Productivity, Employment, and Inventories*

Yongsung Chang, Andreas Hornstein, and Pierre-Daniel Sarte
Research Department
Federal Reserve Bank of Richmond
P.O. Box 27622, Richmond, VA 23261

June 10, 2004

Abstract

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 sticky-price models. In a Taylor-type sticky-price model with inventories, we show that the employment response to a productivity shock depends on the extent to which goods are storable. Whereas in conventional sticky-price models without inventories, productivity shocks reduce employment, the same shocks cause firms in our economy to expand output relative to sales, build up inventories and, as a result, hire more workers. We then estimate the employment response to productivity shocks in disaggregated U.S. manufacturing data from 1958 to 1996. Consistent with our theory, we find that an industry's employment response to productivity shifts is strongly correlated with its inventory holdings and the storability of its products

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

*We would like to thank Bennett McCallum for helpful comments. 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: `yongsung.chang@rich.frb.org`, `andreas.hornstein@rich.frb.org`, and `pierre.sarte@rich.frb.org`.

1. Introduction

The stochastic growth model lies at the core of modern business-cycle theory, and in almost all of its representations, the model predicts that employment increases, at least temporarily, in response to a permanent increase in productivity. Yet a growing body of work argues that the behavior of aggregate labor productivity and employment in the postwar United States, as well as in other OECD countries, is at odds with this prediction. For instance, Galí (1999) argues that following a permanent productivity increase, employment declines in these countries, and this decline often persists over time. Based on this observation, Galí then argues that sticky prices must constitute a key feature of the economic environment. With sticky prices and given a level of nominal demand, higher productivity allows firms to meet the level of real demand with less labor.

The basic intuition for Galí's argument in favor of sticky prices depends on the assumption that all goods are perishable. Since this is clearly not the case in a large number of industries where inventories account for a significant fraction of sales, this paper explores a simple sticky-price model with inventories. We show that the way employment responds to a productivity shock depends crucially on the extent to which goods are storable. Thus, if firms can hold inventories, they may choose to expand output relative to sales in response to a favorable cost shock. They would do so both to exploit relatively low production costs and to increase inventory stocks up to higher anticipated levels of sales. We find that following a productivity shock, employment increases (decreases) when the depreciation rate on goods in storage is sufficiently low (high). We then estimate the employment response to productivity shocks using disaggregated U.S. manufacturing data from 1958 to 1996. Consistent with our theory, we show that an industry's employment response to a productivity shift is strongly correlated with its inventory holdings and the storability of its products.

Our theoretical framework introduces inventories into a model with Taylor (1980)-type staggered prices. In line 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).

In our 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. When goods are storable, however, firms stuck with fixed prices 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 not too high. 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.

To make ideas more concrete, we estimate the employment responses to productivity shocks in manufacturing industries using trivariate vector autoregressions (VARs) consisting of total factor productivity (TFP), hours worked, and inventory holdings derived from the *NBER Manufacturing Productivity Database*. We identify productivity shocks as the permanent components of TFP. We then correlate the magnitude of an industry's short run employment response to productivity shocks with measures of both storability and industry price stickiness. To measure storability, we use data on the average service life of an industry's products obtained from Bils and Klenow (1998), as well as data on an industry's average inventory-sales ratio. To measure price stickiness, we use data on the average duration for which prices remain unchanged from Bils and Klenow (2002).

An analysis of the manufacturing industry data shows that the employment response to productivity shocks varies substantially across industries. Consistent with the predictions of our theoretical model, a positive (negative) response of employment to a productivity shock coincides with a build up (decline) of inventories in many industries. Furthermore, an industry's employment response is strongly correlated with the storability of the industry's output. For example, in an industry producing highly perishable goods such as meat and dairy products, where firms carry very little inventory, hours worked fall significantly in response to a productivity shock. In contrast, in durable goods manufacturing, where the average inventory stock is as large as 18 percent of annual sales, we see a strong employment and inventory build-up following a productivity increase. We also find a small negative correlation between the employment response to a productivity shock and the degree of price stickiness across industries. However, the effects of product storability across industries dominate those of price stickiness in accounting for the cross-sectional heterogeneity of employment responses.

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, Fernald and Kimball, 1998; Shea, 1999; and Chang and Hong, 2004) with the structural

VAR identification of productivity shocks using aggregate data. Basu et al. (1998) report a negative correlation between hours worked and their productivity measure adjusted for cyclical utilization. Shea (1998) finds that an increase in the orthogonal component of R&D and patents tends to increase input use, especially labor, in the short run, but reduces inputs in the long run. Chang and Hong (2004) find that in a majority of U.S. manufacturing industries, hours worked increase in response to a permanent increase in industry TFP.[1]

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 storability and price stickiness. In Section 4, we explore the model implications in disaggregated data from the U.S. manufacturing industry. Section 5 concludes.

2. A Sticky-Price Model with Inventories

We begin by describing a partial-equilibrium industry model where monopolistically competitive firms adjust their prices infrequently in a staggered fashion. Unlike other theoretical work on sticky prices, we assume that goods are storable and that firms can hold inventories. Inventory holdings allow firms to smooth their production over time in response to shocks. In particular, following a productivity shock that lowers its marginal cost, a firm may choose to increase employment in order to build up inventories if it anticipates higher sales in the future, and despite the fact that its current sales may actually decline because of fixed nominal prices.

2.1. Industry Demand

Consider an industry where monopolistically competitive firms produce a continuum of differentiated products $i \in [0, 1]$. Real aggregate demand for the industry products, $\tilde{q}_t$, exhibits constant elasticity of substitution θ between the differentiated products, $\{q_t(i) : i \in [0, 1]\}$:

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

[1] Galí (1999)'s empirical work has recently been disputed on the grounds of mis-specifications resulting from omitted varibale bias (Altig, Christiano, Eichenbaum, and Linde 2002), as well as the over-differencing of hours (Christiano, Eichenbaum, and Vigfusson 2003). While Francis and Ramey (2002) find evidence in support of Galí, and the stationarity of hours remains controversial (Shapiro and Watson 1989), our focus is not on these issues per se but rather accounting for heterogeneity across industries.

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

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

The demand for product i is

$$q_t(i) = [P_t(i)/P_t]^{-\theta}\,\tilde{q}_t. \tag{2.3}$$

For simplicity, we assume that nominal industry demand Q_t is exogenous,

$$P_t\tilde{q}_t = Q_t, \tag{2.4}$$

and nominal demand grows at a constant rate $Q_{t+1}/Q_t = \mu > 1$.

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

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

where z_t represents TFP. 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 helpful to abstract from capital adjustment entirely 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,

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

We further assume that to generate 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 then 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 for 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 $w_t = W_t/P_t$ 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,

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

2.3. Optimal Production and Price Setting

A firm can adjust its price only every J periods, and in any period a fraction $1/J$ of firms can adjust their price. 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 $p_{j,t} = P_{j,t}/P_t$. Because of changes in the aggregate price level, a firm unable to adjust its price sees its relative price change over time, $p_{j+1,t+1} = p_{j,t}P_t/P_{t+1}$ for $j = 0,\ ...,\ J-2$.

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. 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.[2]

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

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, and the aggregate state:

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

Optimal price setting satisfies the following first-order conditions,

$$\begin{aligned} \frac{\partial V_{j,t}}{\partial p_{j,t}} &= q_{j,t} - \phi_{j,t}\theta p_{j,t}^{-\theta-1}\tilde{q}_t + \beta E_t\left[\frac{\partial V_{j+1,t+1}}{\partial p_{j+1,t+1}}\frac{P_t}{P_{t+1}}\right] \quad \text{for } j = 0, .., J-1 \\ \text{with } \frac{\partial V_{0,t}}{\partial p_{0,t}} &= \frac{\partial V_{J,t}}{\partial p_{J,t}} = 0, \end{aligned} \tag{2.10}$$

where $\phi_{j,t}$ is the shadow value of additional sales (i.e. the Lagrange multiplier on the sales constraint (2.3)). The shadow value of sales is equal to its real price (2.3) minus the marginal cost of additional sales, $mc_{j,t}$,

$$\phi_{j,t} = p_{j,t} - \underbrace{c'_{j,t}\left\{1 + \kappa(1+\eta)(q_{j,t}/a_{j,t})^{\eta} - \kappa\eta(q_{j,t}/a_{j,t})^{1+\eta}\right\}}_{mc_{j,t}}. \tag{2.11}$$

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, $\kappa(1+\eta)(q_{j,t}/a_{j,t})^{\eta}$, and the sales-facilitating effect of inventory holdings, $-\kappa\eta(q_{j,t}/a_{j,t})^{1+\eta}$. We can repeatedly substitute for the shadow value of sales in equation (2.10) to obtain an expression for the optimal price as

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

Hence, adjusting firms set their price as a generalized markup over the marginal cost of sales.

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

In a steady state with no inflation and no relative price changes, equation (2.12) reduces to a static markup over marginal cost. This expression 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 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. Market Clearing

We define industry employment, production, sales, and inventory holdings respectively as

$$\bar{h}_t=\frac{1}{J}\sum_{j=0}^{J-1}h_{jt},\ \bar{q}_t=\frac{1}{J}\sum_{j=0}^{J-1}q_{jt},\ \bar{y}_t=\frac{1}{J}\sum_{j=0}^{J-1}y_{jt},\ \bar{n}_t=\frac{1}{J}\sum_{j=0}^{J-1}n_{j,t}. \tag{2.14}$$

To close the model, we posit a real wage that is an increasing function of industry employment,

$$w_t=\bar{w}\bar{h}_t^{\gamma},\ \gamma>0. \tag{2.15}$$

Note that our industry labor supply specification 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 industries of the manufacturing sector.

3. The Hours Response 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 do not depreciate while in storage, employment increases in response to a permanent productivity shock. This finding is robust across various parameterizations 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, industry employment responds negatively to a permanent productivity shock. Finally, we provide a description of how the employment response changes with inventory-sales ratios and the degree of price stickiness.

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 the 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 labor elasticity of output, $w \sum_j n_{j,t} / \sum_j p_{j,t} y_{j,t} = 0.6$. We assume that the nominal demand for industry products increases at a fixed 4 percent annual rate.[3] For our baseline model, we assume that goods do not depreciate in storage, $\delta = 0$, and that nominal prices remain fixed for 4 quarters, $J = 4$. For alternative parameterizations, we consider the range of $0 \leq \delta \leq 0.95$ and $2 \leq J \leq 6$. Productivity, z_t, is assumed to follow a random walk with zero drift.

We use observations on average inventory-sales ratios to determine the scale 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 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,

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

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. In later sections, we check for robustness with respect to the sales-cost elasticity and the inventory-sales ratio. For the sales-cost elasticity, we consider values in the range $\eta \in [0.05, 2.5]$. For the inventory-sales ratio, we consider values in the range $n/q \in [0.05, 0.4]$ at an annual frequency, which covers 96 percent of average inventory-sales ratios across 458 4-digit manufacturing industries over the period 1958-1996.Finally, we assume that the elasticity of industry wage with respect to hours (the

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

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, hours worked, and inventory holdings 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 (during which their relative prices remain low due to a gradual adjustment of average prices). 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 2 and type 3 firms, knowing that they will be able to lower their price in the near future, anticipate higher future sales. Hence, they increase production and employment contemporaneously in an effort to build up inventories to meet these sales. Only firms that have adjusted their price in the last period (type 1 firms) 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 over time. Since goods deteriorate rapidly while in storage, their incentives to hold inventories is limited. 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.

The positive employment response to a permanent productivity increase when goods are storable turns out to be robust to variations in the sales-cost elasticity parameter, η, and to variations in the average inventory-sales ratio, n/q. Conditional on the sales-cost elasticity and the steady-state inventory-sales ratio, we use equation (3.1) to pin down the scale parameter κ. We then solve the model and plot the cumulative one-year response of employment to a one-percent permanent increase in productivity. In this simulation, all other parameters are set to their baseline values, $\delta = 0$ and $J = 4$.

Figure 3A shows that in response to a one percent permanent productivity increase, employment always increases during the first year when goods are highly durable, irrespective of the sales-cost elasticity parameter and the steady-state inventory-sales ratio. The percentage increase of employment ranges between 0.45 and 0.70 depending on the parameter. Given the sales-cost elasticity, employment responds more strongly the larger the industry's average inventory-sales ratio. Given the inventory-sales ratio, a productivity shock induces a larger employment response for smaller values of the sales-cost elasticity.

3.3. Durability and Price Stickiness

Since inflexible prices represent a key feature of our analysis, Figure 3B depicts the cumulative-one-year response of employment to a productivity increase across various degrees of price stickiness and average inventory-sales ratios. Variations in the average inventory-sales ratio are directly governed by the scale parameter κ, while all other parameters are set to their benchmark values. With highly storable goods, the employment response is positive for a wide range of price stickiness ($1 \leq J \leq 6$), as well as annual inventory-sales ratios ($0.05 \leq n/q \leq 0.4$). The strength of the response, however, diminishes the longer firms are stuck with a fixed nominal price. As the degree of price rigidity increases, firms are able to smooth production in response to their increasing marginal costs over a longer period. This effect dampens their desire to build inventories and thus mutes the initial increase in production. This muted increase in production in turn translates into a smaller increase in hours worked.

In addition to the degree of nominal price rigidity, our model relates the hours response to the storability of an industry's products. As goods depreciate more rapidly in storage, firms' incentive to hold inventories to facilitate future sales disappears. Figure 3C shows the model's employment responses to a permanent productivity increase across different inventory-sales ratios ($0.05 \leq n/q \leq 0.4$), and depreciation rates ($0 \leq \delta \leq 0.9$). The more storable goods are, the stronger is the employment response. This finding arises irrespective of the industry's average inventory-sales ratio. As the depreciation rate δ increases, the employment response diminishes and eventually becomes negative. As δ approaches one, goods in storage perish completely and firms can no longer carry inventories. Firms with fixed nominal prices see their relative price rise at the time of a productivity increase, and then decrease production to meet a correspondingly lower demand. Our model thus reverts back to a standard sticky-price framework where productivity shocks are contractionary.

Finally, Figure 3D shows employment responses for different degrees of price stickiness ($1 \leq J \leq 6$), and storability ($0 \leq \delta \leq 0.9$). For each choice of depreciation rate, we adjust κ so as to generate an annual inventory-sales ratio of 0.2, our baseline value. We can see that the importance of storability for how employment reacts to a productivity shock is not independent of the degree of nominal price rigidity. For example, the employment response remains positive for depreciation rates as high as 90 percent when prices are set for two quarters. However, when nominal prices remain fixed for four quarters, hours worked respond positively to a favorable productivity shock only if the quarterly depreciation rate is less than 10 percent.

4. The Hours Response to Productivity in U.S. Manufacturing

We have just argued that with sticky prices a firm's employment response to a productivity shock depends crucially on its ability to carry inventories. To investigate this implication 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 storability and price stickiness.

4.1. Data

Industry productivity, employment, and inventory holdings are obtained 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*. We use the TFP growth series directly available from the *Database* as our measure of productivity. To capture employment, we use industry aggregate hours worked, calculated as the sum of production and non-production worker hours. Unfortunately, data on the average workweek of non-production workers is not available. We follow the *Database*'s convention, therefore, and set the workweek of non-production workers to 40 hours.

For each industry, the Database provides a measure of the total end-of-period nominal inventory stock. Inventories are not disaggregated by stage of production, that is, they include not only finished goods, but also goods in process and materials. This potentially creates two complications for our analysis. First, since the inventory measure is comprehensive, it is not obvious what price deflator to use in order to obtain real inventory stocks. To deflate nominal inventory stocks, we use the Database value of shipments price deflator. Second, our theory emphasizes the role of finished goods inventories, or more specifically, the role of finished goods storability in affecting the employment response to a productivity shock. We should bear in mind, however, that throughout most of our analysis, our measure of inventories captures more than just finished goods. In the Appendix we report results from using an alternative price index to deflate inventories and from a smaller sample of 2-digit SIC industries for which we have finished goods inventory data.

We use two measures of storability. First, we use the average service life of products in Bils and Klenow (1998). The Bils-Klenow service-life data are based on information from the Bureau of Economic Analysis and the interoffice memorandum used by major U.S. property casualty insurance companies. In addition to the Bils-Klenow data, we assume that raw foods (e.g., meat and dairy products) 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 storability measures for 98 manufacturing products at the 4-digit level with an average service life, weighted by the average industry output, of 4.2 years.

We use the average service life of goods only as a proxy for firms' ability to carry inventories. The relevant measure of depreciation based on our theory is that of goods in storage, but service life reflects depreciation of goods in use. Nevertheless, one expects the degree to which a good can be stored to be strongly related to the average service life of that good. We consider average inventory-sales ratios as our second measure of storability. Our theory indeed implies that the inventory-sales ratio is positively related to the storability of goods, see for example equation (3.1).

To measure the degree of price stickiness across industries, we use the average duration for which product prices remain unchanged as documented by Bils and Klenow (2002). 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 price stickiness for 110 manufacturing goods. Across these goods, the average price stickiness, weighted by industry average output, is 4.6 months. Overall, we have measures of both storability and price stickiness for 72 industries.

4.2. Empirical Specification

For each industry, we estimate a structural VAR in industry TFP, z_t, hours worked, h_t, and end-of-period inventory holdings, n_{t+1}, and then calculate the response of hours worked to a permanent increase of productivity. Following earlier empirical work by Galí (1999), we identify permanent productivity shocks through the assumption that measured productivity is non-stationary and that productivity shocks are the only source of movements in long-run measured productivity. Following Chang and Hong (2004), we use TFP rather than labor productivity as our measure of productivity. Within our VAR framework, labor productivity is not an appropriate measure of productivity at the industry level because shocks with permanent effects on labor productivity confound true productivity shocks with other shocks. Put another way, a permanent change in industry labor productivity can result not only from a permanent true productivity change but also from permanent changes in input ratios caused by permanent changes in relative prices.[4]

We do not require hours worked or inventories to be stationary at the industry level. The current empirical evidence on the stationarity of aggregate hours worked is mixed.[5] On theoretical grounds, the assumption that per capita hours worked are stationary at the aggregate level is often motivated by reference to balanced growth path properties. 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. We do not believe that unit root tests are entirely convincing given that our data set contains only about 40 data points per industry. That said, using standard Dickey-Fuller unit root tests, we reject non-stationarity of hours worked and inventory stocks at the 10% significance level for 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 representation of our three variable system,

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

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. Our main identifying restriction, namely

[4] In principle, the same argument could be made for aggregate data as 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. At the industry level one would expect, however, that permanent changes of relative prices and factor rentals are large.

[5] 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).

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, $A(1)^{-1}$, as a way to orthogonalize the reduced form errors $\{e_{h,t}, e_{n,t}\}$. Under the assumption that $A(L)$ in (4.1) is invertible, the triangular nature of $A(1)^{-1}$ and the fact that Σ is diagonal allows us to estimate each equation in $A(L)^{-1}S_t$ recursively. All VARs are estimated using a one year lag and standard errors are computed by bootstrapping based on 500 draws.

4.3. The Short-Run Response of Hours Worked to Productivity Shocks in Manufacturing

Our theoretical framework emphasizes the role of inventories in that they facilitate making sales and provide an opportunity to smooth marginal cost over time. Even though current sales are predetermined from the point of view of firms that cannot adjust their price, these firms may nevertheless choose to increase production in response to a favorable productivity shock, and thus hire more labor, to build up inventories in the anticipation of future sales.

We saw earlier that the theory's prediction depends crucially on firms ability to carry inventories. Hence, we now present two extreme cases with respect to storability: highly perishable raw foods (meat and dairy products) on the one hand, and aggregate durable goods on the other.[6] Given the extreme difference in the nature of storability for the two products, one expects stark differences in their inventory-sales ratios. Indeed, the average inventory-sales ratio is only 3.5 percent of annual sales in the raw food products industry. In contrast, this ratio is 18 percent for aggregate durable goods.[7] In Figures 4A and 4B, we plot the responses of TFP, hours worked, and inventory holdings to a productivity increase based on our VARs estimated for the two industries.[8] In the raw food industry (Figure 4A), hours worked unambiguously fall following a one-standard-deviation productivity shock while inventories show little movement. In sharp contrast, in durable manufacturing industries (Figure 4B), hours worked rise significantly after a productivity shock. Consistent with our

[6] Here, aggregate "durable" industry refers to the industries conventionally classified as durable (SIC 24-25 and 32-39), where the criterium is a durability of products longer than one year. Note, however, that when we use detailed disaggregated data on average service life (the Bils-Klenow measure), some products conventionally classified as non-durable show a service life longer than one year: for example, 4.1 years for men's suits and coats (SIC 2311).

[7] For aggregate manufacturing, the inventory-annual sales ratio is 0.14.

[8] Similar to the industry data we estimate structural VARs for the industry aggregates, 'perishables' and 'durables'. Aggregate 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.

theory, one can also see a significant inventory build-up in this case.

Having illustrated the effect of storability for two extreme examples, we now investigate whether cross-industry differences in storability can more generally account for some of the cross-industry heterogeneity in the way that employment responds to productivity shocks. Since industries may have experienced different degrees of technological change over time, we normalize the productivity shocks across industries. We thus consider a productivity shock that increases TFP by one percent in the long run, instead of the conventional one-standard-deviation shock, in all cases. We denote industry i's short-run response of hours worked (i.e. the first-year response) to this normalized-productivity shock by SRR_i^h.

Figure 5 plots the short-run response of employment to a permanent productivity increase against our two measures of storability: average service life, $\ln DUR_i$, and average inventory-sales ratios, $\ln \left(\overline{n/q}\right)_i$. First, note that following a productivity increase, hours worked rise in the majority of industries: $SRR_i^h > 0$ in 66 out of 98 industries for which we have measures of average service life, and in 342 out of all of the 458 manufacturing industries.[9] Second, it appears that storability, in particular service life more so than the inventory-sales ratio, has significant explanatory power in accounting for substantial cross-industry differences in the employment response to a productivity shock. Table 1 presents findings from various cross-sectional regressions that support this notion. Looking at Column (1), a one-percent increase in average service life is associated with a 0.167 (with t-ratio of 4.39) percentage point increase in the short-run response of employment to a favorable productivity shock. In an industry producing a commodity whose average service life is 4.2 years (the sample mean), hours worked increases by 0.45 percent in the first year following a productivity shock that increases industry TFP by one percent in the long run. In an industry associated with an average service life of 9 years (one-standard deviation longer than the sample mean), the hours increase is distinctly larger, about 0.57 percent over the course of a year.

Our second measure of storability, the inventory-sales ratio, is also positively related to the magnitude of the hours response to a productivity increase. Looking at Column (2), we find that across 458 manufacturing industries, an average inventory-sales ratio that is one percent larger results in a 0.191 (with a t-ratio of 3.06) percentage-point larger short-run response of employment. In Column (3), when both the average service life and the

[9] 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, the use of TFP over labor productivity as a measure of productivity is more appropriate in our structural VARs based on industry data on theoretical grounds alone. Chang and Hong (2004) further 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 they use labor productivity as a measure of productivity. Furthermore, they can attribute this predominantely negative response to permanent changes in relative input use.

average inventory-sales ratio are included in the regression, the effects of service life dominate those of the inventory-sales ratio. This is not entirely surprising since average service lives and inventory-sales ratios are correlated across industries — Table 2 shows a correlation coefficient 0.41 across 98 industries. More importantly, other variables than storability affect the inventory-sales ratio, and at the same time they affect the employment response to productivity shocks. To deal with this endogeneity of the inventory-sales ratio we use the product's service life as an instrument. A product's service life is likely to be exogenous, and our results so far suggest that it is a good proxy for storability. Indeed, when we use the average service life as an instrument for the average inventory-sales ratio, Column (4), we find a significant and strong positive effect of storability, 0.917 (with a t-ratio of 3.23), on the short-run employment response.

Since price rigidity constitutes a key aspect of our theoretical framework, we now add a measure of price stickiness to our regressions of the short run employment response on storability. For our service life measure of storability, Column (5), we find a negative but small and statistically insignificant effect of price stickiness, and the coefficient on storability remains positive and statistically significant. For our inventory-sales ratio measure of storability, we consider two regressions. The results from an OLS cross-sectional regression in Column (6) are similar to the results with the service life measure of storability: the coefficient on price stickiness is negative but small and statistically insignificant and the coefficient on storability does not change much.[10] In the second regression, Column (7), the results improve dramatically once we instrument for the inventory-sales ratio using the service life. The coefficient on the inventory-sales ratio declines somewhat relative to the regression excluding price stickiness, Column (4), but remains positive and significant. Furthermore, the coefficient on price stickiness remains negative but becomes large in absolute value and is statistically significant, -0.452 (with a t-ratio of -1.6).[11] These findings all accord well with the theoretical predictions made earlier (recall Figure 3B).

Our theoretical model suggested that firms could hire more labor, increase production, and build up inventories in response to a favorable productivity shock, even in the face of inflexible prices. We calculate a cross-sectional correlation between the responses of hours and inventories to a productivity shock of 0.42 across 458 manufacturing industries (see Table 2), thus supporting the mechanisms outlined in our model. Based on U.S. manufacturing data, therefore, we conclude that an industry's ability to carry inventories plays a critical

[10] Note that there are only 72 industries for which both storability and price stickiness measures are available. Thus the t-ratio on the coefficient associated with the inventory-sales ratio is smaller due to a restricted sample. A regression using 111 industries for which price stickiness measures are available yields similar results.

[11] Using Italian manufacturing industry data, Marchetti and Nucci (2003) find that industries with stickier prices tend to show more negative response of hours to productivity shocks.

role in the way that it responds to productivity disturbances. In particular, industries that are able to carry large inventories — as a result of their products' storability — tend to display a large positive employment response to a permanent productivity increase.

Previously we have mentioned two potential measurement problems, the choice of price deflator for nominal inventories and the fact that measured inventories not only include finished goods, but also goods in process and materials. For the inventory price deflator we have considered using the materials price index as an alternative deflator. We find that the choice of price deflator does not materially affect our results, see Table A.1 in the Appendix. We have also studied 2-digit SIC manufacturing industries for which we have data on finished goods inventories. The results for finished goods inventories in 2-digit SIC industries are reported in the Appendix. The results for 2-digit industries are not as strong in support of our theory, but we should point out that for 2-digit SIC industries we are essentially looking at two samples with ten and sixteen observations only, compared with our samples of 72 and 98 observations for 4-digit SIC industries. To summarize, these robustness checks do not suggest that the two measurement problems require major qualifications of our results.

5. Conclusion

We have 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 sticky-price models. In a Taylor-type sticky-price model with inventories, we have shown that the employment response to a productivity shock depends in large part on the extent to which goods are storable. Whereas in conventional sticky price models without inventories, productivity shocks reduce employment, the same shocks cause firms in our economy to expand output relative to sales, build up inventories and, as a result, hire more workers. For quantitatively reasonable calibrations, we saw that following a productivity shock, employment increases (decreases) when the depreciation rate on goods in storage is sufficiently low (high). We then estimated the employment response to productivity shocks from disaggregated U.S. manufacturing data over the period 1958 to 1996. Consistent with our theory, we found that an industry's employment response to shifts in productivity was strongly correlated with its inventory holdings and the durability of its products

References

[1] Altig, David, Lawrence 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. 1998. "Are Technology Improvements Contractionary?" International Finance Discussion Paper No. 625, Board of Governors of the Federal Reserve System.

[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 Kahn, James A. 2000. "What Inventory Behavior Tells Us About Business Cycles", *American Economic Review* v.90 no.3, 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 Klenow, Peter 2002. "Some Evidence on the Importance of Sticky Prices," mimeo.

[8] Chang, Yongsung, and Jay H. Hong. 2004. "Do Technological Improvements in Manufacturing Raise or Lower Employment?" mimeo, Federal Reserve Bank of Richmond.

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

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

[11] Francis, Neville, and Valerie Ramey. 2002. "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited" NBER Working Paper No. 8726.

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

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

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

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

[16] Shapiro, Matthew D. and Mark W. Watson. 1988. "Sources of Business Cycle Fluctuations" *NBER Macroeconomics Annual 1988*, 111-148.

[17] 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.

[18] Taylor, John B. 1980. "Aggregate dynamics and staggered contracts." *Journal of Political Economy* **88**, 1-23.

Table 1. Employment Response to Productivity across Manufacturing Industries

	(1) OLS	(2) OLS	(3) OLS	(4) IV	(5) OLS	(6) OLS	(7) IV
constant	0.206	0.790	0.135	2.209	0.253	0.588	2.30
	(3.37)	(6.41)	(0.57)	(3.69)	(1.24)	(1.47)	(2.24)
$\ln DUR_i$	0.167		0.174		0.107		
	(4.39)		(4.04)		(2.56)		
$\ln(\overline{n/q})_i$		0.191	-0.032	0.917		0.134	0.712
		(3.06)	(-.31)	(3.23)		(1.16)	(2.12)
$\ln PStick_i$					-0.075	-0.081	-0.452
					(-0.49)	(-0.47)	(-1.60)
R^2	0.167	0.018	0.168	NA	0.088	0.017	NA
Obs.	98	458	98	98	72	72	72

Note: The dependent variable is SRR_i^h, the short-run response of hours to productivity shock estimated from industry VARs. DUR_i denotes the durability (the average service life) of the industry product; $(\overline{n/q})_i$ is the average inventory-sales ratio of the industry; $PStick_i$ represents the price-stickiness (average duration of the price) of the industry product. The numbers in parenthesis are t ratios. The R^2 of the first-stage regression in IV is 0.2.

Table 2. Cross-Sectional Correlations

SRR_i^h	SRR_i^h	SRR_i^n	$\ln DUR_i$	$\ln\overline{(n/q)}_i$
SRR_i^h	1.00			
SRR_i^n	0.42	1.00		
$\ln DUR_i$	0.41	0.20	1.00	
$\ln\overline{(n/q)}_i$	0.15	0.02	0.45	1.00

Note: The short-run response of hours (SRR_i^h) and inventory holdings (SRR_i^n) of industry i are from industry VARs. DUR_i denotes the durability (the average service life) of the industry product; $\overline{(n/q)}_i$ is the average inventory-sales ratio of the industry. Correlations regarding durability are based on 98 industries only. All other correlations are based on 458 4-digit industries.

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

A. Aggregate Response

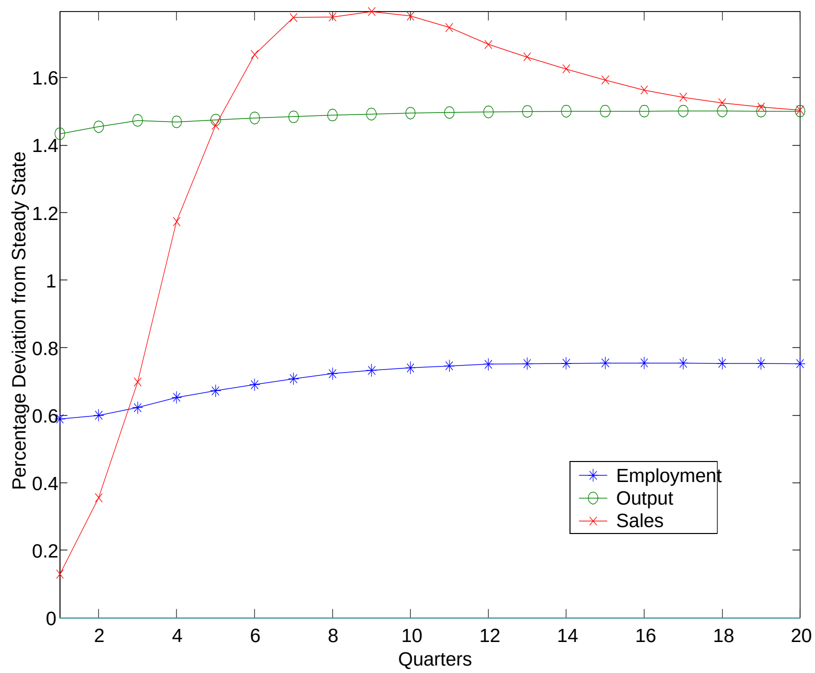

B. Individual Firms' Response

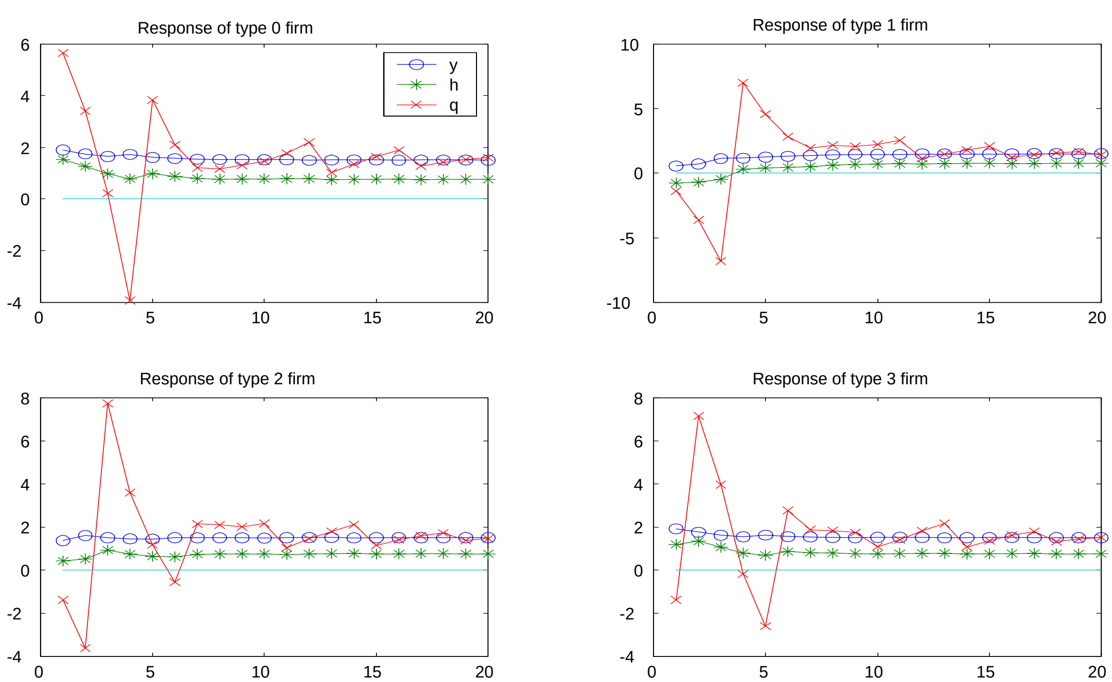

Figure 2. The Response to a Permanent Productivity Increase With Depreciation in Storage, $\delta = 0.9$.

A. Aggregate Response

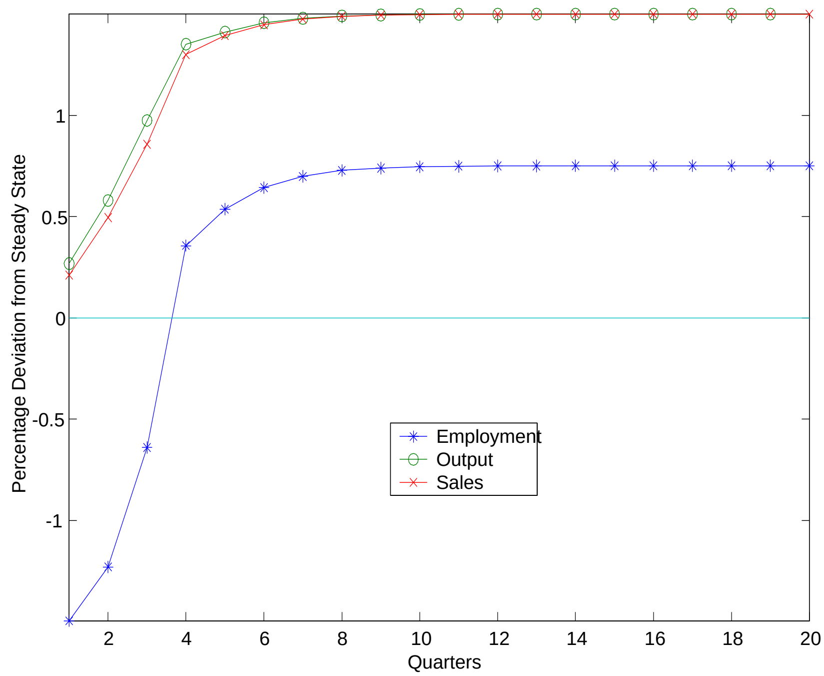

B. Individual Firms' Response

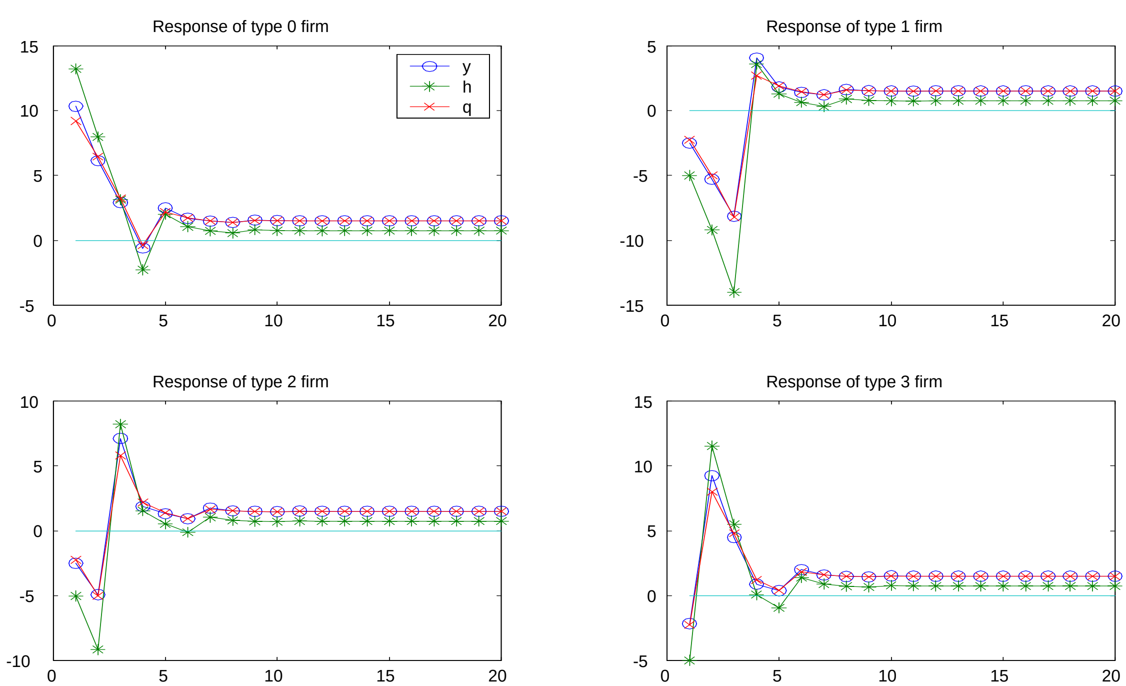

Figure 3 Employment Responses to Productivity Shocks

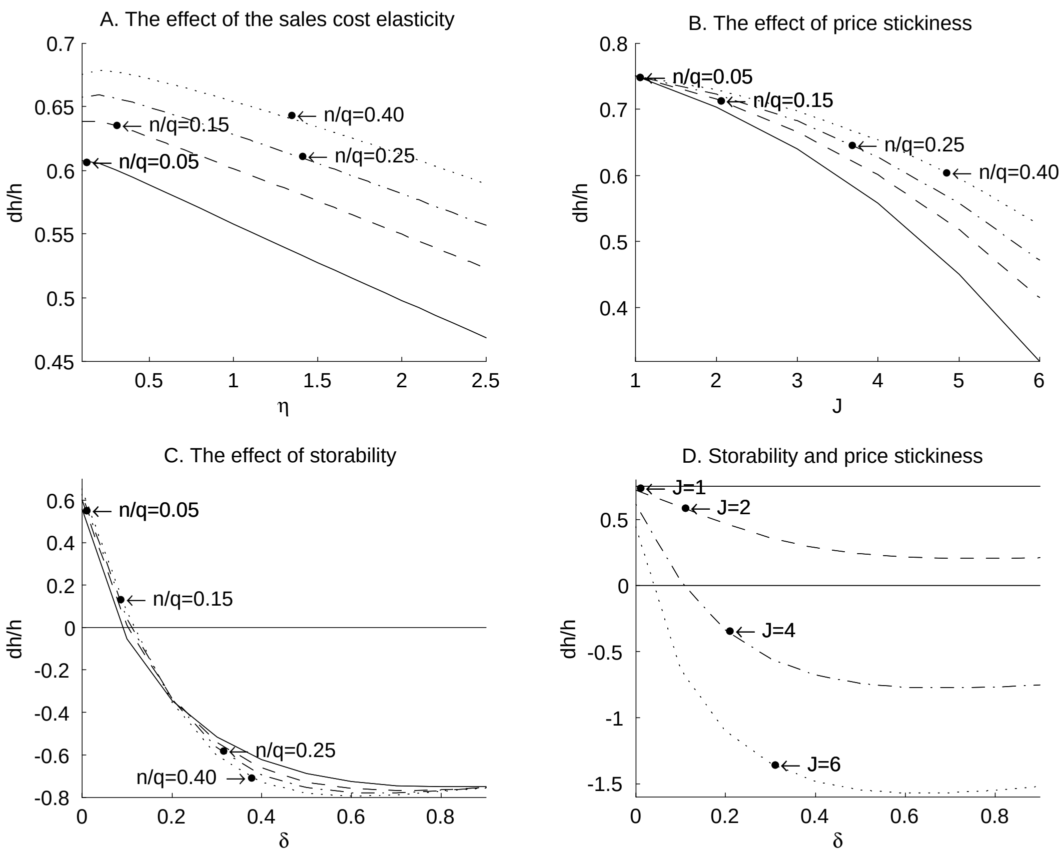

Note: We graph the cumulative one-year response of hours to a one-percent permanent increase in productivity. In Panel A the depreciation rate ($\delta = 0$) and price stickiness ($J = 4$) remain fixed. In Panel B the depreciation rate ($\delta = 0$) and sales-cost elasticity ($\eta = 1$) remain fixed. In Panel C price stickiness ($J = 4$) and sales-cost elasticity ($\eta = 1$) remain fixed. In Panel D, the steady state inventory-sales ratio ($n/q = 0.2$) and sales-cost elasticity ($\eta = 1$) remain fixed. In all panels we adjust the sales-cost scale factor κ to keep the inventory-sales ratio constant as we vary the variable on the x-axis.

Figure 4A. Response to Technology from VAR: Raw Food Manufacturing

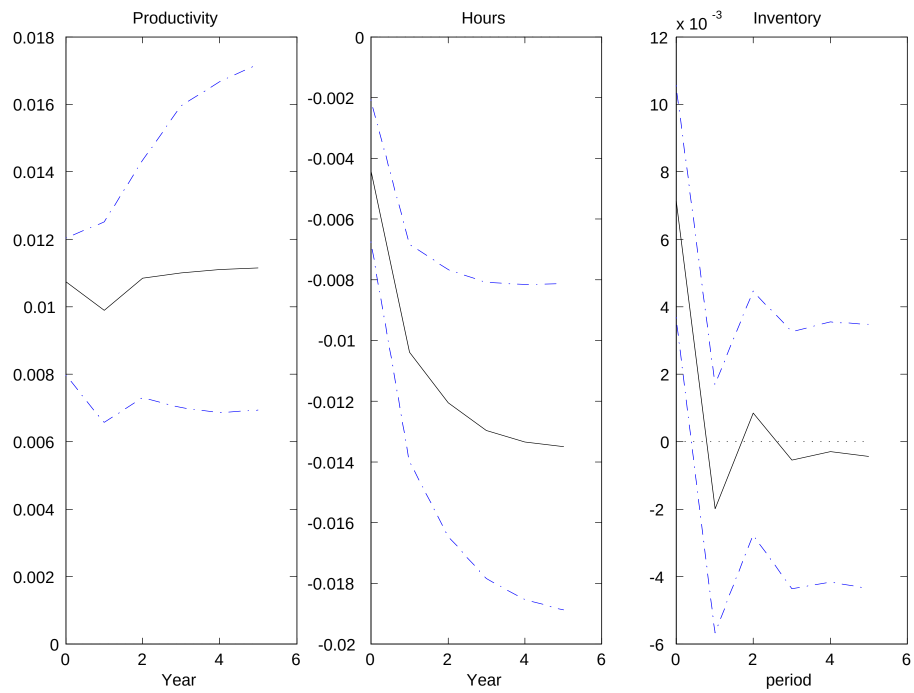

Figure 4B. Response to Technology from VAR: Durable Goods Manufacturing

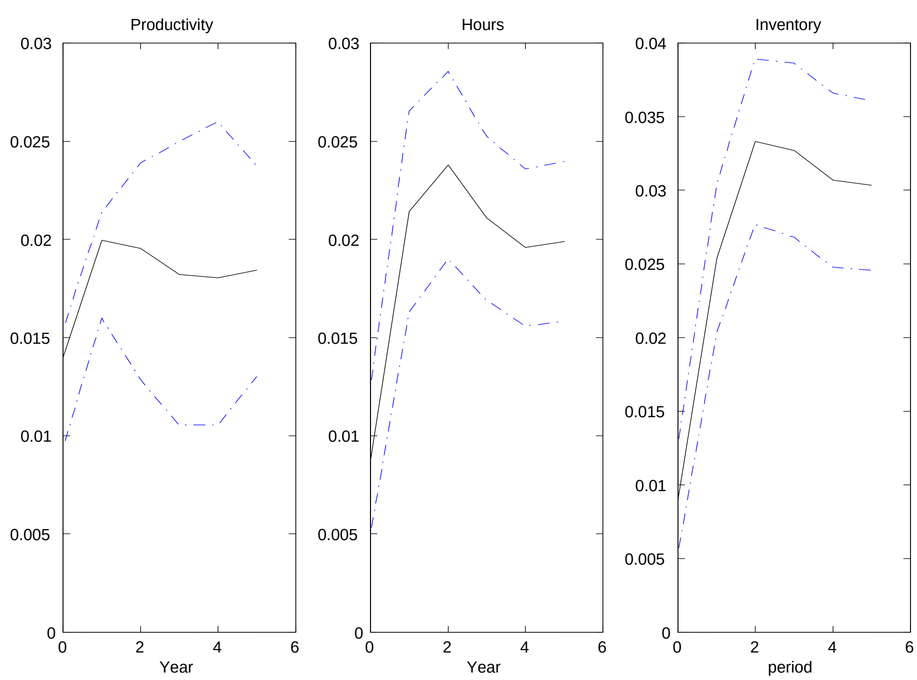

Figure 5A. Employment Response and Storability: Service Life

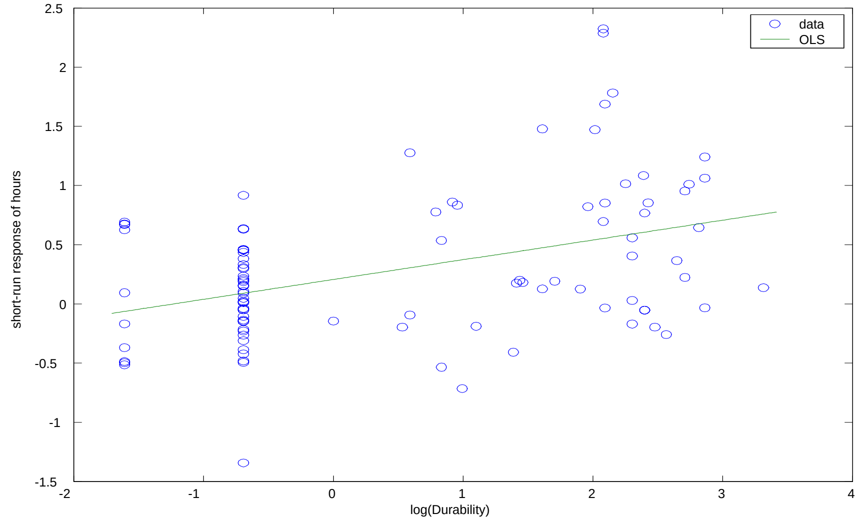

Figure 5B. Employment Response and Storability: Inventory-Sales Ratio

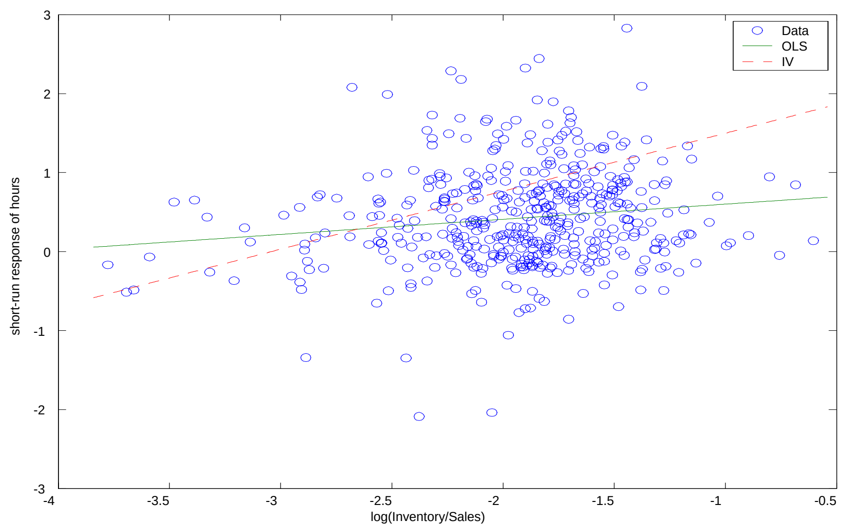

Note: The short-run response of hours represents the first-year response to a productivity shock that increases the industry TFP by one percent in the long run, estimated from industry VARs. Storability represents the average service life of the industry product in Figure 5A with 98 industry observations and the average inventory-sales ratio in Figure 5B with 458 industry observations. Figure 5A includes the OLS regression of Table 1, Column (1), and Figure 5B includes the OLS and IV regressions of Table 1, Columns (2) and (4).

Appendix: Robustness

1. Alternative Measure of Real Inventories for 4-Digit SIC Inventories. The choice of price index to deflate nominal inventories does not affect our results materially. In Table A.1 we display the relation between employment responses and measures of storability when we obtain real inventory stocks by deflating nominal inventories with the price index for materials inputs. The results are essentially the same as when we deflate nominal inventories with the price index for the value of shipments, Table 1.

Table A.1. Employment Response to Productivity across Manufacturing Industries
Alternative Measure of Real Inventories

	(1) OLS	(2) OLS	(3) OLS	(4) IV	(5) OLS	(6) OLS	(7) IV
constant	0.203	0.614	0.06	1.90	0.244	0.478	1.83
	(3.29)	(5.5)	(0.26)	(3.76)	(1.24)	(1.23)	(2.16)
$\ln DUR_i$	0.162		0.176		0.098		
	(4.2)		(4.04)		(2.33)		
$\ln(\overline{n/q})_i$		0.102	-0.065	0.784		0.098	0.554
		(1.78)	(-0.62)	(3.23)		(0.89)	(2.02)
$\ln PStick_i$					-0.072	-0.063	-0.368
					(-0.46)	(-0.36)	(-1.44)
R^2	0.15	0.007	0.16	NA	0.011	0.006	NA
Obs.	98	458	98	98	72	72	72

Note: Nominal inventories have been deflated with the materials price index.
All other variables are as defined in Table 1.

2. Finished Goods Inventories for 2-Digit SIC Industries. We mentioned earlier that the measure of inventories across 4-digit SIC industries is more comprehensive than our theory would imply. Inventories at the 4-digit level include not only finished goods inventories but also materials and goods in process. Inventory data by stage of production is available only for 2-digit SIC industries. Unfortunately, there are only 16 of these industries.[12] The number of 2-digit industries for which we have service-life measures is even smaller, with only ten of these available. While statistical tests on such a small sample are not particularly meaningful, it is nevertheless of interest to explore whether available observations on finished goods inventories at the 2-digit level are roughly consistent with the evidence we have on total inventories at the 4-digit level. Our results are summarized in Figure A.1 and Table A.2.

Figure A.1 is analogous to Figure 1 and displays how the short-run response of hours to a permanent one percent productivity increase relates to the storability of goods. Panel A displays the relation for the average service life of the industry goods.[13] If anything, service life and the employment response seem to be positively related, but with only ten observations no obvious pattern emerges from the graph.Panel B displays the relationship between the short run hours response and the average inventory sales-ratio for finished goods. The simple correlation between the employment response and the inventory-sales ratio is actually negative, Table A.2. This result is driven by one observation only, transportation equipment (SIC 37). Omitting this industry would clearly lead to a positive relationship between inventory-sales ratios and the employment response. Even without omitting this industry, the Spearman rank order correlation is not negative, 0.01. Finally, the correlation between the employment and inventory response to a permanent productivity shock is positive, Table A.2, that is even if prices are sticky, inventories are apparently used to smooth marginal cost over time.

[12] There are twenty 2-digit SIC industries but for two industry groups only aggregates are available. The first group consists of industries 24, 25 and 39 and the second group consists of industries 23, 27 and 31.

[13] We use our observations on service life for 4-digit SIC industries. Therefore for the 2-digit industries where we have any observations on service life we tend to have more than one observation.

Table A.2. Cross-Sectional Correlations for 2-Digit Industries

SRR_i^h	SRR_i^h	SRR_i^n	$\ln(\overline{n/q})_i$
SRR_i^h	1.00		
SRR_i^n	0.11	1.00	
$\ln(\overline{n/q})_i$	-0.19	0.08	1.00

Note: Variables are as defined in Table 2, but for 2-digit SIC industries and inventory-sales ratios are calculated for finished goods inventories. There are 16 industry observations.

Figure A.1A Employment Response and Storability for 2-Digit SIC Industries: Service Life

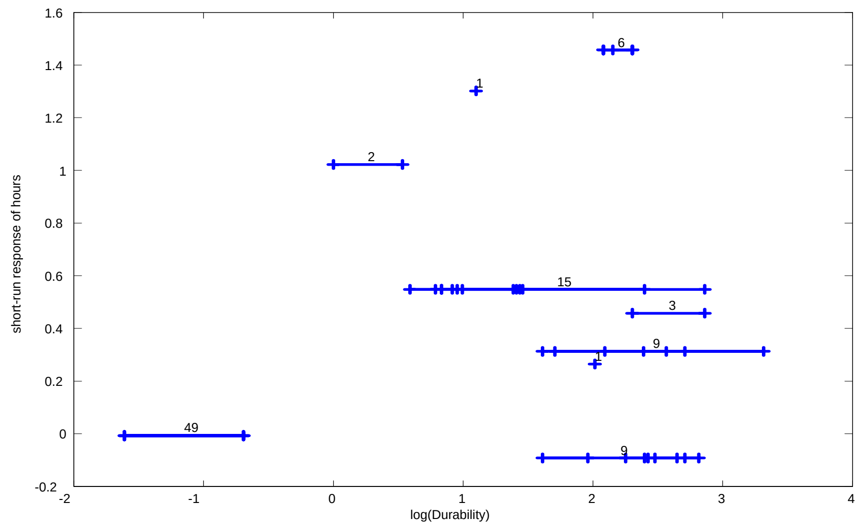

Figure A.1B Employment Response and Storability for 2-Digit SIC Industries: Inventory-Sales Ratio

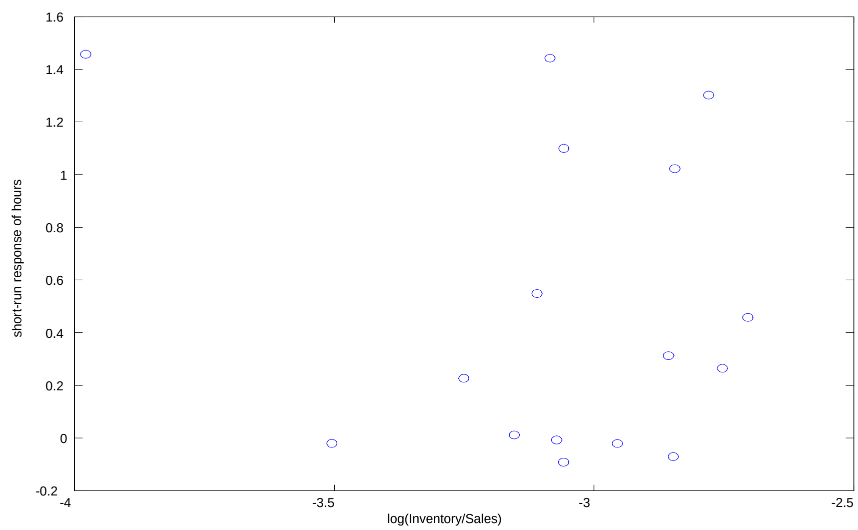

Note: In Panel A we graph the range of service life values for goods produced by the 4-digit industries contained in a particular 2-digit industry. Distinct service life observations are marked on the line, and the number on top of the line indicates the total number of service life observations for a 2-digit industry. The inventory-sales ratio in Panel B is for finished goods inventories only. We have service life observations for ten 2-digit SIC industries and we have average inventory-sales ratios for 16 industries. All other variables are as defined in Figure 1.