Understanding How Employment Responds to Productivity Shocks When Firms Hold Inventories*

Yongsung Chang
Seoul National University

Andreas Hornstein
Federal Reserve Bank of Richmond

Pierre-Daniel Sarte
Federal Reserve Bank of Richmond

August 17, 2006

Abstract

Whether technological progress raises or lowers aggregate employment in the short run has been the subject of much debate in recent years. Using a simple model of industry employment, we show that cross-industry differences of inventory holding costs, demand elasticities, and price rigidities potentially all affect employment decisions in the face of productivity shocks. In particular, the employment response to a permanent productivity shock is more likely to be positive the less costly it is to hold inventories, the more elastic industry demand is, and the more flexible prices are. Using data on 458 4-digit U.S. manufacturing industries over the period 1958-1996, we find statistically significant effects of variations in inventory holdings and demand elasticities on short-run employment responses, but find less evidence pertaining to the effects of measured price stickiness.

Keywords: Productivity, Employment, Inventory Investment, Sticky Prices
JEL Classification: E13, E22, E31

*This is a substantially revised version of the Federal Reserve Bank of Richmond Working Paper 2004-09, "Productivity, Employment, and Inventories." We wish to thank Bennett McCallum, John Fernald, Mark Watson, as well as seminar participants at the 2004 NBER Economic Fluctuations Group meeting and the NBER summer workshop for their comments. We further thank Valerie Ramey for calling our attention to the importance of industry demand elasticities. Her suggestions have led to substantial re-writing of an earlier draft. We thank Dan Hurley for excellent research assistance. Chang acknowledges financial support from the Korea Research Foundation grant KRF-2005-0035. Any opinions expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

Our e-mail addresses are: `yohg@snu.ac.kr`, `andreas.hornstein@rich.frb.org`, and `pierre.sarte@rich.frb.org`.

1. Introduction

In 2004, inventories held throughout the U.S. economy represented approximately 120 percent of final sales of goods and structures.[1] Industries that produce goods for inventory, or for which inventories are a significant component of the sales process, currently account for half of all private sector employment. Furthermore, during recessions, changes in inventories account for about 70 percent of the peak-to-trough decline in real GDP.[2] In spite of these facts, modern business cycle analysis has traditionally placed little emphasis on the role of inventories.[3]

The nature of business cycles changes considerably when firms carry inventories. Firms are able to make current production differ from current sales, speculate on future price movements, and absorb shocks to demand and costs. Using a simple industry model, we describe how inventories affect firms' employment response to productivity shocks. In a cross-section of U.S. manufacturing industries, consistent with the theory we develop, we find statistically significant effects of variations in inventory holdings and demand elasticities in the employment response to changes in technology. In particular, industries with larger inventory holdings and more elastic demand see a more pronounced increase in employment following a productivity improvement.

Motivated by Gali (1999), a considerable body of work in contemporary macroeconomics has argued that aggregate employment falls in response to permanent increases in productivity and, moreover, that this fall persists over time. While this notion remains controversial and the subject of an active research agenda, a negative response of employment to an improvement in productivity is often interpreted as evidence in favor of sticky prices. While most of this work is confined to the analysis of aggregate data, Chang and Hong (2006) find substantial variation in the way that employment responds to technology shocks in U.S. manufacturing industry data. When we shift the analysis to industry data, the number of observations on the employment effects of changes in technology increases significantly. Hence, one can more deeply explore how the heterogeneity in employment responses is affected by industry characteristics.

Sticky price models typically imply that employment falls in response to a positive technology shock. When productivity increases, fixed prices imply unchanged real sales so that less labor is required to meet a given level of nominal demand. Contrary to this conventional prediction, we show that when firms hold inventories, they can choose to expand output relative to sales in response to a favorable cost shock despite having sticky prices. They

[1] This ratio refers to quarterly final sales and inventories from the NIA Table 5.7.5.B and is taken from Haver.

[2] See Blinder (1981).

[3] Notable exceptions include Blinder (1981), Ramey (1991), and Bils and Kahn (2000).

would do so both to exploit relatively low production costs and to increase inventory stocks in anticipation of higher future sales.

To study the role of inventories in a model that allows for price rigidities, we introduce inventories into a standard Taylor (1980)-type framework. Consistent with the inventory literature, we assume that firms value inventories because they reduce the cost of possible stock-outs while allowing production smoothing over time (Ramey and West, 1999; Bils and Kahn, 2000). Within this setting, the flexible price model is nested as the case with no price staggering. In addition, we explore the role of differences in demand elasticities across industries. Specifically, we formalize the notion that when industry demand is sufficiently inelastic, permanent technological improvements lead to decreases in employment even when prices are fully flexible. Ultimately, we clarify that both the sticky price and flexible price framework are consistent with either decreases or increases in employment following a productivity improvement. Employment increases associated with such an improvement, however, always coincide with a build-up in inventories.

Given our model's predictions, we estimate the employment responses to productivity shocks in U.S. manufacturing industries using structural vector autoregressions (VARs) in the spirit of Gali (1999), and more recently Christiano et al. (2004). The VARs consist of total factor productivity (TFP), hours worked, and inventory holdings derived from the *NBER Manufacturing Productivity Database,* where productivity shocks are identified as the permanent component of TFP. We also estimate demand elasticities for all manufacturing industries. We then correlate the magnitude of an industry's short-run employment and inventory responses to productivity shocks with industry measures of demand elasticity, inventory-sales ratios, and price stickiness. To measure price stickiness, we use data on the average duration for which prices remain unchanged provided by Bils and Klenow (2004).

As predicted by our theoretical model, a positive (negative) response of employment to technological improvements coincides with a build up (decline) of inventories in many industries. Furthermore, we find that both the employment and inventory responses to a productivity shift are positively correlated with industry inventory holdings and demand elasticity in the cross section. We find less evidence, however, in support of price stickiness. Specifically, in the sample of industries for which the Bils and Klenow data are available, the effects of price stickiness are not statistically significant. We also find that the price decline associated with a permanent productivity increase is generally larger in the long run than in the short run.

Our analysis of inventories and how they affect the employment response to productivity shocks builds on a suggestion by Bils (1998). Accounting for changes in industry employment, Bils finds a positive and significant effect of the inventory-sales ratio interacted with changes in labor productivity. Besides providing an explicit theoretical justification for this argument, we connect the empirical work that emphasizes the use of disaggregated industry data (see Basu et al. 2004; Shea, 1999; Marchetti and Nucci, 2004; and Chang and Hong, 2006) with the structural VAR identification of productivity shocks using aggregate data.

This paper is organized as follows. Section 2 presents an industry model with staggered prices augmented to include inventory investment. In Section 3, we calibrate the model and study the employment response to productivity shocks across different degrees of goods storability, price stickiness, and demand elasticity. In Section 4, we explore the model implications in disaggregated data from the U.S. manufacturing industry. Section 5 concludes.

2. A Model with Staggered Prices and Inventories

We introduce inventories into a standard Taylor (1980)-type model with staggered prices. We study a partial-equilibrium industry model where monopolistically competitive firms may adjust their prices infrequently. Firms value inventories because they reduce the cost of possible stock-outs while allowing for production smoothing over time (Ramey and West 1999, Bils and Kahn 2000).

In a standard sticky price environment, a favorable productivity shock reduces the marginal cost of production and, consequently, prices. Due to the staggered price setting across firms, this adjustment is spread out over time so that the price level (output) only gradually decreases (increases) to its new long-run equilibrium path. Along the transition path, firms that can adjust their price set lower relative prices and, consequently, increase both their sales and production. In contrast, firms that cannot adjust their nominal prices see their relative prices increase and face declining sales. When goods are not storable, these firms then correspondingly reduce production and employment as in Gali (1999). When goods are storable, however, firms have an incentive to increase production, both to take advantage of the current relatively low marginal cost of production and to build up inventories in anticipation of higher future sales; indeed, these firms know that they will be able to lower their relative price at some future date. Therefore, despite prices being rigid, employment can increase across all firms in response to a positive productivity shock when the cost of holding inventories is sufficiently low. When holding inventories is very costly, because of high depreciation in storage for instance, work hours decrease in response to a productivity shock, as in the simple sticky-price model without inventories.

Without any price staggering, firms adjust their price at every date. In that case, a permanent improvement in productivity can lead to either an increase or decrease in employment depending on the elasticity of industry demand. In particular, if the demand curve is inelastic, sales do not change much with respect to prices and, therefore, improvements in productivity allow lower levels of labor input to be used. The reverse is true when demand is elastic so that employment increases after a positive technology shock. In principle, this mechanism is also operational when prices are sticky but we argue that without inventories, implausibly large demand elasticities are needed in order that improvements in technology generate increases in employment. Simply put, by virtue of fixing real demand in the short run, sticky prices largely mitigate the effects of a very elastic demand.

2.1. Industry Demand

Consider an industry where monopolistically competitive firms produce a continuum of differentiated products, $i \in [0,1]$. The industry output, q_t^I, is a CES aggregate of differentiated products, $\{q_t(i) : i \in [0,1]\}$:

$$q_t^I = \left[\int_0^1 q_t(i)^{(\theta-1)/\theta}\, di\right]^{\theta/(\theta-1)}, \qquad \theta > 1. \tag{2.1}$$

Given nominal prices for the differentiated products, $\{P_t(i) : i \in [0,1]\}$, the industry-price index, that is, the unit cost of industry output, is

$$P_t^I = \left[\int_0^1 P_t(i)^{1-\theta} di\right]^{1/(1-\theta)}, \tag{2.2}$$

and the demand for product i is

$$q_t(i) = \left[P_t(i) / P_t^I\right]^{-\theta} q_t^I. \tag{2.3}$$

We assume that an industry's demand depends on economy-wide aggregate demand, q_t^A, as well as its relative price, $p_t^I = P_t^I / P_t^A$, where P_t^A is the economy-wide aggregate price index,

$$q_t^I = \left(p_t^I\right)^{-\phi} q_t^A, \qquad \phi > 0. \tag{2.4}$$

Hence, we differentiate between an industry's demand elasticity, ϕ, and an individual firm's demand elasticity, θ, (which measures the elasticity of substitution across differentiated products in the industry). In particular, while profit maximization, in order to be well defined, requires the firm's demand to be elastic ($\theta > 1$), industry demand can be inelastic. We assume that economy-wide demand, q_t^A, is exogenous and that the economy-wide price index grows at a constant rate, $P_{t+1}^A / P_t^A = \mu > 1$.

2.2. Production

We base our model of production and inventory holdings on the empirical inventory investment literature (see Ramey and West, 1999). Firms choose a production path that minimizes total cost, or more specifically, production cost and the cost of making sales. In this subsection, we drop the firm index for ease of exposition.

A firm produces x_t units of a commodity using labor, h_t, as the single input to a decreasing-returns-to-scale production function

$$x_t = z_t h_t^{\alpha}, \qquad 0 < \alpha \leq 1, \tag{2.5}$$

with productivity z_t. We can think of production as being constant-returns-to-scale in capital and labor, with capital being costly to adjust relative to labor in the short run. Since our focus is on short-run adjustments in production, it is convenient to abstract from capital adjustment in our model.

We assume that making sales is costly and that only part of current production, denoted by y_t, makes up output that can potentially be sold. This output, along with beginning-of-period inventory, n_t, add up to total goods available for sales in the current period,

$$a_t = n_t + y_t. \tag{2.6}$$

Generating q_t units of sales requires $\kappa q_t (q_t/a_t)^{\eta}$, $\kappa, \eta \geq 0$, units of the commodity produced. The resource constraint on output that can potentially be sold is given by

$$y_t = x_t - \kappa q_t \left(\frac{q_t}{a_t} \right)^{\eta}. \tag{2.7}$$

Observe that having more goods available for sales, a_t, reduces the cost of making an actual sale (say, by reducing the probability of stock-outs). The particular functional form adopted for the cost of making sales implies an average inventory-sales ratio that is independent of scale, both with respect to production and sales. Bils and Kahn (2000), as well as Ramey and West (1999), document and emphasize this feature in a number of industries.

Firms hire labor in a competitive labor market at wage rate W_t. Together with the production function (2.5), this implies the production-cost function

$$c(x_t) = w_t \left(x_t / z_t \right)^{1/\alpha}, \tag{2.8}$$

where $w_t = W_t / P_t^A$ is the real wage. With $\alpha < 1$, this cost function is strictly increasing and strictly convex—that is, the marginal cost of production is strictly increasing. Goods not sold in a given period can be stored, but depreciate at rate δ,

$$n_{t+1} = (1-\delta)(a_t - q_t), \qquad 0 \leq \delta \leq 1. \tag{2.9}$$

Alternatively, we could have assumed that firms pay an inventory holding cost proportional to the inventory stock, but the depreciation in storage framework has the advantage of nesting the standard sticky price model with non-storable goods as a special case.

2.3. Optimal Production and Price Setting

A firm can only adjust its price every J periods, and in any period a fraction $1/J$ of firms can adjust their prices. Prices are fully flexible when $J = 1$. Each producer is indexed according to how much time, j, has elapsed as of time t since its previous price change. We restrict our attention to symmetric outcomes in which producers that adjust their price at the same time are identical. Thus there are J types of firms, and we index firms by $j \in \{0, 1, ..., J-1\}$. We denote by $n_{j,t}$ the beginning-of-period t inventory holdings of a firm that, j periods ago, set its price to $P_{0,t-j}$. The relative price of a type j firm at time t is denoted by $p_{j,t} = P_{j,t}/P_t^I$. Because of changes in the industry price level, a firm unable to adjust its price sees its relative price change over time, $p_{j+1,t+1} = p_{j,t}P_t^I/P_{t+1}^I$ for $j = 0$, ..., $J-1$.

Firms maximize the discounted present value of current and future real profits. These profits are discounted at the constant real interest rate $1/\beta$, $0 < \beta < 1$, and nominal profits are deflated using the aggregate price index P_t^A. The value of a type $j = 0$ firm that currently adjusts its nominal price, $V_{0,t}(n_0)$, depends on its beginning-of-period inventory holdings and the aggregate state of the industry.[4]

$$\begin{aligned} V_{0,t}(n_{0,t}) = \max_{p_{0t},\, q_{0t},\, a_{0t},\, x_{0t}, n_{1,t+1}} \quad & \left\{p_t^I p_{0,t} q_{0,t} - c_t(x_{0,t}) + \beta E_t V_{1,t+1}(p_{1,t+1}, n_{1,t+1})\right\} \\ \text{subject to} \quad & (2.3), \text{ and } (2.5) \text{ to } (2.9). \end{aligned}$$

The value of firms that cannot change the price of their products, $V_{j,t}(p_j, n_j)$, $j = 1, \ldots, J-1$, depends on their current relative prices, p_j, their beginning-of-period inventory holdings, n_j,

[4]Hereafter, we subsume the dependence on the aggregate state in the time index t.

and the aggregate state:

$$V_{j,t}(p_{j,t}, n_{j,t}) = \max_{q_{jt},\ a_{jt},\ x_{jt},\ n_{j+1,t+1}} \left\{ p_t^I p_{j,t} q_{j,t} - c_t(x_{j,t}) + \beta E_t V_{j+1,t+1}\left(p_{j+1,t+1}, n_{j+1,t+1}\right) \right\},$$
$$\text{subject to } \ (2.3), \text{ and } (2.5) \text{ to } (2.9),$$

$$V_{J-1,t}(p_{J-1,t}, n_{J-1,t}) = \max_{q_{J-1,t},\ a_{J-1,t},\ x_{J-1,t}\ , n_{0,t+1}} \left\{ p_t^I p_{J-1,t} q_{J-1,t} - c_t(x_{J-1,t}) + \beta E_t V_{0,t+1}\left(n_{0,t+1}\right) \right\}$$
$$\text{subject to } (2.3), \text{ and } (2.5) \text{ to } (2.9).$$

Optimal price setting satisfies the following first-order conditions,

$$\frac{\partial V_{j,t}}{\partial p_{j,t}} = p_t^I \left[q_{j,t} - \lambda_{j,t} \theta p_{j,t}^{-\theta-1} \tilde{q}_t \right] + \beta E_t \left[\frac{\partial V_{j+1,t+1}}{\partial p_{j+1,t+1}} \frac{P_t^I}{P_{t+1}^I} \right] \text{ for } j = 0, .., J-1 \quad (2.10)$$
$$\text{with } \frac{\partial V_{0,t}}{\partial p_{0,t}} = \frac{\partial V_{J,t}}{\partial p_{J,t}} = 0,$$

where $\lambda_{j,t}$ is the shadow value of additional sales (that is, the Lagrange multiplier on the sales constraint (2.3)). The shadow value of sales is equal to the real price minus the marginal cost of an additional unit sold, $\psi_{j,t}$,

$$\lambda_{j,t} = p_t^I p_{j,t} - \underbrace{c'_{j,t} \left\{ 1 + \kappa (1+\eta) \left(q_{j,t}/a_{j,t} \right)^{\eta} - \kappa \eta \left(q_{j,t}/a_{j,t} \right)^{1+\eta} \right\}}_{\psi_{j,t}}. \quad (2.11)$$

Equation (2.11) combines the first-order conditions for optimal sales and inventory holdings. The marginal cost of an additional sale reflects three components: the marginal cost of producing the commodity, $c'_{j,t}$, the marginal cost of making sales, $\kappa (1+\eta) \left(q_{j,t}/a_{j,t} \right)^{\eta}$, and the sales-facilitating effect of inventory holdings, $-\kappa \eta \left(q_{j,t}/a_{j,t} \right)^{1+\eta}$. We can repeatedly substitute expression (2.11) for the shadow value of sales in equation (2.10) to obtain an expression for the optimal price as

$$p_{0,t} = \left(\frac{\theta}{\theta - 1} \right) \frac{E_t \left[\sum_{j=0}^{J-1} \beta^j \psi_{j,t+j} q_{j,t+j} \right]}{E_t \left[\sum_{j=0}^{J-1} \beta^j \left(P_t^I / P_{t+j}^I \right) q_{j,t+j} \right]}. \quad (2.12)$$

Hence, price-adjusting firms set their price as a generalized markup over the marginal cost of sales. In a steady state with no inflation and no relative price changes, equation (2.12) reduces to a static markup over marginal cost. Expression (2.12) for the optimal price differs from the one that emerges in a standard Taylor-type sticky-price model only with respect to the definition of marginal cost. Thus, our analysis takes into consideration the marginal cost of sales rather than the marginal cost of production.

In contrast to conventional sticky-price models, inventories affect the marginal cost of additional sales in our framework. Specifically, the production-smoothing role of inventories can be observed in the optimal choice of $a_{j,t}$ and $n_{j+1,t+1}$,

$$c'_{j,t}\left\{1-\kappa\eta\left(q_{j,t}/a_{j,t}\right)^{1+\eta}\right\}=\beta\left(1-\delta\right)E_t\left[c'_{j+1,t+1}\right], \tag{2.13}$$

which equates the marginal increase in cost from producing one additional unit today, given no change in current sales, to the decrease in the discounted present value of marginal cost from producing one less unit tomorrow. Aside from production smoothing, note also that today's marginal cost of holding additional inventories implicitly reflects the sales facilitating effect of inventories, $-\kappa\eta\left(q_{j,t}/a_{j,t}\right)^{1+\eta}$.

2.4. Industry Equilibrium

We close the model by assuming that the real wage in the industry labor market is an increasing function of industry employment,

$$w_t=\bar{w}\bar{h}_t^{\gamma},\ \gamma>0. \tag{2.14}$$

Our industry labor supply specification, therefore, allows a permanent increase in industry productivity to cause a permanent increase in hours worked. We shall see in section 4.2 that hours appear to be non-stationary in most manufacturing industries.

The following section explores a recursive industry equilibrium where prices and quantities depend on the state of the industry. The state of the industry is characterized by the relative prices inherited from the previous period (by the firms that cannot adjust their nominal prices), $\{p_{j,t}:j=1,...,J-1\}$, the beginning-of-period inventory holdings of all firm types, $\{n_{j,t}:j=0,...,J-1\}$, and productivity, z_t. We define aggregate industry employment, production, sales, and inventory holdings respectively as

$$\bar{h}_t=\frac{1}{J}\sum_{j=0}^{J-1}h_{jt},\ \bar{y}_t=\frac{1}{J}\sum_{j=0}^{J-1}y_{jt},\ \bar{q}_t=\frac{1}{J}\sum_{j=0}^{J-1}q_{jt},\ \bar{n}_t=\frac{1}{J}\sum_{j=0}^{J-1}n_{j,t}. \tag{2.15}$$

3. The Response of Employment and Inventories to Productivity Shocks in the Model

We now study the industry's response to a permanent productivity shock. For this purpose, we linearize our model around its deterministic steady state. We show that when goods are storable, firms use inventories to smooth production and employment increases in response to a permanent productivity shock. This finding applies to various parametrizations of price rigidity and the sales-cost function. Consistent with conventional sticky-price models, we also show that if commodities depreciate rapidly while in storage, that is, if it is costly for firms to use inventories, industry employment responds negatively to a permanent productivity shock. Finally, we provide a description of how the employment response changes with inventory-sales ratios, the degree of price stickiness, and demand elasticities.

3.1. Calibration

Our benchmark parameter values are selected along the lines of other quantitative studies on business cycles. A time period represents a quarter. We assume a 4 percent annual real interest rate, $\beta = 0.99$. The labor elasticity of synthetic output, α, equals 2/3 and a firm's price elasticity of product demand, θ, is set to 10, which implies an 11 percent markup of price over marginal cost when prices are flexible. Because firms set their price as a markup over marginal cost, the labor-income share is lower than the output elasticity with respect to labor, $w \sum_j h_{j,t} / \sum_j p_{j,t} y_{j,t} = 0.6$. We normalize both economy-wide real aggregate demand and the industry relative price at one, and assume that the economy-wide price index increases at a fixed 4 percent annual rate.[5] For our baseline model, we assume that goods do not depreciate in storage, $\delta = 0$, that industry demand is elastic, $\phi = 1.5$, and that nominal prices remain fixed for 4 quarters, $J = 4$. Alternative parameterizations are also considered for robustness, in particular $0 \leq \delta \leq 0.95$, $0.5 \leq \phi \leq 5$, and $1 \leq J \leq 6$. Productivity, z_t, is assumed to follow a random walk with zero drift.

We use observations on average inventory-sales ratios to determine the scaling parameter κ in the sales-cost function, conditional on the sales-cost elasticity, η. For the baseline model, we use $\eta = 1$ and set the steady state inventory-sales ratio, n/q, to 0.2 at an annual frequency. If marginal cost were constant and equal across firms, the first-order conditions for inventories (2.13) would imply that the steady state sales-inventory ratio, q/a, is the same for all firm types,

$$\kappa\eta \left(q/a\right)^{1+\eta} = 1 - \left(1 - \delta\right)\beta. \tag{3.1}$$

Given the assumptions on the sales-cost elasticity and the annual inventory-sales ratio, which implies $q/a = 0.55$ at a quarterly frequency, this yields $\kappa = 0.0329$. In our economy, firms face increasing marginal costs so that sales-inventory ratios differ across firms, but the average annual sales-inventory ratio remains close to 0.2 under this parameterization. For alternative parameterizations of the inventory-sales ratio, we consider values in the range $0.05 \leq n/q \leq 0.4$, which cover 96 percent of average inventory-sales ratios across 458 4-digit

[5] Since real aggregate demand and the aggregate price index are exogenous, monetary policy does not accommodate productivity disturbances at the industry level. Dotsey (2002), and Galí et al. (2003), have shown that even with sticky prices, employment can respond positively to a productivity increase if monetary policy accommodates productivity shocks.

manufacturing industries over the period 1958-1996.[6] Finally, we assume that the elasticity of industry wage with respect to hours (the inverse of the labor-supply elasticity), γ, is 1.

3.2. Dynamic Effects of Productivity Shocks

Figure 1 shows the aggregate and individual-firm responses of sales, output, and employment to a permanent productivity increase when goods do not depreciate in storage ($\delta = 0$), and prices remain fixed for 4 periods. Observe first that in response to a productivity shock, industry employment increases on impact and converges to a higher steady state. Since the productivity shock lowers firms' marginal costs, firms that adjust prices at the time of the shock (type 0 firms) lower their nominal price and see their sales increase in the current as well as upcoming periods. To meet the immediate increase in demand, as well as to build inventories so as to lower the cost of making future sales, price-adjusting firms hire more labor and increase production in the current period (see the upper-left panel of Figure 1B). In contrast, firms that cannot change their nominal price see their relative prices rise and their current sales decline. However, type 3 firms (those that adjusted their price three periods ago), knowing that they will be able to lower their price in the near future, anticipate higher future sales. Hence, in an effort to build up inventories to meet these sales, they increase production and employment contemporaneously. Only firms that have adjusted their price in the last period (type 1 firms) significantly lower their employment on impact following the productivity shock. At the aggregate industry level, therefore, output, employment, and inventories all increase following a positive productivity shock when goods are durable.

In contrast, Figure 2 shows that when it is very costly to store goods due to a high depreciation rate in storage ($\delta = 0.9$), industry employment declines following a permanent increase in productivity. In this case, even anticipating higher future sales, firms whose prices are fixed at the time of the shock (types 1, 2, and 3) cannot effectively use inventories to smooth marginal cost since goods deteriorate rapidly while in storage. As their relative price increases on impact, sales and production decline and, consequently, they hire less labor. As depicted in Figure 2, output closely tracks sales under this scenario. Furthermore, output and employment increase only for firms that can change their price at the time of the productivity shock. At the industry level, therefore, employment declines. In this sense, our model with costly storage essentially reproduces the intuition underlying standard sticky-price models without inventories, including Gali (1999) and others.

[6] With regards to the sales-cost elasticity, we study parameterizations in the range $\eta \in [0.05, 2.5]$. Results are not sensitive with respect to this parameter.

3.3. Storability, Demand Elasticity, and Price Stickiness

Given the framework we have just laid out, the short-run hours response to a permanent technology disturbance varies with an industry's use of inventories, its demand elasticity, and the degree of nominal price rigidity in that industry. Figure 3 shows the cumulative one-year response of employment to a one-percent permanent increase in productivity for model solutions computed using various combinations of the parameters δ, κ, ϕ, and J.

Panels A.1 and A.2 of Figure 3 contrast two cases with different industry demand elasticities. In particular, employment responses are plotted for different price durations ($1 \leq J \leq 4$), and depreciation in storage ($0 \leq \delta \leq 0.9$), according to whether industry demand is inelastic (panel A.1, $\phi = 0.75$), or elastic (panel A.2, $\phi = 1.5$). For each choice of depreciation rate, δ, we adjust κ so as to generate an annual inventory-sales ratio of 0.2, as in our baseline case. When demand is relatively inelastic in Panel A.1, the short-run response of hours to a productivity increase is always negative irrespective of price duration and depreciation in storage. As the depreciation rate increases, the hours response becomes more negative the greater the degree of price stickiness. When demand is relatively elastic in Panel A.2, the short-run hours response to a productivity increase is always positive when goods do not depreciate in storage, $\delta = 0$, irrespective of price stickiness. The hours response is always positive when prices are fully flexible, $J = 1$. When prices are sticky, $J = 4$, small values of the depreciation rate are enough to generate negative employment responses.

To sum up, when prices are fully flexible, permanent improvements in productivity can either increase or decrease short-run employment depending on the elasticity of industry demand. When prices are sticky, improvements in productivity can either increase or decrease short-run employment depending on firms' ability to carry inventories.

Panels B.1 and B.2 plot employment responses for different degrees of price duration, $J = 1, 2, 3$, and annual inventory-sales ratios, $0.05 \leq n/q \leq 0.4$, when industry demand is elastic and inelastic, respectively. We adjust the sales cost κ in order to generate different annual inventory-sales ratios. When demand is inelastic, $\phi = 0.75$ in panel B.1, the employment response is negative whether prices are sticky or flexible. In contrast, when demand is elastic, $\phi = 1.5$ in panel B.2, an industry's employment response is always positive irrespective of price stickiness. Observe that in both cases, however, an industry's short-run employment response to a productivity improvement always increases with its average inventory-sales ratio.

Panels C.1 and C.2 depict the effects of industry demand elasticity, ϕ, and price stickiness, J, on hours worked when goods do not depreciate in storage, $\delta = 0$, and when goods perish rapidly in storage, $\delta = 0.9$. When $\delta = 0$, irrespective of price stickiness, employment responds positively to an improvement in productivity as long as industry demand is sufficiently elastic,

$\phi > 1$. When goods are more perishable, with a price duration of 2 quarters, an industry demand elasticity of 2 is required to obtain a positive employment response. When prices are fixed for a year, $J = 4$, a larger demand elasticity, $\phi = 6$, is required in order that employment increase in response to a positive technology shock. Irrespective of depreciation in storage and price stickiness, note that the employment response always increases with the magnitude of industry demand elasticity.

4. The Response of Employment and Inventories to Productivity in U.S. Manufacturing

As we have just argued, an industry's employment response to a productivity shock depends in principle not only on the extent of price stickiness in that industry, but also on its ability to carry inventories and the elasticity of demand for the industry's product. To investigate these implications in the data, we now turn our attention to the joint behavior of productivity, employment, and inventories in U.S. manufacturing industries. For each industry, we estimate the short-run employment response to a permanent productivity increase. We then relate that employment response to measures of demand elasticity, inventory holdings, and price stickiness.

4.1. Data

We obtain data on industry productivity, employment, and inventory holdings from the *NBER Manufacturing Productivity Database* (see Bartelsman and Gray, 1996). The *Database* includes annual information on 458 4-digit SIC manufacturing industries over the period 1958-1996, and largely reflects information from the *Annual Survey of Manufacturing.*

To measure industry productivity, some adjustment to the industry TFP growth series of the *Database* is necessary. The cost of labor in the *Database* only includes wages and salaries but does not include fringe benefits. Therefore, we scale each 4-digit industry's payments to labor based on the ratio of employer FICA payments to wages and salaries in the corresponding 2-digit industry. Given the adjusted wage income share, we recalculate the industry's TFP growth series following the procedure described in Bartelsman and Gray (1996).

Industry employment is the sum of aggregate hours worked by production and non-production workers in that industry. Unfortunately, data on the average workweek of non-production workers is not available. We follow the *Database*'s convention and set the work-week of non-production workers to 40 hours. Our measure of real inventory stocks is the total end-of-period nominal inventory stock deflated by the value of shipments price deflator.

We estimate the price elasticity of industry demand, $\hat{\phi}$, from a regression of industry real sales growth on the rate of industry price change. To account for the simultaneity problem intrinsic to such regressions, we instrument for the rate of price change using variables that reflect supply shocks, namely TFP growth and energy price changes. While energy prices are most likely exogenous relative to demand shifts, their correlation with output prices is limited since they make up only a small portion of total cost. In contrast, TFP growth rates are more strongly correlated with industry output price changes but, due to unobserved factor utilization, they may not represent strictly exogenous supply shocks. Wages and material prices are likely to be poor instruments since they are non-trivially affected by industry demand shifts when the industry has a significant share in factor markets. While our price elasticity measures vary with the choice of instruments, our cross-industry regressions below are not significantly affected by the particular choice of demand elasticity estimate. As a practical matter, the most statistically significant estimates are obtained with our benchmark instruments, TFP growth rates, and energy price changes.

We obtain statistically significant demand elasticity estimates in 343 out of the 458 industries available. There are only 18 industries (or less than 4 percent of all industries) whose demand elasticities are estimated with an incorrect sign (i.e., the demand elasticity is positive). For those industries, we set the price elasticity to 0 in our cross-industry regressions.[7] For the 440 industries whose estimated demand elasticities are negative, estimates range from -7.30 to -0.04, with a mean of -1.52 and a median of -1.26.

We control for the effects of inventory holding costs on the employment response to changes in technology by using industry average inventory-sales ratios, $\overline{n/q}$. While inventory-sales ratios are endogenous in the model, without a production smoothing motive, long-run inventory-sales ratios reflect only inventory technology parameters that are unrelated to demand (i.e., unrelated to ϕ and θ), or average price duration, J.[8] This feature of the model limits the potential endogeneity problem associated with using average inventory-sales ratios as a control variable rather than some more direct measure of inventory holding costs.

Finally, we measure the length of time over which prices are fixed (J) using the average duration for which product prices remain unchanged as documented by Bils and Klenow (2004). The Bils-Klenow price data are based on unpublished price quotes collected by the Bureau of Labor Statistics over the period 1995-1997. At the 4-digit level, we obtain measures of average price duration for 110 manufacturing goods. Across these goods, price duration, weighted by industry average output, is on average 4.6 months.

[7] Our results are not affected by this transformation.

[8] Recall equation (3.1) which holds when firms face constant marginal costs.

4.2. Empirical Specification

For each industry, we estimate a baseline structural VAR in the growth rates of industry TFP, $\Delta\ z_t$, hours worked, Δh_t, and end-of-period inventory holdings, Δn_{t+1}. We then calculate the response of hours worked to a permanent increase in productivity. Following the original work of Galí (1999), we identify permanent productivity shocks by assuming that measured productivity is non-stationary and that productivity shocks are the only source of movements in long-run productivity. Following Chang and Hong (2006), we use TFP rather than labor productivity as our measure of productivity. At the industry level, labor productivity is not an appropriate measure of productivity since shocks with permanent effects on labor productivity confound true productivity shocks with other shocks. Put another way, permanent changes in industry labor productivity result not only from permanent productivity changes, but also from permanent changes in input ratios stemming from permanent changes in relative prices.[9]

We do not require hours worked or inventories to be stationary at the industry level. While current empirical evidence on the stationarity of aggregate hours worked is mixed, the assumption that aggregate per capita hours worked are stationary is often motivated by reference to balanced growth path considerations.[10] At the industry level, however, a permanent change in the productivity of a given industry leads to a permanent change in relative productivity across industries. Therefore, for aggregate hours to remain unchanged, industry hours must change permanently to reflect the change in relative productivity. Using standard Dickey-Fuller unit root tests across all 458 industries, we reject non-stationarity of hours worked and inventory stocks at the 10 percent significance level in only one out of twenty industries. Thus, we do not restrict hours or inventories to be stationary and write our industry VARs in first-differences of TFP, hours worked, and inventory holdings.

Our estimation procedure can be summarized as follows. Consider the vector MA repre-

---

[9] The same argument can be made for the aggregate data used by Galí (1999). The identification assumption that only permanent productivity shocks have permanent effects on labor productivity is based on a balanced growth argument within the neoclassical growth model. Implicitly, one assumes that other permanent changes to the return on capital, such as permanent changes to the capital income tax rate, have only a minor impact on the (productivity-normalized) capital-labor ratio. See for instance Uhlig (2004). At the industry level, however, one expects that permanent changes in relative prices and factor rentals are large.

[10] This was first discussed by Shapiro and Watson (1988). For some recent and opposing views, see Christiano et al. (2003), and Francis and Ramey (2002).

sentation of our three variable system,

$$\underbrace{\begin{bmatrix} \Delta z_t \\ \Delta h_t \\ \Delta n_t \end{bmatrix}}_{Y_t} = \underbrace{\begin{bmatrix} a_{11}(L) & a_{12}(L) & a_{13}(L) \\ a_{21}(L) & a_{22}(L) & a_{23}(L) \\ a_{31}(L) & a_{32}(L) & a_{33}(L) \end{bmatrix}}_{A(L)} \underbrace{\begin{bmatrix} \varepsilon_{z,t} \\ e_{h,t} \\ e_{n,t} \end{bmatrix}}_{\varepsilon_t} \tag{4.1}$$

where $\varepsilon_{z,t}$ and $\{e_{h,t}, e_{n,t}\}$ denote productivity and non-productivity shocks respectively. The disturbance terms are assumed to be orthogonal, that is $E\varepsilon_t\varepsilon_t' = \Sigma$ is diagonal, and the polynomial in the lag operator need not be finite. We do, however, assume that a VAR representation exists, that is $A(L)$ in (4.1) is invertible, $B(L) = A(L)^{-1}$. Our main identifying restriction – that only productivity shocks affect industry TFP in the long run – implies that $a_{12}(1) = a_{13}(1) = 0$. The responses of hours worked and inventories to a productivity shock identified in this way will be independent of the other shocks. Thus, there is no need to disentangle the non-productivity shocks through additional identifying assumptions. As in Gali (1999), we simply impose a block triangular structure on the inverse matrix of long-run multipliers in the VAR, $B(1)$, as a way to orthogonalize the reduced form errors $\{e_{h,t}, e_{n,t}\}$. The triangular nature of $B(1)$, and the fact that Σ is diagonal, allows us to estimate each equation in $B(L)Y_t$ recursively. All VARs are estimated using a one-year lag and standard errors are computed by bootstrapping based on 500 draws.

Since industries may have experienced different degrees of technological change over time, we normalize the productivity shocks across industries. In particular, for each industry we scale the productivity shock such that it increases TFP by one percent in the long run. We denote industry i's short-run response of hours worked (i.e. the first-year response) to this normalized productivity shock by SR^h.

We then perform cross-industry regressions that relate the estimated short-run response of employment, SR^h, to industry characteristics, X, such as average inventory-sales ratios, demand elasticity, and the average duration of output prices,

$$SR^h = \alpha + \mathbf{Xb} + u. \tag{4.2}$$

Note that the dependent variable in (4.2), SR^h, is only an estimate of the true short-run employment response to a technology shock. Since this uncertainty introduces heteroskedasticity into our cross-industry regressions, we use Huber-White robust standard errors for inference.[11]

[11] In addition to the use of Huber-White standard errors, we also estimated equation (4.2) with industry observations weighted by the inverse of the standard error corresponding to that industry's short-run employment response. In other words, industries with less precise estimates of SR^h are given less weight. This

4.3. Employment and Inventories in a Cross-Section of Manufacturing Industries

The results from our estimated structural VARs are best illustrated by comparing the impulse response functions for two industry aggregates that differ widely in their ability to store their output and in their estimated demand elasticities: raw foods (meat and dairy products) on the one hand, and durable goods on the other hand. The raw foods industry has relatively low inventory holdings and faces a relatively inelastic demand: industry inventories are 3.5 percent of annual sales, and the mean estimated demand elasticity is 0.5.[12] For comparison, in aggregate manufacturing, inventories represent 14 percent of annual sales and the mean estimated demand elasticity is 1.04. Conversely, the durable goods industry has relatively high inventory holdings and faces a relatively elastic demand: inventories represent 18 percent of annual sales, and the mean estimated demand elasticity is 1.64.[13]

Figure 4 shows the responses of TFP, hours worked, and inventory holdings to a permanent productivity increase in each of our benchmark industries.[14] In the raw foods industry (Figure 4A), hours worked unambiguously fall following a one-standard-deviation productivity shock, while inventories show little movement. Our theory suggests that these responses for raw foods arise because of their inelastic demand and, to the degree that prices are sticky, because inventory holding costs are high in that industry. Conversely, in the durable goods industry (Figure 4B), hours worked increase significantly after a productivity shock and a significant inventory build-up takes place. Given the theory we have laid out, we naturally expect these findings in that the relative ease with which durable goods can be carried as inventories mitigates any negative effects of price stickiness on employment. Moreover, even if prices were fully flexible, the fact that durable goods have elastic demand helps produce a positive response of employment to productivity shocks.

According to our industry VARs estimated across 458 4-digit manufacturing industries, following a productivity shock that raises industry long-run TFP by 1 percent, hours worked increase in 302 industries. While the estimated hours response vary considerably across industries, from -2.77 to 3.59 percent, the mean employment increase is 0.32 percent, (see the first row of Table 1).[15] In Table 2, the short-run response in hours, SR^h, is highly

additional weighting scheme does not have a substantive effect on our regression results.

[12] The estimated demand elasticities range from 0.2 to 1.19 (in absolute value) with a median of 0.41.

[13] For the durable goods aggregate, we estimate negative demand elasticities in 253 out of 260 industries. Estimated demand elasticities range from 0.06 to 7.26 in absolute value with a median of 1.39.

[14] When we estimate structural VARs for the two industry aggregates, "raw foods" and "durable goods," TFP growth, Δz_t, is aggregated as: $\Delta z_t = \sum_i \mu_{i,t} \Delta z_{i,t}$, where $\Delta z_{i,t}$ and $\mu_{i,t}$ are TFP growth and the output share of industry i in the industry aggregate. Industry aggregate employment and inventory stocks are the sum of the component industries employment and inventories.

[15] This finding stands in sharp contrast to Kiley (1998) who finds a negative correlation between hours and the permanent component of labor productivity across 2-digit U.S. manufacturing industries. As pointed out above, in our structural VARs with industry data, TFP rather than labor productivity is the theoretically

correlated with the short-run response of inventories, SR^n, with a correlation coefficient of 0.57 across all industries. The employment response is also positively correlated with both average inventory-sales ratios (0.13) and industry demand elasticities (0.22), as predicted by the theory.

Results from cross-industry regressions, shown in Table 3, indicate that both average inventory-sales ratios and estimated demand elasticities have significant explanatory power in accounting for differences in employment responses to technology shocks. Moving from an industry with an average inventory-sales ratio of 0.16 (the sample mean), to one with an average inventory-sales ratio of 0.43 (a one-unit log increase), increases the employment response to an improvement in technology (which ultimately raises industry TFP by one percent) from 0.32 percent to 0.52 percent. Similarly, moving from an industry with a demand elasticity of 1.04 (the sample mean) to 2.04 (a one-unit increase), raises the employment response to a normalized technology shock from 0.32 to 0.54 percent. Both these effects are highly significant, with p values well below 1 percent.

To explore the effects of price stickiness, we are forced to use the smaller sample size for which the Bils and Klenow (2004) measure of average price duration is available, 111 out of 458 industries. In this smaller sample, a regression of the short-run employment response to technology shocks against average inventory-sales ratios, demand elasticities, and price stickiness, shown in Table 3, Column (2), continues to identify average inventory-sales ratios as a significant explanatory variable, with a coefficient of 0.14 compared to 0.20 in the larger sample (Table 3, Column 1). The coefficient on demand elasticity, however, is now close to zero and ceases to be significant. This result reflects the fact that the correlation between the employment response to technology shocks and industry demand elasticity declines from 0.22 in the large sample to 0.04 in the restricted sample. In excluding more than 3/4 of our original data sample, therefore, important information is omitted. In this regression, the average price duration of industry products appears unrelated to the way employment responds to technology shocks. While this finding somewhat undermines the importance of price stickiness in determining firms' employment response to technology shocks, the small sample size of the Bils and Klenow (2004) data set prevents us from viewing this finding as conclusive.

In general, we retain from our findings, based on the full sample of 458 observations, that an industry's ability to carry inventories and its demand elasticity significantly affect the way employment and inventories respond to productivity disturbances. In particular, industries

preferred measure of productivity. Chang and Hong (2006) show that this distinction is also empirically relevant. Indeed, for the same data, they find mostly negative employment responses to identified permanent productivity changes when labor productivity is used as a measure of technology. They attribute this predominantly negative response to permanent changes in relative input use.

that are either able to carry large inventories or in which demand elasticity is high tend to display large positive employment and inventory responses to permanent productivity increases.

5. Further Considerations

5.1. Alternative VAR Specifications

It should be clear that our model has implications not only for employment but also for such variables as sales, inventory-sales ratios, and prices that are not included in our baseline 3-variable VAR. Although our data encompasses a large number of industries, the fact that each industry contains only about forty observations confines us to relatively small scale VARs. Therefore, to evaluate the robustness of our results, we estimate three alternative 3-variable VARs that continue to include productivity and employment, but we substitute either sales, inventory-sales ratios, or prices for inventories. Data on industry sales and prices can be directly obtained from the *NBER Database*. Real sales is defined as the value of industry shipments deflated by the value of shipments price deflator, and we use the latter deflator as our industry price measure.

We find that the short-run employment responses to a permanent productivity increase in these alternative 3-variable VARs generally mimic those of our baseline specification. First, the distribution of short-run employment responses across industries is very similar for the four different VAR specifications. Employment tends to increase in most industries, and the mean employment increase in response to a 1 percent permanent TFP increase is about 0.3 percent across VAR specifications. Second, the industry short-run employment responses are highly correlated among the different VAR specifications, with correlations ranging from 0.80 to 0.85. Third, the impact of inventory-sales ratios and estimated demand elasticities on the cross-industry variation of short-run employment responses is qualitatively and quantitatively very similar to those shown in Table 2. In Table 4, Columns (1)-(4), practically all coefficients are significant at the 1 percent level. We see the only exception to this pattern in the alternative VAR estimates with the price series as the third variable (Column 4). The employment responses from this price-based VAR tend to be smaller, but even this VAR yields predominantly positive employment responses.

We find that in response to a permanent productivity increase, sales tend to increase more than inventories, and inventory-sales ratios tend to fall. This feature replicates the stylized facts for unconditional correlations of sales and inventory-sales ratios, cf. Ramey and West (1999) and Bils and Kahn (2000). We also find, however, that the countercyclical behavior of inventory-sales ratios is not uniform: in 94 out of the 458 industries, inventory-sales ratios rise following a productivity increase.[16]

Finally, in most industries, prices fall in response to a permanent increase of TFP. Overall, prices decrease in both the short run and the long run in 411 out of 458 industries. Furthermore, in 361 out of these 411 industries, the long-run price decline exceeds the decline in short-run prices. The mean decline in prices following a normalized technology shock is 0.56 percent in the short run, and 0.98 percent in the long run across all industries. Note that the long-run magnitude of the mean fall in prices is almost equal to the long-run increase in TFP. These observations are suggestive of an industry supply that is more elastic in the short run than in the long run.

5.2. Searching for a More Direct Measure of Inventory Holding Costs

We have argued that inventory holdings play a significant role in accounting for cross-sectional differences in firms' employment response to technology shocks in manufacturing data. While long-run inventory-sales ratios are independent of demand considerations absent a production-smoothing motive in our model, variations in industry demand elasticity and the volatility of demand likely affect these ratios in practice. To avoid this potential endogeneity problem requires that one obtain a more direct exogenous measure of inventory holding costs.

To address this issue, we examine a data set provided by Bils and Klenow (1998) on the average service-life of products. This data set is based on information from the Bureau of Economic Analysis and the interoffice memorandum used by major U.S. property casualty insurance companies. We augment the Bils-Klenow service-life data by assuming that raw foods and processed foods (e.g., canned and frozen foods) can be stored for 0.2 years and 0.5 years, respectively. Specifically, raw foods include meat products (SIC 2011, 2013, 2015), dairy products (2021, 2022, 2023, 2024, 2026), and bakery products (2051, 2052). Processed foods include the rest of food industries under SIC 20. In total, we are able to obtain service life measures for 98 manufacturing products at the 4-digit level, with a mean weighted by average industry output of 4.5 years, Table 1. In the sample of industries for which we have average service life data, the correlation between service life and inventory-sales ratios is high and positive (0.44, Table 2), which suggests that service life might constitute a useful instrument for average inventory-sales ratios.

In Table 5, Column (1), we re-estimate the regression of employment responses to technology shocks on inventory-sales ratios and industry demand elasticity but, in doing so, we instrument for inventory-sales ratios using Bils and Klenow's service life data. We find that the coefficient on average inventory-sales ratios increases considerably, from 0.2 to 0.68.

[16] Additional tables are available upon request.

However, the coefficient on demand elasticity now becomes statistically insignificant. The increase in the standard error associated with the demand elasticity coefficient arises from the fact that, while the service life data is correlated with average inventory-sales ratios, it is even more strongly correlated with industry demand elasticity (0.59, Table 2). The second stage regression, therefore, is subject to non-trivial multicollinearity. In a model with durable goods consumption, goods that are more durable are typically associated with a larger investment response to a change in their price. The reason is that following a price change, a lower depreciation rate helps create a larger change in the implicit rental rate of the durable good. As shown in Table 5, Column (2), the multicollinearity issue we have just described also applies if we were to use average service life as an instrument for demand elasticity instead of average inventory-sales ratios. In this case, of course, it is the coefficient on demand elasticity that increases considerably, from 0.22 to 0.53, but so does the estimated standard error on the average inventory-sales ratios coefficient. Ultimately, the ideal instrument for our model should be correlated with inventory-sales ratios but uncorrelated with industry demand elasticity.

6. Conclusion

Recent studies have interpreted the employment response to productivity shifts as a litmus test for the existence of frictions in nominal price adjustment. This paper has shown that whether or not inventories can be used to break the link between production and sales is crucial for understanding firms' employment response to productivity shocks. In conventional sticky price models without inventories, productivity shocks reduce employment. In contrast, the same shocks cause firms that can carry inventories to expand output relative to sales and, as a result, to hire more workers. For quantitatively reasonable calibrations, we show that following a productivity shock, employment increases when the inventory holding costs are sufficiently low and industry demand is elastic.

Based on 458 4-digit disaggregated U.S. manufacturing data over the period 1958 to 1996, we estimate how firms respond to productivity shocks. Consistent with our theory, we find that an industry's employment response to shifts in productivity is strongly correlated with its standard characteristics, such as inventories and demand elasticities, that are entirely within the domain of flexible price models. On the other hand, using direct evidence on the degree of price rigidities across industries, we could not find a link between price stickiness and the short-run employment response to productivity shocks.

References

[1] Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde. 2002. "Technology Shocks and Aggregate Fluctuations," mimeo, Northwestern University.

[2] Bartelsman, Eric J., and Wayne Gray. 1996. "The NBER Manufacturing Productivity Database," NBER Technical Working Paper 205.

[3] Basu, Susanto, Miles Kimball, and John Fernald. 2004. "Are Technology Improvements Contractionary?" Federal Reserve Bank of Chicago Working Paper 2004-20.

[4] Bils, Mark 1998. "Discussion." In *Beyond Shocks: What Causes Business Cycles* Conference Series No. 42, Federal Reserve Bank of Boston, 256-263.

[5] Bils, Mark, and James A. Kahn. 2000. "What Inventory Behavior Tells Us About Business Cycles," *American Economic Review* 90, 458-481.

[6] Bils, Mark, and Peter Klenow. 1998. "Using Consumer Theory to Test Competing Business Cycle Models," *Journal of Political Economy* 106, 233-261.

[7] Bils, Mark, and Peter Klenow. 2004. "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy* 112, 947-985.

[8] Blinder, Alan. 1981. "Inventories and the Structure of Macro Models," *American Economic Review* 71, 11-16.

[9] Chang, Yongsung, and Jay H. Hong. 2006. "Do Technological Improvements in Manufacturing Raise or Lower Employment?" *American Economic Review* 96, 352-368.

[10] Chari, V.V., Patrick Kehoe, and Ellen McGrattan. 2004. "Are structural VARs useful guides for developing business cycle theories?" Unpublished Manuscript, Federal Reserve Bank of Minneapolis.

[11] Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson. 2003. "What Happens After a Technology Shock?" mimeo, Northwestern University.

[12] Dotsey, Michael. 2002. "Structure from Shocks," *Federal Reserve Bank of Richmond Economic Quarterly* 88, 36-48.

[13] Francis, Neville, and Valerie A. Ramey. 2005. "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited," *Journal of Monetary Economics* 52, 1379-99.

[14] Galí, Jordi 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89, 249-271.

[15] Galí, Jordi, J. David Lopéz-Salido, and Javier Vallés 2003. "Technology Shocks and Monetary Policy: Assessing the Fed's Performance," *Journal of Monetary Economics* 50, 723-743.

[16] Kiley, Michael. 1998. "Labor Productivity in U.S. Manufacturing: Does Sectoral Comovement Reflect Technology Shocks?" mimeo.

[17] Marchetti, Domenico, and Francesco Nucci. 2004. "Price Stickiness and the Contractionary Effect of Technology Shocks," *European Economic Review*, forthcoming.

[18] Ramey, Valerie A., and Kenneth D. West. 1999. Inventories. In John B. Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics* 1B, 863-923, Elsevier Science: Amsterdam.

[19] Shapiro, Matthew D., and Mark W. Watson. 1988. "Sources of Business Cycle Fluctuations." In Stanley Fisher (ed.), *NBER Macroeconomics Annual 1988*, Cambridge and London: MIT Press, 111-148.

[20] Shea, John. 1999. "What Do Technology Shocks Do?" In Ben S. Bernanke and Julio J. Rotemberg (eds.), *NBER Macroeconomics Annual 1998*, Cambridge and London: MIT Press, 275-310.

[21] Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy* 88, 1-23.

[22] Uhlig, Harald. 2004. "Do Technology Shocks Lead to a Fall in Total Hours Worked?" *Journal of the European Economic Association* 2, 361-371.

Table 1. Summary Statistics

	Mean	S.D.	Min	Max	obs.
SR^h	.32	.77	-2.77	3.59	458
SR^n	.85	.98	-2.05	6.32	458
$\overline{n/q}$	.16	.07	.02	.55	458
$\hat{\phi}$	1.04	.76	0	4.85	458
J	4.3 mo.	3.3	.77	29.9	111
δ^{-1}	4.5 yrs.	5.64	.2	27.5	98

Note: SR^h and SR^n are the short-run responses of hours and inventories obtained from our baseline industry VAR. $\overline{n/q}$ is the average inventory-sales ratio; $\hat{\phi}$ the estimated industry demand elasticity; J: the average duration of output price (in months) for 111 industries; δ^{-1} the average service life (in years) of the output for 98 industries.

Table 2 Cross-Industry Correlations

	SR^h	SR^n	$\ln \overline{n/q}$	$\hat{\phi}$	$\ln J$	$\ln \delta^{-1}$
SR^h	1.00					
SR^n	0.57	1.00				
$\ln \overline{n/q}$	0.13	0.05	1.00			
$\hat{\phi}$	0.22	0.14	0.03	1.00		
$\ln J$	0.07	0.01	0.1	0.03	1.00	
$\ln \delta^{-1}$	0.21	0.10	0.44	0.59	0.2	1.00

Note: The variables are as defined in Table 1. Correlations with the log average price duration ($\ln J$) are based on 111 industries. Correlations for the log average service life ($\ln \delta^{-1}$) are based on 98 industries. Correlation between $\ln J$ and $\ln \delta^{-1}$ is based on 72 industry products.

Table 3. The Employment Response to Productivity Shocks:
The Effects of Inventory-Sales Ratio, Demand Elasticity, and Price Stickiness

	(1)	(2)
$\ln(\overline{n/q})$	0.20^{***}	0.14^{*}
	(0.07)	(0.07)
$\hat{\phi}$	0.22^{***}	0.01
	(0.06)	(0.07)
$\ln J$		0.07
		(0.09)
Nobs	458	111
R^2	0.06	0.02

Note: The dependent variable is SR^h, an industry's short-run response of hours to a 1 percent permanent productivity increase estimated from our baseline 3-variable industry VAR. The right-hand side industry variables are as defined in Table 1. The numbers in parenthesis are standard deviations based on robust Huber-White sandwich estimators of variance. Asterisks denote p-values of less than 10% (*), less than 5% (**), and less than 1% (***). All regressions include a constant term, not displayed in the Table.

Table 4. Alternative VARs: The Response of Employment to Productivity Shocks:

	(1)	(2)	(3)	(4)
Dep. Var.	SR_n^h	SR_q^h	$SR_{n/q}^h$	SR_p^h
$\ln(\overline{n/q})_i$	0.20***	0.14**	0.15***	0.05
	(0.07)	(0.06)	(0.06)	(0.05)
$\hat{\phi}$	0.22***	0.18***	0.18***	0.19***
	(0.06)	(0.05)	(0.05)	(0.04)
R^2	0.06	0.05	0.06	0.06

Note: The dependent variable SR_i^h is an industry's short-run response of hours to a 1 percent permanent productivity increase in a 3-variable VAR when the third variable is inventories ($i = n$), sales ($i = q$), inventory-sales ratio ($i = n/q$), or price ($i = p$). The right-hand side variables are as defined in Table 1. The regressions are as described in Table 3.

Table 5. The Employment Response to Productivity Shocks:
The Effects of Inventory-Sales Ratio and Demand Elasticity

	(1)	(2)
$\ln(\overline{n/q})$	0.68**	0.03
	(0.32)	(0.13)
$\hat{\phi}$	0.12	0.53**
	(0.17)	(0.19)
Nobs	98	98

Note: See Table 3. In Column (1), the log of average service life ($\ln \delta^{-1}$) is used to instrument for $\ln \overline{n/q}$. In Column (2), the log of average service life is used as an instrument for demand elasticity ϕ.

Figure 1. The Response to a Permanent Productivity Increase Without Depreciation in Storage, $\delta = 0$

A. Aggregate Response

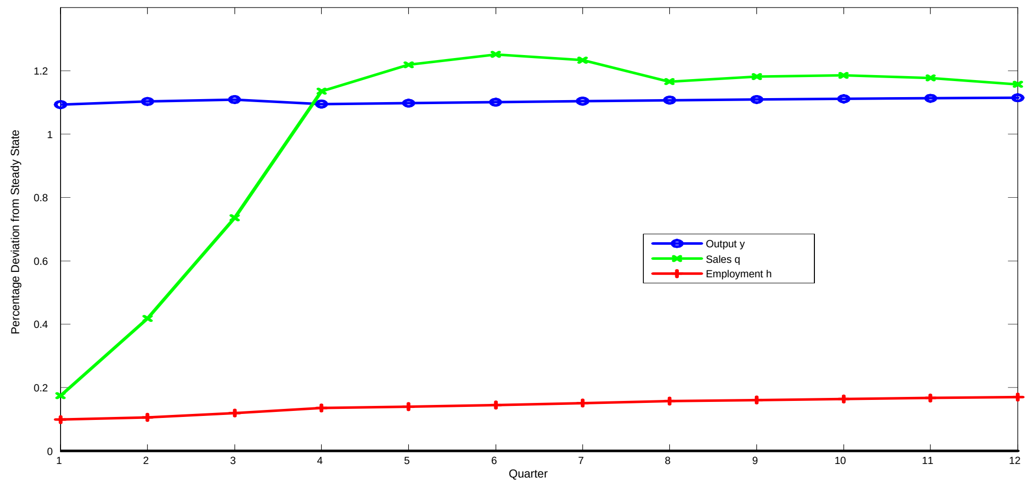

B. Individual Firms' Response

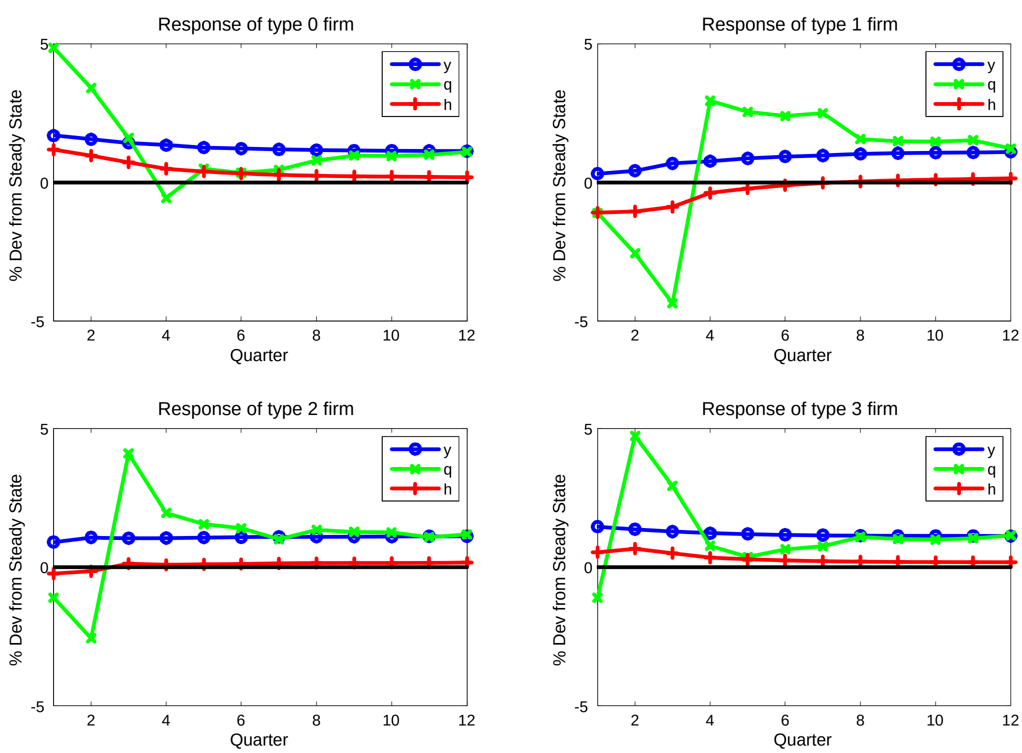

Figure 2. The Response to a Permanent Productivity Increase With Depreciation in Storage, $\delta = 0.9$

A. Aggregate Response

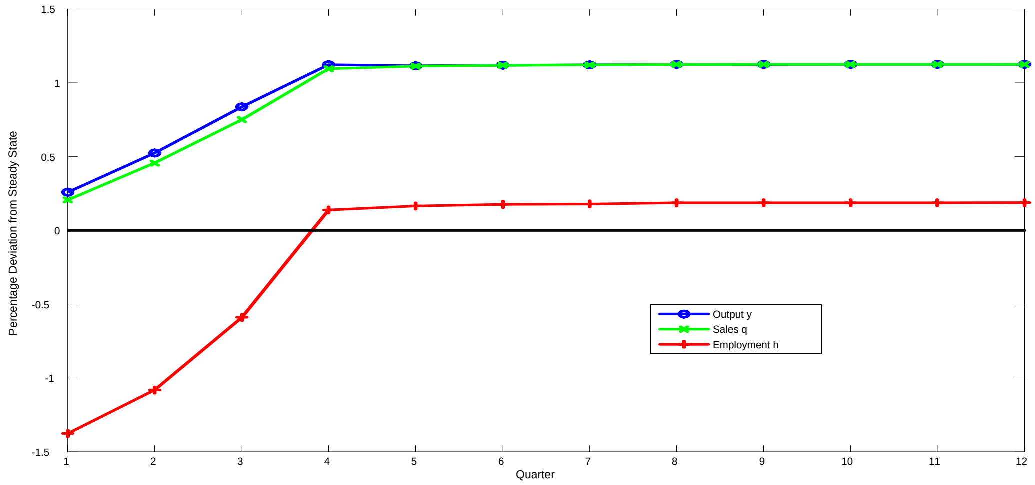

B. Individual Firms' Response

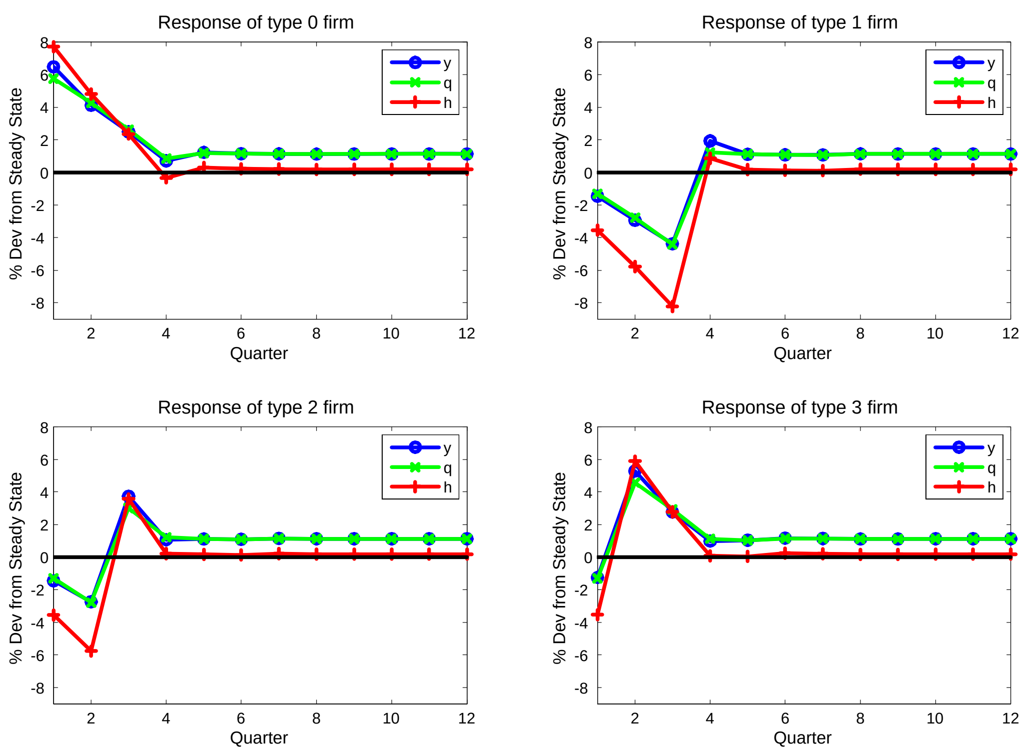

Figure 3. Employment Responses to Productivity Shocks

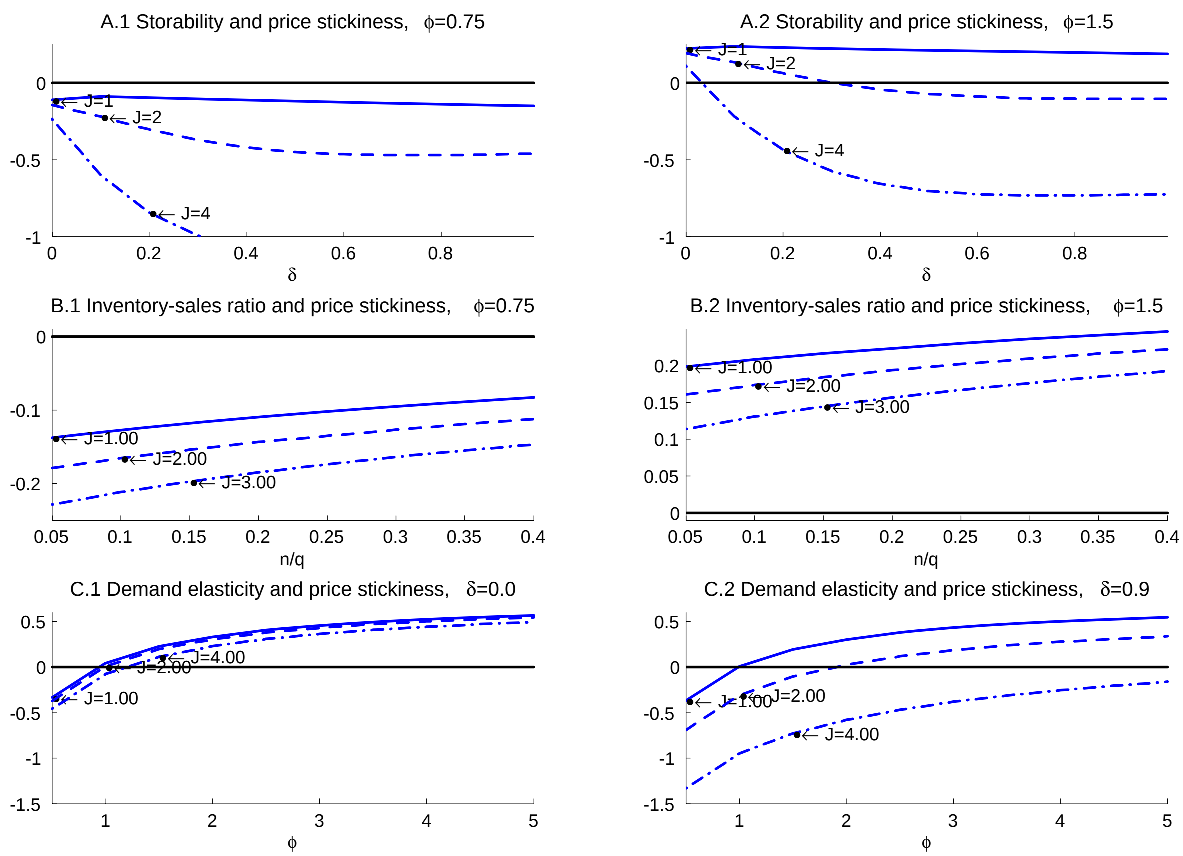

Note: We graph the cumulative one-year response of hours to a 1 percent permanent increase in productivity. All parameter values are set at their baseline values with the exception of $(\delta, \phi, J, \kappa)$ as described in the text. In Panel A, we vary δ and J for inelastic (A.1, $\phi = 0.75$) and elastic (A.2, $\phi = 1.5$) demands. In Panel B, we vary the steady state inventory-sales ratio n/q and J for inelastic (B.1, $\phi = 0.75$) and elastic (B.2, $\phi = 1.5$) demands. In Panel C, we vary ϕ and J for low (C.1, $\delta = 0$) and high (C.2, $\delta = 0.9$) depreciation rates.

Figure 4. Response to Technology from VAR

Panel A. Raw Food Manufacturing

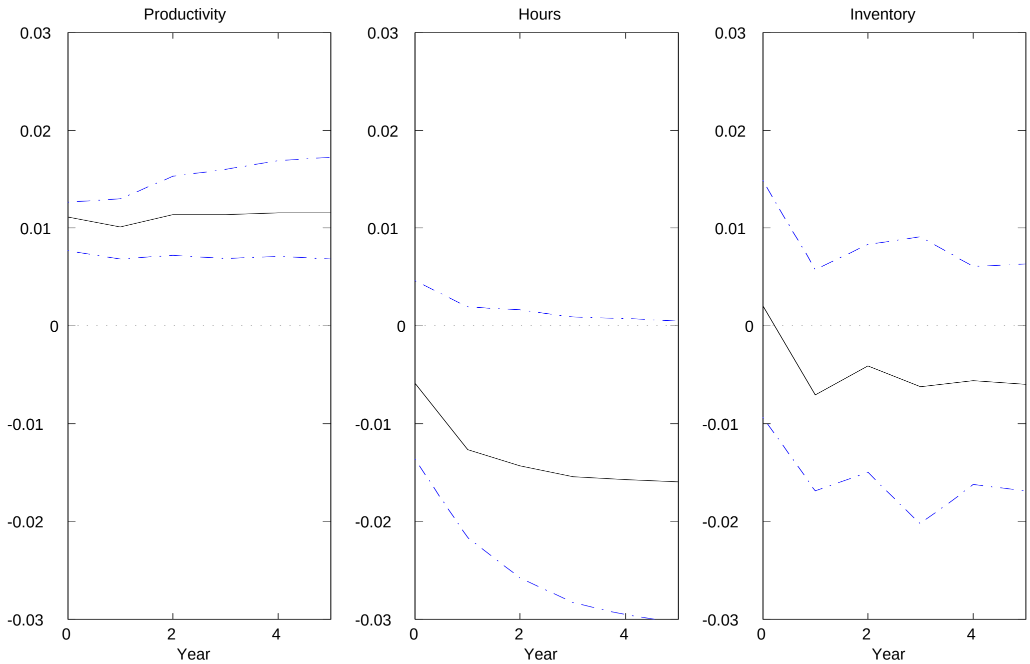

Panel B. Durable Goods Manufacturing