &= \left[\hat{z}_{-1}^{1-\theta} - c \hat{z}_{-1}^{-\theta} \right. \\ &\quad \left. + \beta \int_{\varepsilon \times u \times \eta} e^{(\theta-1)(\varepsilon+u)} V(\hat{z}'_{-1}, g') dF(\varepsilon, u, \eta) \right], \end{aligned} \quad (2)$$

³ I suppress sector and firm subscripts to conserve notation and revert to them below when needed.

where $F(\cdot)$ is the joint cdf of the three shocks. The law of motion for the firm's (normalized) relative price is $\hat{z}'_{-1} = \frac{\hat{z}}{\exp(g)^{\theta} \exp(\eta)} \exp(\varepsilon + u)$ if the firm adjusts its price to \hat{z} and $\hat{z}'_{-1} = \frac{\hat{z}_{-1}}{\exp(g)^{\theta} \exp(\eta)} \exp(\varepsilon + u)$ if it leaves its price unchanged. The term $e^{(\theta-1)(\varepsilon+u)}$ that multiplies the firm's continuation value is the growth rate of the firm's technology $\left(\frac{a_{t+1}}{a_t}\right)^{\theta-1}$, which enters the profit function in the original problem.

B. Calvo Time-Dependent Pricing

In this exercise, I assume that firms have no control over the timing of their price changes. Rather, the probability that a firm adjusts in a given period is constant and equal to λ . The functional equations characterizing the firm's problem in this setup are

$$V(\hat{z}, g) = (1 - \lambda) V^a(\hat{z}, g) + \lambda V^n(\hat{z}, g),$$

where V^a and V^n satisfy

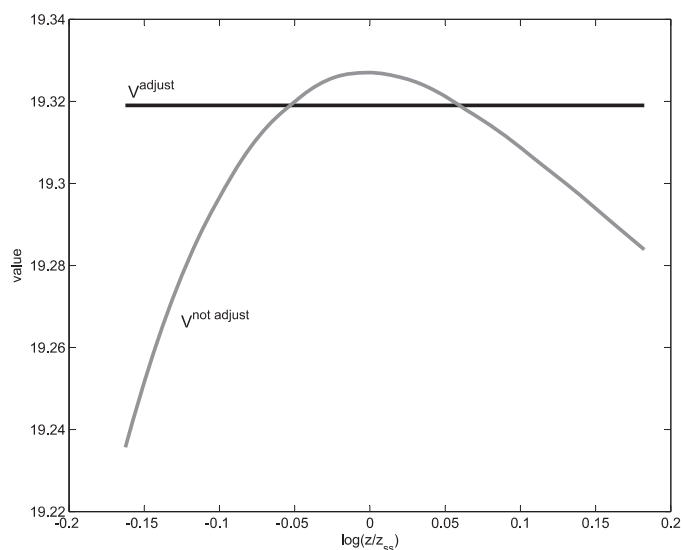
$$\begin{aligned} V^a(\hat{z}_{-1}, g) &= \max_{\hat{z}} \left[\hat{z}^{1-\theta} - c \hat{z}^{-\theta} \right. \\ &\quad \left. + \beta \int_{\varepsilon \times u \times \eta} e^{(\theta-1)(\varepsilon+u)} V(\hat{z}'_{-1}, g') dF(\varepsilon, u, \eta) \right], \\ V^n(s) &= \left[\hat{z}_{-1}^{1-\theta} - c \hat{z}_{-1}^{-\theta} \right. \\ &\quad \left. + \beta \int_{\varepsilon \times u \times \eta} e^{(\theta-1)(\varepsilon+u)} V(\hat{z}'_{-1}, g') dF(\varepsilon, u, \eta) \right]. \end{aligned} \quad (3)$$

To solve these problems, I employ collocation, a functional approximation technique. The idea behind this method is to approximate the two value functions with a linear combination of orthogonal polynomials and solve for the unknown coefficients by requiring that the two equations are satisfied exactly at a number of nodes along the state-space. (For a more detailed discussion of the solution method and its accuracy, see the technical appendix available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00016.)

C. Calibration

I choose parameter values to ensure that the predictions of the state-dependent model match certain features of the U.S. economy. The length of the period is set to one quarter, but the model will be evaluated against annual data from the NBER Productivity database. I therefore choose the parameters that characterize the process for the growth rate of the aggregate price level to match the annual mean, serial correlation, and volatility of inflation in the manufacturing

FIGURE 1.—FIRM VALUES



sector. This gives $\alpha = 0.98 \times 10^{-3}$, $\delta = 0.89$, $\sigma_\eta^2 = 2.32 \times 10^{-5}$. The elasticity of demand, θ , is equal to 5, so that the steady-state markup is equal to 25%.⁴ This leaves three additional parameters that must be calibrated in the state-dependent economy: ξ , the menu costs, as well as the volatility of sectoral σ_ϵ^2 and firm-specific σ_u^2 technology disturbances. The three targets that pin down these parameters are (1) an average duration of contract lengths of 15.3 months, which corresponds to an estimate of Leith and Malley (2007) regarding the frequency with which firms in the NBER productivity database change prices; (2) an average size of nonzero price changes of 9%, consistent with findings from Nakamura and Steinsson (2008) and Bils and Klenow (2004); and (3) a standard deviation of annual sectoral inflation rates of 4.9% in the sectoral price data available from the NBER Productivity Database. The parameters that render the model consistent with these targets are a menu cost, ξ , equal to 1.44% of the firm's steady-state revenues, a standard deviation of firm-specific shocks equal to $\sigma_u = 0.0254$ and of sector-specific shocks, $\sigma_\epsilon = 0.0122$.

As for the Calvo model, I choose λ to match a frequency of price changes of 0.196, as in the state-dependent model, and assign the same volatility of technology shocks.

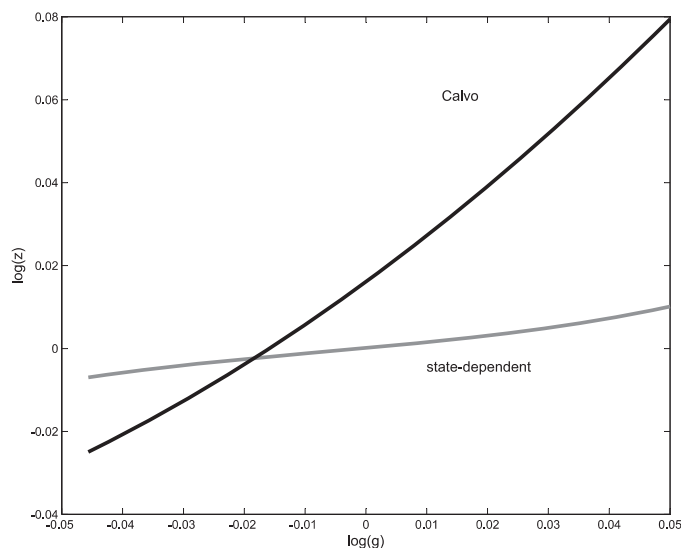
D. Optimal Pricing Rules

As is typical in economies with menu costs, firms follow generalized (S, s) rules and reprice whenever shocks, whether aggregate, sectoral, or idiosyncratic, force \hat{z}_{-1} to drift away from its optimum (which turns out to be close to $\frac{\theta}{\theta-1}c$, the frictionless optimum). Figure 1 illustrates the firm's value of changing and that of not changing its price as a function of $\log\left(\hat{z}_{-1}/\left(\frac{\theta}{\theta-1}c\right)\right)$: the deviation of the firm's price from its frictionless optimum. If a firm adjusts its

price, its value is independent of z_{-1} , by inspection of the firm's problem in equation (2). In contrast, if the firm does not adjust its price, it sells at z_{-1} , and the further away z_{-1} is from its optimum, the lower the firm's value is. The intersection of these two value functions determines the firm's inaction and adjustment regions; whenever $\log\left(\hat{z}_{-1}/\left(\frac{\theta}{\theta-1}c\right)\right)$ is sufficiently far from 0, the firm finds it worthwhile to pay the menu cost and adjust. In contrast, the Calvo firms' adjustment decisions is exogenous: firms adjust with a constant hazard $1 - \lambda$.

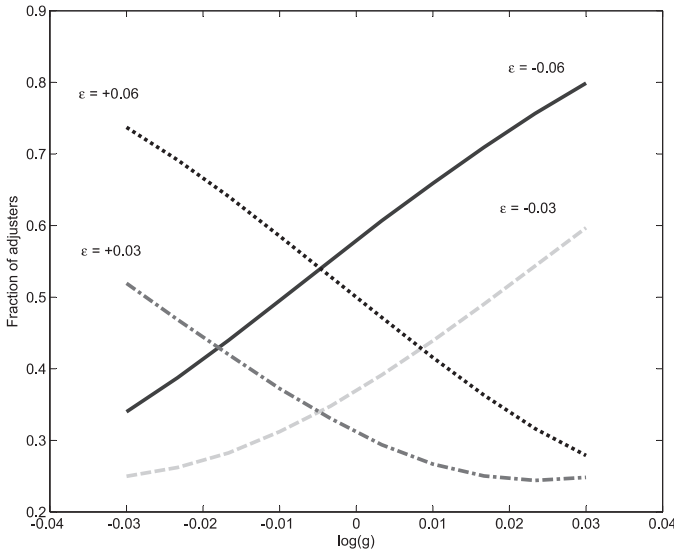
Consider next the firms' pricing rules. As the recursive representation of the problem indicates, the firm's price, conditional on adjustment, depends solely on the growth rate of the aggregate price level. Given that this growth rate, g , is persistent, it helps forecast future changes in the growth rate of the price level and therefore affects the adjusting firm's optimal price. As figure 2 indicates, Calvo firms respond more aggressively to an increase in the growth rate of the price level than state-dependent firms do. These differences in price functions arise because of the type of nominal frictions Calvo and menu cost firms are subject to. If a Calvo firm finds itself with a suboptimal price in a given period in the future, it pays dearly: given that it will not readjust its price for an average of five quarters, it will incur losses from the suboptimal price for a number of periods to come. In contrast, a state-dependent firm can always choose to pay the menu cost and reprice: its losses from having a suboptimal price in future periods are smaller than those of a time-dependent firm. This in turn implies that a Calvo firm has a stronger incentive to offset future expected changes in its marginal cost every time it adjusts than a state-dependent firm does. A similar argument explains why Calvo firms choose higher prices on average than state-dependent firms do: they have a stronger incentive to respond to the trend growth in the aggregate price level, as captured by α .

FIGURE 2.—OPTIMAL PRICE FUNCTIONS



⁴ See Obstfeld and Rogoff (2000) for a brief survey of estimates of θ .

FIGURE 3.—FRACTION OF ADJUSTERS IN A SECTOR



How do firms respond to technology disturbances? The fact that adjusting firms choose to return their normalized relative price to a target level \hat{z}^* , which depends on g , implies only that the firm's nominal price p responds one-for-one to a technology disturbance in both the Calvo and menu-cost world. Thus, an increase in a , because of an increase in the sectoral technology ψ or the idiosyncratic technology ϕ , leads firms to lower their nominal prices one-for-one.

I next aggregate firm decision rules into sectors in order to derive several moments of the sectoral price data that can be used empirically. Figure 3 plots the fraction of adjusters $\mathcal{F}r(\epsilon, g)$ in a sector that starts at the ergodic distribution of prices and is then subject to a sectoral technology shock ϵ , and a growth rate of the price level equal to g , in the state-dependent model. In addition, firms are subject to idiosyncratic technology shocks u . This fraction is, by assumption, constant in a time-dependent economy like Calvo.

Consider first sectors subject to negative technology disturbances ($\epsilon = -0.06$) and $\epsilon = -0.03$). Firms in these sectors desire, on average, to increase their prices in order to respond to the higher marginal costs of production. This incentive to change prices is reinforced if the economy-wide nominal shock is also positive. For this reason, the fraction of firms that adjusts in sectors with negative technology shocks increases in g , the growth rate of the general price level.

In contrast, firms in sectors with positive technology shocks see their marginal cost falling and desire price decreases. Their desire to decrease real prices is automatically satisfied if the aggregate price level increases, thereby eroding the firm's real price. The fraction of firms that adjusts in sectors with positive technology disturbances therefore decreases as the growth rate of the economy-wide price level rises.

The results presented in figure 3 are not useful for empirical purposes, as I do not directly observe the fraction of firms that adjust in a given sector. This fraction, however, strongly influences the elasticity of a sector's inflation rate to its technology disturbance. Given that adjusters respond one-for-one to a sectoral technology shock, the elasticity of sectoral inflation rates to sectoral shocks is a function of the fraction of firms that are adjusting prices. To see this, suppose that adjusting firm j in sector i chooses an inflation rate in period t that is equal to

$$\pi_{ijt}^* = \epsilon_{it} + g_t + \tilde{u}_{ijt},$$

which, in light of our discussion above, is a good approximation to a menu-cost firm's optimal price function. Here \tilde{u}_{ijt} captures the idiosyncratic shock to the firm's desired price and includes the contemporaneous cost shock u_{ijt} as well as the cumulative history of idiosyncratic and aggregate disturbances since the firm has previously adjusted. Letting $\Theta(\epsilon_{it}, g_t)$ denote the firm's adjustment region in the u_{ijt} space (the set of u_{ijt} for which a firm adjusts), the sectoral inflation rate, defined as $\pi_{it} = \int \pi_{ijt} dj$, is then equal to

$$\pi_{it} = \mathcal{F}r(\epsilon_{it}, g_t)(\epsilon_{it} + g_t) + \int_{\tilde{u}_{ijt} \in \Theta(\epsilon_{it}, g_t)} \tilde{u}_{ijt} dj.$$

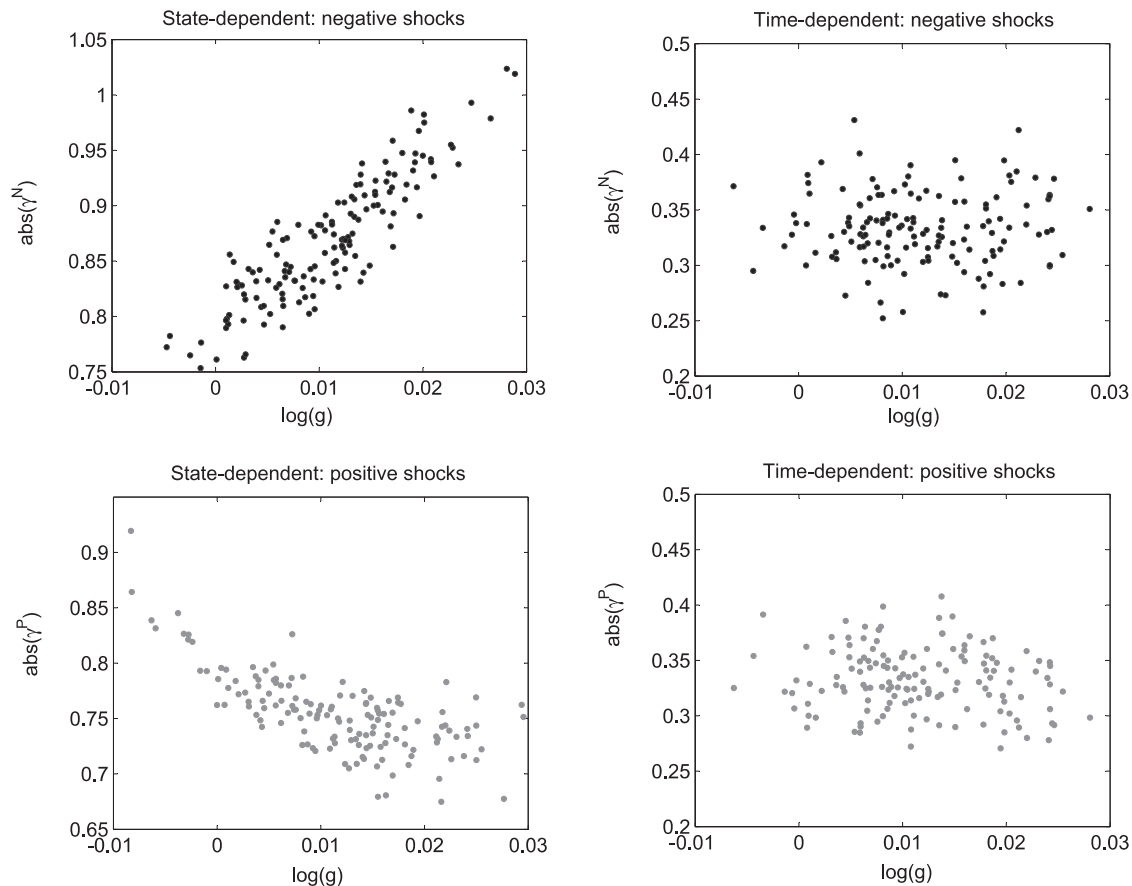
The second term in this expression captures the selection effect at the sectoral level: firms with \tilde{u}_{ijt} aligned with $\epsilon_{it} + g_t$ are more likely to adjust prices. Thus, although a regression of π_{it} on $\epsilon_{it} + g_t$ provides an upward-biased estimate of $\mathcal{F}r(\epsilon_{it}, g_t)$ because of the selection bias, the regression coefficient is nevertheless correlated with the fraction of adjusting firms.

Given this discussion, one way to test the state-dependent model is to estimate, for each period t , the following cross-sectional regressions of sectoral inflation rates on sectoral technology shocks in which the coefficients on negative, (γ_t^N), and positive, (γ_t^P), shocks, are allowed to differ:

$$\pi_{it} = \xi_t + \gamma_t^N \epsilon_{it} \mathcal{I}_{\epsilon_{it} < 0} + \gamma_t^P \epsilon_{it} \mathcal{I}_{\epsilon_{it} > 0} + u_{it}.$$

If pricing is indeed state dependent, γ_t^N should increase in absolute value (become more negative) in periods in which the economy is experiencing an aggregate shock that raises all firms' prices. These are as sectors in which more firms are adjusting since the sectoral shock reinforces the aggregate shock and calls for larger price increases. Similarly, γ_t^P should fall in absolute value (become less positive) in periods in which the economy is hit by a positive aggregate shock. To illustrate this, I aggregate individual firm decision rules and simulate the state- and time-dependent economies above and compute these elasticities for which period. Figure 4 plots γ_t^N and γ_t^P against the aggregate disturbance, g_t across periods, for the state- and time-dependent models.

FIGURE 4.—EFFECT OF AGGREGATE SHOCKS ON NEGATIVE AND POSITIVE ELASTICITIES



As anticipated, the absolute value of γ_t^N increases with the aggregate shock: when g_t is close to 0, the elasticity is close to 0.75, and as g_t increases to 0.03, this elasticity rises to 1 in absolute value. Similarly, the absolute value of γ_t^P falls from 0.9 to 0.7 as g_t increases from 0 to 0.03. Notice also that γ_t^P is somewhat flatter in g_t than γ_t^N is and that this slope flattens as g_t increases. This is an outcome of the trend growth rate in the aggregate price level, which implies that in simulations, g_t is mostly positive. As figure 3 indicates, the fraction of adjusting firms for sectors with positive shocks is flatter than for sectors with negative shocks in the region in which g is positive, as in this region the aggregate shock is canceled by the sectoral shock. This is a typical feature of state-dependent models: the hazard of price changes is flatter for smaller deviation of the desired price from its target than for larger deviations.⁵ Also note that if the aggregate shock increases even further, firms in all sectors of the economy, regardless of their sectoral disturbances, would find it optimal to increase their prices. Thus, given that the fractions in figure 3 or elasticities in figure 4 are drawn for small values of the aggregate disturbance, as in the U.S. data, all statements above hold locally rather than globally.

In practice, the two-stage procedure outlined above of estimating elasticities for each period using cross-sectional regressions and relating these elasticities to measures of aggregate disturbances is inefficient. One can instead parameterize γ_t^N and γ_t^P directly as functions of the aggregate shock:

$$\gamma_t^N = \beta_0 + \beta_1 g_t,$$

$$\gamma_t^P = \alpha_0 + \alpha_1 g_t,$$

and infer the size of the coefficients β_1 and α_1 from panel regressions that pool observations across sectors and time periods together,

$$\pi_{it} = \xi_i + \gamma_t^P \mathcal{J}_{(\epsilon_{it} > 0)} \epsilon_{it} + \gamma_t^N \mathcal{J}_{(\epsilon_{it} < 0)} \epsilon_{it} + \rho g_t + u_{it},$$

where ξ_i are sector-specific fixed effects that account for differences in the trend growth rate of prices across sectors.

As earlier, if pricing is state dependent, higher aggregate inflation, g_t , should increase the fraction of adjusting firms in sectors in which the technology shocks are negative and decrease the fraction of adjusters in sectors in which technology shocks are positive. As a result, higher aggregate inflation should increase the (absolute value of, that is, make it more negative) elasticity of sectoral inflation rates to

⁵ See, for example, Caballero and Engel (2007).

TABLE 1.—NONLINEARITIES IN THE MODEL

Variable	Coefficient	State-Dependent Pricing	Time-Dependent Pricing
g_t	ρ	0.92 (0.14)	0.74 (0.11)
$\varepsilon_{it}I(\varepsilon_{it} > 0)$	α_0	-1.03 (0.06)	-0.93 (0.03)
$\varepsilon_{it}I(\varepsilon_{it} < 0)$	β_0	-1.18 (0.14)	-1.02 (0.06)
$g_t\varepsilon_{it}I(\varepsilon_{it} > 0)$	α_1	-0.80 (2.03)	1.10 (1.08)
$g_t\varepsilon_{it}I(\varepsilon_{it} < 0)$	β_1	-5.39 (1.94)	-0.99 (1.78)

Note: Means and standard deviations over 500 estimates of equation (4) using model-generated data reported.

sectoral technology shocks, in sectors in which technology shocks are negative and decrease (in absolute value, that is, make it less negative) in sectors with positive technology shocks. The state-dependent model thus suggests that $\beta_1 < 0$ and $\alpha_1 > 0$. Substituting out the definitions of γ_t^N and γ_t^P , the regression I propose to estimate is

$$\begin{aligned} \pi_{it} = & \xi_i + \alpha_0 \mathcal{J}_{(\varepsilon_{it} > 0)} \varepsilon_{it} + \alpha_1 (g_t \times \mathcal{J}_{(\varepsilon_{it} > 0)}) \varepsilon_{it} \\ & + \beta_0 \mathcal{J}_{(\varepsilon_{it} < 0)} \varepsilon_{it} + \beta_1 (g_t \times \mathcal{J}_{(\varepsilon_{it} < 0)}) \varepsilon_{it} \\ & + \rho g_t + s_{it}. \end{aligned} \quad (4)$$

In table 1, I report the coefficient estimates in this regression computed using model-simulated data. I construct sectoral inflation rates in a manner that attempts to mimic the nature of the empirical data I study in the next section. In particular, I simulate firm-decision rules for 446 sectors, as in the NBER Productivity Database, for 36×4 quarters, as 36 years of data are available for empirical analysis. Each sector is made up of 125 firms: this number is the weighted (according to each sector's sales share) average of the number of firms in each four-digit SIC manufacturing sector as measured by the inverse of the Herfindhal-Hirschmann concentration ratio for 1992 reported by the U.S. Census Bureau. The period in the model is one quarter, while the NBER Productivity Database reports annual sectoral observations. To allow comparison between the model and the data, I construct annual sectoral inflation rates in the model, π_{it} , by computing a Paasche index using price and quantity data for individual firms. To compute sectoral productivity shocks, I divide total output produced in a given year by all firms in an industry by their total labor input. I use changes in log of this measure of labor productivity as a measure of industry-specific shocks, ε_{it} . Price stickiness at the firm level makes this an imperfect measure of the ε_{it} given that output is demand determined, but by mimicking the empirical exercise of the next section, I can quantitatively compare coefficient estimates in the model and in the data. Finally, g_t at the annual level is constructed by summing consecutive four-quarter nonoverlapping sets of the quarterly growth rate of the model's exogenous driving process, g_t , and filtering out low-

frequency variations using an HP(10) filter in the spirit of the business cycle literature. The consequences of filtering (and estimates without filtering) are discussed in the data section below, but shortly, filtering increases the strength of the selection effect, in both the model and the data, presumably because it allows us to isolate unexpected shocks to inflation that are not yet incorporated in the firms' prices. I employ 500 rounds of model simulations and report means and standard deviations of coefficient estimates across the different simulations in parentheses. I repeat this exercise for both the state- and time-dependent models.

Time aggregation to annual data clearly washes out some of the nonlinearities reported in figures 3 and 4 at the quarterly frequencies. To the extent to which most firms are able to respond to aggregate and sectoral shocks from one year to another, the importance of price stickiness, whether time or state dependent, is reduced. As a result, as the first column ("State Dependent") of table 1 illustrates, the coefficient on positive technology shocks, α_1 , is insignificantly different from 0 (although negative, the mean across 500 simulations is twice less than the standard deviation) at the annual frequency. In contrast, the average coefficient on negative shocks is large in absolute value and more than twice larger its standard deviation. The intuition for why the coefficient on positive shocks is virtually 0 is the same as in my discussion of figures 3 and 4 above and in Ball and Mankiw (1994): firms in sectors with positive shocks to their technology are less willing to change prices given that their incentive to lower the price is offset by trend aggregate inflation, and they are thus in a flatter region of their adjustment hazard, making their elasticities to technology shocks less responsive to aggregate inflation.

In contrast to the state-dependent model, all coefficients on the interaction terms are insignificantly different from 0 in my simulations of the time-dependent model, consistent with the evidence in figure 4. Together the two columns of table 1 suggest that one can measure the strength of the selection effect in the data using estimates of equation (4). Finally, notice that simulations of the state-dependent model consistently produce elasticities of inflation rates to aggregate (g_t) and sectoral (ε_{it}) shocks that are greater than those in the time-dependent models. This is because of the selection effect at the individual firm level and at the sectoral level (that is not completely soaked up by the nonlinear terms).

III. Data

I test the predictions of the state-dependent pricing model using annual data from 1958 to 1996 for 446 four-digit SIC industries from the NBER Manufacturing Productivity Database.⁶ The data are derived from various government

⁶ Available at <http://www.nber.org/nberces/nbprod96.htm> and discussed in Bartelsman and Gray (1996).

sources, notably the Census Bureau's Annual Survey of Manufacturing, and contains information on total shipments, materials expenditure, investment, capital stock, number of production and nonproduction workers, payroll, production worker hours and wages, and price deflators for shipments, materials, and so on for each industry. Material expenditures include expenditure on energy, and the deflator for materials accounts for movements in the price of energy. Bils and Chang (1999) is a recent example that uses this data set in order to ask how industry prices respond to variations in costs and production, although, given my focus on asymmetries in response to purely technological shocks, my approach differs from theirs along several dimensions. I use these data in order to conduct my empirical exercises as discussed below.

A. Measures of Technology Shocks

My measures of technology shocks are Solow (1957) residuals estimated using the methodology developed by Hall (1990) and Basu and Kimball (1997) in order to account for the possibility of increasing returns, imperfect competition, and variable input utilization, respectively.

I assume a differentiable production function in which firms produce output Y , using capital services \tilde{K} , labor services \tilde{L} , intermediate inputs of materials, and energy M according to

$$Y = F(\tilde{K}, \tilde{L}, M, A).$$

Capital services depend on the stock of capital K , but also capital utilization Z : $\tilde{K} = ZK$, while labor services depend on the number of workers N , hours worked per employee H , and each worker's effort level, E : $\tilde{L} = ENH$. Taking logarithms of this production function, totally differentiating, and invoking cost minimization, one obtains

$$dy = \mu[s_k dk + s_L(dn + dh) + s_m dm] \\ + \mu[s_k dz + s_L de] + da,$$

where lowercase letters denote logs, s_j is the share of factor j in total revenue, and μ is the markup. The difficulty in estimating this equation directly is that effort and capital utilization are not observed. I follow Basu and Kimball (1997) and proxy the unobserved input utilization with hours per worker dh .⁷ The justification for this approach is that firms operate along all margins simultaneously and, given convex costs of changing hours worked, effort, and capital utilization, will choose to change them simultaneously in response to a shock. Changes in hours worked are therefore correlated with unobserved capital utilization and effort. More formally, Basu and Fernald (2000) solve a dynamic cost minimization problem of a firm subject to

costs of changing employment levels, hours worked, and capital utilization; they show that as long as capital's depreciation rate does not depend on its utilization level and the production function is Cobb-Douglas, a log-linear approximation to the firm's optimality conditions implies that dz and de depend on dh only. I therefore estimate

$$\Delta y_{it} = c_i + \mu \Delta x_{it} + \gamma \Delta h_{it} + \xi_{it}, \quad (5)$$

where Δy_{it} is the change in the log output of industry i , Δx_{it} is the share-weighted sum of the growth rate of real inputs (labor, capital, materials, and energy). I calculate total output as shipments plus change in end-of-period inventories and deflate it using the price deflator for shipments. The Productivity Database distinguishes between production and nonproduction workers in reporting industry employment and reports hours data only for production workers. I use the two as separate inputs in the production function and assume that hours per worker are time invariant for nonproduction workers. My results are robust to an alternative measure of inputs that includes only production workers. My proxy for variable input utilization, Δh_i , is the log difference in hours per worker reported for production workers.

I calculate the share of each factor of production as the time-series average of total payments to each factor divided by total revenues in each industry. One could in principle depart from this Cobb-Douglas assumption of constant shares and allow shares to vary over time, but as Basu and Fernald (2000) argue, this approach increases the likelihood of misspecification because observed factor prices are not allocative period by period in a world with implicit contracts or quasi-fixity.⁸ To calculate payments to capital, I first calculate the Hall and Jorgenson (1967) user cost of capital, R , according to

$$R = (r + \delta) \frac{1 - ITC - \tau d}{1 - \tau},$$

where r is the required rate of return on capital (I follow Hall, 1990, and assume it equal to the S&P 500 dividend yield), δ is the depreciation rate, ITC is the investment tax credit, d is the present value of depreciation allowances, and τ is the corporate income tax rate. Jorgenson and Yun (1991) provide data on ITC , d , and δ for 53 types of capital goods, while the tax data are provided by the Bureau of Economic Analysis at the two-digit level of disaggregation. I calculate the user cost of capital for each asset and a weighted average over the different types of assets for each SIC 2 industry in the data set, with the weights reflecting the relative importance of each type of asset in each industry. I judge the relative importance of the different types of assets in each industry by using Bureau of Economic Analysis data on the 1982 Distribution of New Structures and Equipment

⁷ Conley and Dupor (1999) use electricity consumption to proxy for capital utilization.

⁸ Results are robust to allowing instead time-varying shares.

to using industries. The required payment to capital is finally calculated as RP_kK , where P_kK is the current-dollar value of the industry's stock of capital. Given that the database reports only wage and salary costs of labor, I follow Bils and Chang (1999) and magnify both production and nonproduction labor costs to account for employer pension payments and compensation benefits. These data are again based on information available in the underlying NIPA tables at the two-digit level of disaggregation. In addition, I magnify total labor costs (for both production and nonproduction workers) by 9% to account for the database's exclusion of payments to auxiliary and support personnel. Bartelsman and Gray (1996) report that these costs account for 7.9% and 10.7% of total payroll in manufacturing in 1972 and 1986, respectively.

OLS estimates of equation (5) are likely to be biased because of the correlation between technology shocks and input choices. I therefore instrument the right-hand-side variables using current and one-period lags of deflated oil price changes, changes in government spending, changes in the U.S. effective nominal exchange rate, and monetary policy shocks estimated using a seven-variable VAR according to the Christiano, Eichenbaum, and Evans (1999) block-recursive identification procedure. My instruments are similar to those used by Basu and Kimball (1997), to which I add a measure of changes in nominal exchange rates of the United States against its trading partners. Given the exchange rate disconnect puzzle documented in open-economy macroeconomics, it is unlikely that sectoral technology shocks are correlated with this variable.⁹

The relatively short span of time-series observations renders industry-by-industry estimates of the coefficients in equation (5) rather imprecise. I therefore pool two-digit industries together and estimate equation (5) using a panel (fixed-effects) 2SLS estimator for each SIC 2 industry.

Although my interest is not in estimates of equation (5) per se, I briefly compare my estimates to those in earlier work. The time-series standard deviation (average across all sectors) of the purified Solow residuals is 0.063, whereas that of changes in the TFP (the difference between the growth rate of real output and hare-weighted sum of the growth rates of inputs) is 0.076. The two series are strongly correlated (0.63). In contrast, Basu and Fernald (2000) estimate that technology shocks in the entire manufacturing sector are almost twice less volatile: the standard deviation of the Solow residual is 0.035 and that of the purified series is 0.028 according to their estimates. My estimates of returns to scale are also similar to those of Basu, Fernald, and Kimball (2004), who use the Jorgenson data set of 29 industries (including 21 industries at roughly the SIC 2 level) from 1949 to 1996. Their estimation strategy differs slightly from mine as they restrict the coefficient on the proxy for input utilization to be constant across industries,

but despite the differences in the level of aggregation underlying the two sets of estimates, my results and theirs are not too dissimilar. For durable goods, the median return to scale estimate is 1.11, compared to 1.07 in their work, with a correlation of 0.71 across coefficient estimates in the different industries. For nondurables, the correlation is 0.77 if one excludes two industries (food and leather) for which these parameters are imprecisely estimated, while the degree of returns to scale is higher in my work (1.07) than theirs (0.89).¹⁰

IV. Empirics

Before I discuss formal estimates from panel regressions, I use my estimates of technology shocks $\varepsilon_{it} = \tilde{\varepsilon}_{it} + c_i$, constructed above, to estimate the following cross-sectional regressions,

$$\pi_{it} = \xi_t + \gamma_t^N \varepsilon_{it} \mathcal{I}_{\varepsilon_{it} < 0} + \gamma_t^P \varepsilon_{it} \mathcal{I}_{\varepsilon_{it} > 0} + u_{it}, \quad (6)$$

for each time period using ordinary least squares, where π_{it} are sectoral inflation rates. Recall from figure 4 that the state-dependent model predicts that γ_t^N should increase in absolute value and γ_t^P should fall in absolute value in periods of higher aggregate disturbances. In figure 5 I plot the (absolute value of) two elasticities against three measures of aggregate shocks that increase firms' desired prices: HP-filtered inflation in the manufacturing sector, commodity price changes, and a (two-year lagged) measure of monetary policy disturbances described above. The model is silent as to what measures of aggregate shocks should be included: any disturbance that affects all firms' desired nominal prices should raise the fraction of adjusters in sectors with negative shocks. Commodity price changes and monetary disturbances are natural proxies for g_t in the model. So is manufacturing inflation, as in the model firm's prices on average respond strongly to the nominal disturbance g_t , for arguments discussed in Caplin-Spulber (1987), Golosov-Lucas (2006), and Midrigan (2006).

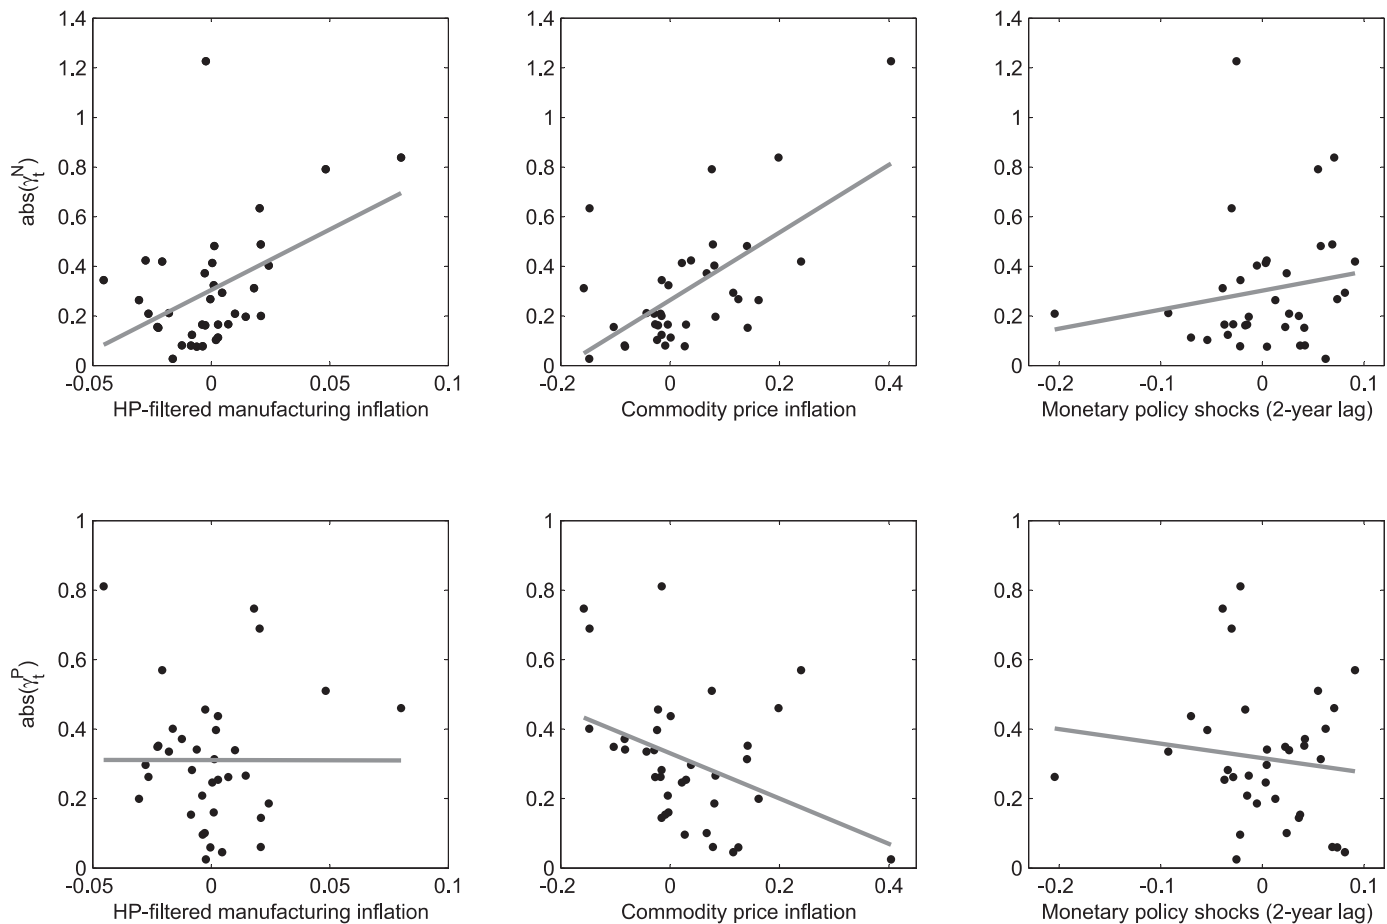
The line through these scatter plots is from a fitted OLS regression of $\gamma_t^N(\gamma_t^P)$ against the variable on the x -axis. For all three measures of shocks, the elasticity on negative sectoral shocks increases with the aggregate disturbance. Similarly, the elasticity on positive sectoral shocks decreases with commodity price inflation and monetary policy shocks, although it does not vary with HP-filtered inflation. This last fact should not be of much concern as simulations of the model reported in table 1 do not suggest a strong relationship between the elasticity on positive shocks and the aggregate disturbance at annual frequencies.

I next proceed to formally measuring the size of the selection effect in the U.S. data employing a panel specification in which elasticities to positive and negative shocks

⁹ Results are robust to excluding the nominal exchange rate as an instrument.

¹⁰ My estimates are close to those of Burnside, Eichenbaum, and Rebelo (1995).

FIGURE 5.—ELASTICITY OF SECTORAL INFLATION RATES TO SECTORAL TECHNOLOGY SHOCKS: U.S. MANUFACTURING (1961–1996)



are directly parameterized as functions of aggregate disturbances. In particular, I employ the same regression specifications as in equation (4), which I repeat here for convenience:

$$\pi_{it} = \xi_i + \alpha_0 \mathcal{I}_{(\varepsilon_{it} > 0)} \varepsilon_{it} + \alpha_1 (g_t \times \mathcal{I}_{(\varepsilon_{it} > 0)}) \varepsilon_{it} \\ + \beta_0 \mathcal{I}_{(\varepsilon_{it} < 0)} \varepsilon_{it} + \beta_1 (g_t \times \mathcal{I}_{(\varepsilon_{it} < 0)}) \varepsilon_{it} + \rho g_t + s_{it}.$$

As in figure 5, I use three alternative measures of aggregate disturbances, g_t : manufacturing inflation (HP filtered), commodity price inflation, and a measure of monetary policy shocks due to Christiano et al. (1999), which uses a recursive identification assumption and nonborrowed reserves as the postulated instrument. I use these alternative measures of shocks, instead of focusing solely on CPI inflation, in order to ascertain the robustness of my results but also to establish causality, as exogenous variations in γ_t^N and γ_t^P can themselves trigger variation in aggregate inflation. In particular, the two elasticities can fluctuate, and thereby affect inflation in the presence of exogenous changes in higher-order moments of the distribution of idiosyncratic cost shocks, as in Ball and Mankiw (1995).

I present the results in table 2. For comparison, the first column of the table reports the size of these coefficients

predicted by the state-dependent model, together with the standard deviation, across different simulations of the model, of these coefficients in parentheses. The next four columns, labeled I–III, correspond to the different measures of g_t in the regressions. I report, in parentheses, standard errors for these coefficients. These standard errors are corrected for the bias that arises because of my use of a two-stage procedure that imparts uncertainty to my estimates of the technology shocks, ε_{it} ,¹¹ as well as for heteroskedasticity and serial correlation across industries of arbitrary form, by employing an Arellano (1987)-type correction. I describe the approach used to correct for the two-stage bias in the supplemental appendix, available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00016.

Notice first in columns I to III of table 2 that sectoral inflation rates increase with all measures of aggregate shocks and decrease with sector-specific technology shocks. The size of the coefficient on the aggregate shock, g_t , varies substantially from one specification to another, which is not surprising, given that these alternative measures are characterized by different degrees of persistence and the fact

¹¹ In the appendix, I describe the exact correction employed.

TABLE 2.—NONLINEARITIES IN THE U.S. DATA

Variable	Coefficient	Model	Data		
			I HP-Filtered Manufacturing Inflation	II Commodity Price Inflation	III CEE Money Shock, Two-Year Lag
g_t	ρ	0.92 (0.14)	0.74 (0.05)	0.03 (0.01)	0.15 (0.02)
$\varepsilon_{it}I(\varepsilon_{it} > 0)$	α_0	-1.03 (0.06)	-0.23 (0.03)	-0.25 (0.03)	-0.22 (0.03)
$\varepsilon_{it}I(\varepsilon_{it} < 0)$	β_0	-1.18 (0.14)	-0.38 (0.04)	-0.36 (0.04)	-0.41 (0.04)
$g_t\varepsilon_{it}I(\varepsilon_{it} > 0)$	α_1	-0.80 (2.03)	-0.27 (0.74)	0.79 (0.22)	1.06 (0.38)
$g_t\varepsilon_{it}I(\varepsilon_{it} < 0)$	β_1	-5.39 (1.94)	-6.56 (1.15)	-1.66 (0.31)	-2.25 (0.50)
R^2			0.37	0.17	0.17
Number of observations		16,056	16,056	16,056	16,056

Note: Estimates of equation (4) for model and data. Standard errors, corrected for bias from two-stage estimates and robust to heteroskedasticity and serial correlation (Arellano, 1987, correction) reported in parentheses for data. Standard deviations across 500 simulations reported for the model.

that firms in a sticky price environment respond more aggressively to more persistent disturbances, which are expected to last longer. The size of the coefficient on technology shocks is similar across the last three columns, and always of a negative sign. The absolute value of these coefficients is, however, smaller in the data, which is perhaps evidence of strategic complementarities that prevent firms from fully responding to sectoral shocks. Alternatively, there may be more mean reversion in the process for sectoral technology shocks than imposed by the unit root assumption in the model. Finally, notice that the coefficient that captures most strongly the selection effect, β_1 , is of the sign and magnitude predicted by my simulations of the state-dependent model. Firms in sectors with negative technology shocks raise prices faster in periods of positive aggregate disturbances, suggesting state dependence in their pricing decisions. As seen in the “Model” column, the model suggests that the nonlinear response to positive shocks, as measured by α_1 , should wash out at the annual frequency. The data produce mixed implications regarding the size of this effect. Using manufacturing inflation as a measure of disturbances, I estimate a coefficient $\hat{\alpha}_1$ that is

indistinguishable from 0 statistically. Using commodity price changes and monetary policy shocks, I obtain positive coefficients, in line with the predictions of the model at quarterly frequencies (figure 4). What is more important for the aggregate consequences of these nonlinearities is that for all three measures of shocks, the elasticity on negative shocks rises (in absolute value) relative to that on positive shocks, thereby increasing the flexibility of the aggregate price level in excess of that in time-dependent models that assume a hazard of price adjustment independent of the size of the shock.

In table 3 I report a robustness check in which I use actual U.S. manufacturing inflation as a measure of aggregate disturbances. I report results for filtered and unfiltered measures of inflation in both the model and the data. First, notice that the model simulations show a stronger selection effect—as measured by the absolute value of β_1 (5.39 versus 3.63) or the difference between β_1 and α_1 (4.59 versus 2.72)—when HP-filtered inflation is used as a measure of aggregate disturbances. An argument in Ahlin and Shintani (2007) suggests that this is the case because in environments with persistent inflation, transition from a

TABLE 3.—CONSEQUENCES OF HP FILTERING AND INSTRUMENTING

Variable	Coefficient	HP-Filtered Inflation		Unfiltered Inflation		
		Model	Data	Model	Data: Raw Inflation	Data: Instrumented Inflation
g_t	ρ	0.92 (0.14)	0.74 (0.05)	0.86 (0.06)	0.75 (0.02)	0.66 (0.04)
$\varepsilon_{it}I(\varepsilon_{it} > 0)$	α_0	-1.03 (0.06)	-0.23 (0.03)	-0.98 (0.06)	-0.19 (0.03)	-0.25 (0.04)
$\varepsilon_{it}I(\varepsilon_{it} < 0)$	β_0	-1.18 (0.14)	-0.38 (0.04)	-1.03 (0.04)	-0.15 (0.02)	-0.20 (0.04)
$g_t\varepsilon_{it}I(\varepsilon_{it} > 0)$	α_1	-0.80 (2.03)	-0.27 (0.74)	-0.91 (1.30)	-1.97 (0.43)	0.01 (0.61)
$g_t\varepsilon_{it}I(\varepsilon_{it} < 0)$	β_1	-5.39 (1.94)	-6.56 (1.15)	-3.63 (0.89)	-4.49 (0.48)	-4.77 (0.75)
R^2			0.25		0.42	0.26
Number of observations		16,056	16,056	16,056	16,056	16,056

Note: Estimates of equation (4) for model and data. Standard errors, corrected for bias from two-stage estimates and robust to heteroskedasticity and serial correlation (Arellano, 1987, correction) reported in parentheses for data. Standard deviations across 500 simulations reported for the model.

TABLE 4.—QUANTITATIVE IMPORTANCE OF STATE-DEPENDENT TERMS

Elasticity of Sectoral Inflation to Technology Shocks Evaluated at:		$\pi_t = \text{Mean}$	$\pi_t = \text{Mean} + 1 \text{ s.d.}$
(I) HP-filtered π_t	Positive shocks	−0.23 (0.03)	−0.23 (0.03)
	Negative shocks	−0.38 (0.04)	−0.53 (0.05)
		$\Delta Pcom_t = \text{Mean}$	$\Delta Pcom_t = \text{Mean} + 1 \text{ s.d.}$
(II) $\Delta Pcom_t$	Positive shocks	−0.23 (0.03)	−0.14 (0.03)
	Negative shocks	−0.41 (0.04)	−0.59 (0.05)
		$\Delta shock_{t-2} = \text{Mean}$	$\Delta shock_{t-2} = \text{Mean} + 1 \text{ s.d.}$
(III) CEE monetary shock (2-year lag)	Positive shocks	−0.22 (0.03)	−0.16 (0.03)
	Negative shocks	−0.42 (0.04)	−0.54 (0.04)

Note: Elasticity computed using estimates in table 2. Standard errors in parentheses.

period of high inflation (in which firms charge higher prices in expectation of high inflation in future periods) to low inflation (in which firms charge lower prices as the incentive to front-load future desired price increases is reduced) may lower firms' desired prices. HP-filtered measures of inflation may thus account for this effect by conditioning the elasticities on deviations of inflation from its trend and thus controlling for the effect of past inflation on firms' desired prices.

As in the model, use of actual inflation provides weaker, but nevertheless significant, evidence of a selection effect in the data. The absolute value of β_1 drops from 6.56 to 4.49, while the difference between the two coefficients drops from 6.29 to 2.52. The value of α_1 (−1.97) is negative and significantly different from 0 which is somewhat inconsistent with the model's intuition, but recall that the model has no tight implications regarding the sign of α_1 at the annual frequency. The large standard deviation of estimates of α_1 in simulations puts the coefficient of −1.97 in the data well in the range of the estimates in the model.

A final robustness check I perform in the last column of table 3 is to instrument inflation to correct for potential endogeneity of manufacturing inflation, which is itself endogenous to the elasticities on sectoral shocks. I instrument inflation with oil price changes, monetary policy shocks, and changes in the nominal exchange rates. Results are very similar to those reported in the adjacent column and, if anything, suggest an even stronger selection effect. In particular, the coefficient on positive shocks is now virtually 0.

A. Interpreting the Results

I have established above the statistical significance of state-dependent pricing terms in explaining fluctuations in sectoral inflation rates. I next ask whether their effect is

quantitatively large. I use my estimates of equation (4) and calculate, in table 4, the effect of a 1 standard deviation increase in the different measures of aggregate shocks on the elasticity of sectoral inflation to negative or positive technology shocks. I first calculate what the elasticities γ_t^N and γ_t^P would be in the absence of economy-wide disturbances, when the aggregate variables are at their time-series means: these are the estimates of $\alpha_0 + \alpha_1 \times \text{mean}(g)$ and $\beta_0 + \beta_1 \times \text{mean}(g)$ in equation (4). The three different sets of estimates in table 4 correspond to different specifications of the aggregate disturbances: HP-filtered manufacturing inflation, commodity price inflation, and monetary policy shocks. Note first that on average, firms are more willing to increase prices in response to adverse technology disturbances than lower prices in response to favorable shocks: the elasticity on positive shocks is close to −0.2, while that on negative shocks is close to −0.4 when the aggregate variables are at their steady-state means. I thus corroborate, although in a different environment, the results of Peltzman (2000), who finds that output prices are more likely to respond to cost increases than decreases. I next compute the effect of a 1 standard deviation aggregate shock on the elasticities γ_t^N and γ_t^P (for example, $\alpha_0 + \alpha_1 \times [\text{mean}(g) + \text{std. dev.}(g)]$). Notice that for all measures of aggregate disturbances, with the exception of manufacturing inflation rates, an increase in the size of the nominal disturbances reduces the elasticity of sectoral inflation to positive technology shocks by around 40% (for example, from −0.23 to −0.14 for commodity price inflation), while increasing that on negative technology shocks by 30% (for example, from −0.41 to −0.59 for commodity price inflation). An increase in HP-filtered inflation rate itself increases the responsiveness to technology shocks in sectors with negative technology shocks, while leaving the elasticity in sectors with positive shocks unchanged.

How important are these changes in elasticities quantitatively? To answer this question, I resort to the following experiment. Note in equation (4) that the response of sectoral inflation rates to an aggregate shock is $\frac{\partial \pi_{it}}{\partial g_t} = \rho + \alpha_1 \varepsilon_{it}$ if $\varepsilon_{it} > 0$ and $\frac{\partial \pi_{it}}{\partial g_t} = \rho + \beta_1 \varepsilon_{it}$ if $\varepsilon_{it} < 0$. I compute these derivatives for all periods and sectors in my sample for the Christiano et al. (1999) measure of monetary policy shocks and average them across periods and sectors for a measure of how the aggregate manufacturing industry responds to an expansionary nominal disturbance. I also calculate what these “impulse responses” would have been in the absence of state-dependent terms by imposing $\alpha_1 = \beta_1 = 0$. I find that the elasticity of inflation to a monetary policy shock is equal to 0.058 (standard error, s.e., equal to 0.013) in the absence of state dependence in the first period following the shock, 0.148 (s.e. equal to 0.013) in the second period following the shock and 0.167 (s.e. equal to 0.012) in the third period after the shock. In contrast, the actual responses (in the presence of state-dependent terms) are equal to 0.091 (s.e. equal to 0.016), 0.222 (s.e. 0.016) and 0.173 (s.e. 0.016), respectively. Hence, the state-dependent terms increase the responsiveness of inflation rates to monetary shocks by almost 50%, suggesting that endogenous variation in the identity and fraction of adjusting firms is an important source of movements in the overall inflation rate of the U.S. manufacturing sector.

B. Robustness Checks

I have performed several checks to ensure the robustness of these results. To conserve space, these are not reported here, but are available in an earlier version of the paper available on my Web page (Midrigan, 2005). In particular, I have found strong evidence of nonlinearities in the response of sectoral inflation rates to aggregate and sectoral shocks using a continuous parameterization in which a sector’s responsiveness to technology shocks is allowed to vary smoothly with the size of the sectoral and aggregate disturbances. I have redone all analysis using an alternative measure of technology shocks, one based on long-run restrictions, as in Blanchard and Quah (1989) and Gali (1999). I have shown that the responsiveness of negative or positive elasticities to aggregate shocks does not vary much with an industry’s market concentration ratio, suggesting that changes in firms’ ability to collude over the business cycle, as in Rotemberg and Saloner (1986), are not what accounts for the large selection effect I document. I have allowed for heterogeneity, across sectors, in the responsiveness of sectoral inflation rates to technology and aggregate shocks and have found evidence of a large (although reduced relative to estimates that do not control for heterogeneity) selection effect. Finally, I have shown that my results are robust across subsamples (1961–1981 and 1982–1996), although, not surprisingly, nonlinearities are easier to identify during the pre-Volcker era of higher and more volatile inflation.

V. Conclusion

In state-dependent sticky price models, firms are more likely to pay the adjustment costs and change prices if idiosyncratic and aggregate shocks reinforce each other and trigger desired price changes in the same direction. The state-dependent model therefore predicts that the distribution of idiosyncratic shocks, conditional on adjustment, varies endogenously in response to aggregate disturbances, thereby giving rise to a selection or extensive margin effect that imparts a greater degree of flexibility to the aggregate price level in response to nominal disturbances. This paper suggests that this selection effect is quantitatively important in U.S. manufacturing. Using highly disaggregated four-digit data on sectoral input, output, and inflation rates in the U.S. manufacturing sector, I find that sectors that are hit by negative technology shocks adjust more readily in times of greater than expected inflation, increases in commodity prices and larger monetary policy shocks. These results suggest that firm pricing indeed has an important state-dependent element.

REFERENCES

- Ahlin, Christian, and Mototsugu Shintani, “Menu Costs and Markov Inflation: A Theoretical Revision with New Evidence,” *Journal of Monetary Economics* 54 (2007), 753–784.
- Arellano, Manuel, “Computing Robust Standard Errors for Within-Group Estimators,” *Oxford Bulletin of Economics and Statistics* 49 (1987), 431–434.
- Ball, Lawrence, and N. Gregory Mankiw, “Asymmetric Price Adjustment and Economic Fluctuations,” *Economic Journal* 104 (1994), 247–261.
- , “Relative Price Changes as Aggregate Supply Shocks,” *Quarterly Journal of Economics* 110 (1995), 161–193.
- Ball, Lawrence, N. Gregory Mankiw, and David Romer, “The New Keynesian Economics and the Output-Inflation Trade-Off,” *Brookings Papers on Economic Activity* 1 (1988), 1–65.
- Barro, Robert, “A Theory of Monopolistic Price Adjustment,” *Review of Economic Studies* 39 (1972), 17–26.
- Bartelsman, Eric, and Wayne Gray, “The NBER Manufacturing Productivity Database,” NBER technical working paper no. 205 (1996).
- Basu, Susanto, and John Fernald, “Why Is Productivity Procyclical? Why Do We Care?” FRB Chicago working paper (2000).
- Basu, Susanto, John Fernald, and Miles Kimball, “Are Technology Improvements Contractionary?” NBER working paper no. 10592 (2004).
- Basu, Susanto, and Miles Kimball, “Cyclical Productivity with Unobserved Input Variation,” NBER working paper no. 5915 (1997).
- Bils, Mark, and Yongsung Chang, “Understanding How Price Responds to Cost and Production,” NBER working paper no. 7311 (1999).
- Bils, Mark, and Peter Klenow, “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy* 112 (2004), 947–985.
- Blanchard, Oliver, and Danny Quah, “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *American Economic Review* 79 (1989), 657–673.
- Blinder, Alan, Elie Canetti, David Lebow, and Jeremy Rudd, *Asking about Prices: A New Approach to Understanding Price Stickiness* (New York: Russell Stage Foundation, 1998).
- Bonomo, Marco, and Carlos Viana de Carvalho, “Endogenous Time-Dependent Rules and Inflation Inertia,” *Journal of Money, Credit, and Banking* 36 (2004), 1015–1041.
- Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, “Capital Utilization and Returns to Scale,” in Ben Bernanke and Julio Rotemberg (Eds.), *NBER Macroeconomics Annual* (Cambridge, MA: MIT Press, 1995).
- Caballero, Ricardo, and Eduardo Engel, “Heterogeneity and Output Fluctuations in a Dynamic Menu Cost Economy,” *Review of Economic Studies* 60 (1993), 95–119.

- "Price Stickiness in Ss models: New Interpretations of Old Results," mimeograph (2007).
- Calvo, Guillermo, "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics* 12 (1983), 383–398.
- Caplin, Andrew, and John Leahy, "State-Dependent Pricing and the Dynamics of Money and Output," *Quarterly Journal of Economics* 106 (1991), 683–708.
- Caplin, Andrew, and Daniel Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics* 102 (1987), 703–725.
- Cecchetti, Stephen, "The Frequency of Price Adjustment: A Study of Newsstand Prices of Magazines," *Journal of Econometrics* 31 (1986), 255–274.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, "Monetary Policy Shocks: What Have We Learned and to What End?" NBER working paper no. 6400 (1999).
- Conley, Timothy, and Bill Dupor, "A Spatial Analysis of Sectoral Complementarity," *Journal of Political Economy* 111 (2003), 311–352.
- Dotsey, Michael, Robert King, and Alexander Wolman, "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics* 114 (1999), 655–690.
- Gali, Jordi, "Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89 (1999), 249–271.
- Golosov, Mikhail, and Robert E. Lucas, Jr., "Menu Costs and Phillips Curves," NBER working paper no. 10187 (2003).
- Hall, Robert, "Invariance Properties and Solow's Productivity Residuals" (pp. 71–112), in P. A. Diamond (Ed.), *Growth/Productivity Employment* (Cambridge, MA: MIT Press, 1990).
- Hall, Robert, and Dale Jorgenson, "Tax Policy and Investment Behavior," *American Economic Review* 57 (1967), 391–414.
- Jorgenson, Dale, and Kun-Young Yun, *Tax Reform and the Cost of Capital* (New York: Oxford University Press, 1991).
- Kashyap, Anil, "Sticky Prices: New Evidence from Retail Catalogs," *Quarterly Journal of Economics* 110 (1995), 245–274.
- Leith, Campbell, and Jim Malley, "A Sectoral Analysis of Price-Setting Behavior in U.S. Manufacturing Industries," this REVIEW 89 (2007), 335–342.
- Midrigan, Virgiliu, "Is Firm Pricing State or Time-Dependent? Evidence from US Manufacturing and General Equilibrium Implications," mimeograph (2005).
- "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations," mimeograph (2006).
- Nakamura, Emi, and Jon Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics* 123 (2008), 1415–1464.
- Obstfeld, Maurice, and Kenneth Rogoff, "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?" NBER working paper no. 7777 (2000).
- Peltzman, Sam, "Prices Rise Faster Than They Fall," *Journal of Political Economy* 108 (2000), 466–502.
- Rotemberg, Julio, and Garth Saloner, "A Supergame-Theoretic Model of Price Wars during Booms," *American Economic Review* 76 (1986), 390–407.
- Sheshinski, Eytan, and Yoram Weiss, "Inflation and Costs of Price Adjustment," *Review of Economic Studies* 44 (1977), 287–303.
- "Optimum Pricing under Stochastic Inflation," *Review of Economic Studies* 50 (1986), 513–529.
- Solow, Robert, "Technological Change and the Aggregate Production Function," this REVIEW 39 (1957), 312–320.
- Taylor, John, "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy* 88 (1980), 1–23.
- Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen, "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," this REVIEW 86 (2004), 514–533.