

Inventories, Markups, and Real Rigidities in Menu Cost Models*

Oleksiy Kryvtsov
Bank of Canada

Virgiliu Midrigan[†]
New York University

January 2009

Abstract

Real rigidities that limit the responsiveness of real marginal cost to output are a key ingredient of sticky price models necessary to account for the dynamics of output and inflation. We argue here, in the spirit of Bilal and Kahn (2000), that the behavior of marginal cost over the cycle is directly related to that of inventories, data on which is readily available. We study a menu cost economy in which firms hold inventories in order to avoid stockouts and to economize on fixed ordering costs. We find that, for low rates of depreciation similar to those in the data, inventories are highly sensitive to changes in the cost of holding and acquiring them over the cycle. This implies that the model requires an elasticity of real marginal cost to output approximately equal to the inverse of the elasticity of intertemporal substitution in order to account for the countercyclical inventory-to-sales ratio in the data. Stronger real rigidities lower the cost of acquiring and holding inventories during booms and counterfactually predict a procyclical inventory-to-sales ratio.

JEL classifications: E31, F12.

Keywords: Inventory, menu costs, real rigidities.

*The views expressed here are those of the authors and do not reflect those of the Bank of Canada. We thank George Alessandria, Mike Dotsey, Mark Getler, Boyan Jovanovic, James Kahn and Huntley Schaller for useful discussions, as well as seminar participants at the NBER-SITE Price Dynamics Conference at Stanford, Duke Macroeconomics Jamboree, Université du Québec à Montréal, Université de Montréal, Université Laval, University of Western Ontario, Chicago GSB, Federal Reserve Board, the Federal Reserve Banks of Atlanta, Philadelphia, Richmond and St. Louis, as well as at the Bank of Canada.

[†]Corresponding author: virgiliu.midrigan@nyu.edu, 19 W 4th St, 6th Floor, New York, NY 10012.

1. Introduction

Real rigidities, that is, factors that dampen the responsiveness of firms' desired prices to aggregate shocks, are key ingredients in New Keynesian sticky price models necessary to account for the inertia of aggregate prices, inflation and output in the data with the fairly rapid rate at which individual price setters update their nominal prices (once every two to three quarters)¹.

Two classes of models with real rigidities have received attention in recent work². The first class is characterized by a set of assumptions on preferences (non-constant demand elasticities) or technology (decreasing returns to scale) that make it costly for firms to allow their prices to deviate too far away from those of their competitors. These assumptions generate a complementarity in firms' pricing so that firms that reset their nominal prices find it optimal to not fully respond to an aggregate monetary shock³. Although measuring price elasticities or returns to scale in the production function is the subject of ongoing research, a number of recent papers have argued that simple versions of models that feature these types of pricing complementarities are difficult to reconcile with the observed dispersion of relative prices, even in very narrowly defined product groups within outlets⁴.

In this paper we focus on a second class of models with real rigidities in which inertia in aggregate prices and output is driven by factors that reduce the responsiveness of real marginal cost at the aggregate level to business cycle fluctuations. In this second class of models researchers make assumptions on preferences, the degree to which factor utilization can vary, or frictions in the labor and intermediate inputs market that make it optimal for firms to incompletely adjust their prices in response to a monetary shock⁵. Effectively, this second class of models is a set of assumptions about aggregate quantities, that is, about the firms' collective ability to hire additional labor during booms (labor supply elasticity), hoard labor during recessions, purchase additional intermediate inputs, or vary the degree to which they utilize the existing stock of capital and labor. Moreover, even when real rigidities take the form of sticky wages or intermediate goods' prices, an important assumption made is that

¹Bils and Klenow (2004), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008).

²Ball and Romer (1990) have originally suggested this classification.

³Several well-known contributions include Kimball (1995), Eichenbaum and Fisher (2004), Altig, Christiano, Eichenbaum and Linde (2004).

⁴Klenow and Willis (2006), Dotsey and King (2005), Burstein and Hellwig (2007).

⁵Erceg, Henderson and Levin (2000), Dotsey and King (2002), Christiano, Eichenbaum and Evans (2002), Nakamura and Steinsson (2007).

these prices are allocative and quantities are demand-determined. Thus, the assumption that factor inputs can be relatively flexibly varied over the cycle is a key feature in this second class of models. Indeed, Dotsey and King (2005) use the term “real flexibilities” to characterize this class of models.

Inferring the elasticity of real marginal cost to output, a key measure of the strength of real rigidities in this second class of models, is difficult in practice. In particular, one must take a stand on the technology with which goods are produced, the relative importance of factor adjustment costs, the degree to which prices in long-term relationships are allocative⁶, the cost of varying the work-week of capital and labor, as well as the degree of frictions in the labor and intermediate goods markets⁷.

In this paper we use the insights of Bils and Kahn (2000), who relate the behavior of inventories to that of costs, in order to study the implications of models of this second class of real rigidities for the behavior of inventories. If the marginal cost of acquiring and holding inventories is indeed lower in times of a monetary expansion, as in this second class of models of real rigidities, we should see this lower cost reflected not only in a slow response of prices to the monetary expansion, but also in an increase in firms’ inventory holdings. Moreover, since the firm’s price reflects its shadow valuation of inventories, an increase in the stock of inventories is necessary for the firm to find it optimal keep its price low. Thus, for storable goods, real rigidities that dampen the elasticity of marginal cost to output must operate through increases in the stock of inventories. If the firm is unable to purchase more inventories, either because of quantity restrictions by suppliers, or because of other costs of adjusting the stock of inventories, the relatively lower factor prices do not translate into a lower shadow valuation of inventories, and the firm finds it optimal to change its price fully in response to the monetary shock. Therefore we argue that a key empirical test of this class of models is their ability to account for the observed behavior of inventories.

We start by documenting that monetary expansions and business cycles in general are associated with increases in the stock of inventories and declines in the inventory-sales ratio: the elasticity of inventory-sales ratios to sales is equal to -0.8. This is reminiscent of the findings of Ramey and West (1999), and Bils and Khan (2000) that, at business cycle

⁶Stigler and Kindahl (1970), Barro (1977), Hall (2006) are several important references.

⁷See also Solon, Barsky and Parker (1994) who document compositional biases in the measurement of real wages.

frequencies, inventories are procyclical, but less so than sales. Moreover, we use detailed data from the NBER productivity database and the Bureau of Economic Analysis to show that this decline is not simply a compositional artifact resulting from a tendency for industries that have on average lower inventory-to-sales ratios to expand more during booms.

We then formulate and calibrate a model in which nominal prices change infrequently due to fixed costs of price adjustment. Firms hold inventories in our model because of two frictions on the technology of ordering goods. In particular, we assume that ordering a new batch of inventories entails payment of an ordering cost that is independent of the shipment's size. Furthermore, the firm is assumed to order before the realization of a taste shock that determines the consumers' demand for its goods. As a result of these two assumptions, the firm holds excess inventories in order to minimize both the probability of a stockout and the ordering costs. We study these two particular inventory-holding motives because of the attention that they have received in recent work, and we use micro-economic evidence to inform their relative strength.⁸

We use our model economy to study how the response of inventories to monetary policy shocks depends on the elasticity of real marginal cost to output. We vary the responsiveness of real marginal cost to output in the model by introducing a wedge in the optimal labor-leisure choice that allows us to mimic a range of responses for the real wage without modeling the frictions that account for these responses explicitly. Our main finding is that in the model inventories are highly sensitive to changes in the costs of acquiring (the price of intermediate goods or labor) and holding (the interest rate) inventories. As a result, even modest fluctuations in the cost of acquiring and holding inventories generate substantial fluctuations in inventories over the cycle. Accounting for the -0.8 elasticity of the inventory-sales ratio to sales in the data thus requires that the cost of acquiring and holding inventories is roughly constant over the cycle. This implies, in particular, that the elasticity of real marginal cost to consumption must be approximately equal to the inverse of the intertemporal elasticity of substitution (IES). When the elasticity of real marginal cost to consumption is lower than the inverse IES, the combined cost of acquiring and holding inventories decreases during a monetary expansion, and the firm finds it optimal to raise its inventory stock by a large

⁸See Caplin (1985), Ramey (1991), Khan and Thomas (2007) who emphasize non-convexities in the cost of acquiring or producing inventories. Kahn (1992), Bils and Kahn (2000), Khan and Thomas (2007b), Wen (2008) study economies with a stockout-avoidance motive.

amount. This in turn leads, counterfactually, to an increase in aggregate inventory-to-sales ratio. We thus find, in our setup, that it is difficult to reconcile low elasticities of real marginal cost to output with the behavior of inventories in the data, unless one also assumes a high elasticity of intertemporal substitution. As a comparison, elasticities used in earlier work on real rigidities range from 0.10-0.15 as argued for by Woodford (2003), 0.33 in Dotsey and King (2005), and 0.34 in Smets and Wouters (2007) to 2.25 – the baseline parameterization used by Chari, Kehoe, McGrattan (2002).

Our work is related to a number of recent papers studying the behavior of inventories, costs and markups over the business cycle. Our starting point is the observation of Bils and Kahn (2000) that inventories are closely linked to markups and marginal costs and thus may provide important information about the cyclicity of the latter. Khan and Thomas (2007) and Wen (2008) study real business cycle models in which inventories arise respectively due to fixed ordering costs, and a stockout-avoidance motive. These papers find that the model is capable of accounting for the countercyclicality of the inventory-sales ratio observed in the data. Our conjecture is that they do so because investment in capital in times of expansions in these models drives up the cost of purchasing (through higher elasticity of real marginal cost to output) and holding (through higher interest rates) inventories. Two closely related papers, Khan and Thomas (2008) and Jung and Yun (2005), also study sticky price economies with inventories and ask whether the model can account for the behavior of inventories in the data: Khan and Thomas (2008) argue the presence of capital is a necessary ingredient to match the countercyclicality of the inventory-to-sales ratio, while Jung and Yun (2005) emphasize the role of high depreciation rates and adjustment costs. Finally, Chang, Hornstein and Sarte (2007) study the responses to a productivity shock in a sticky price model with inventories. They find that whether an industry expands or contracts employment depends on the carrying costs of inventories: higher carrying costs prevent firms from responding to the productivity shock by investing in inventories and as a result cut employment given that prices are sticky and quantities are demand-determined.

This paper proceeds as follows. Section 2 documents a number of facts about inventories and computes the response of inventory-to-sales ratio to measures of exogenous monetary policy shocks. We also show that the decline in inventory-to-sales ratio is pervasive across all industries and stages of fabrication, and not simply a compositional artifact. Section 3 presents the model and discusses optimal decision rules in this environment. Section 4 com-

putes impulse responses to monetary policy shocks under several assumptions regarding the elasticity of real marginal cost with respect to output (i.e, stickiness of nominal wages) and asks what elasticity is consistent with the decline in the inventory-to-sales ratio observed in the data. Section 5 performs additional robustness checks and uses data on inventories, sales and prices from a Spanish supermarket⁹, in order to evaluate the microeconomic implications of variations of the model capable of reducing the sensitivity of inventories to costs. Section 6 concludes.

2. Data

In this section we document several salient facts regarding the cyclical behavior of inventories. We employ two datasets: the (annual) NBER Manufacturing Productivity Database with data on output, sales, inventories and price deflators in 459 4-digit manufacturing industries from 1957 to 1996; and the Bureau of Economic Analysis (NIPA) data on monthly final sales, inventories, and inventory-to-sales ratios for Manufacturing and Trade sectors from January 1967 to December 1997¹⁰. These data sets complement each other as they allow us to study inventory behaviour for two business cycle frequencies - monthly and annual; and for different levels of aggregation - industry-level and aggregate. Although we focus on a subset of the US economy, our data accounts for most of the U.S. inventory stock. Manufacturing and Trade inventories (value added) comprise 85% (74%) of the total private nonfarm inventory stock (value added); the remaining industries are mining, utilities, and construction.

Sales are defined as real final sales to domestic purchasers in NIPA data and as the real value of industry shipments in case of the NBER data. Output is the sum of final sales and the change in the end-of-period inventory stock. Inventory-to-sales ratio is defined as the ratio of the end-of-period inventory stock to final sales in that period. Shipments deflators are used for industry-level data and CPI (less food and energy) is used for aggregate data to deflate nominal variables. All data are HP filtered.¹¹ Output, sales and inventory-sale ratios

⁹Aguirregabiria (1999).

¹⁰The Bureau of Economic Analysis uses an inventory valuation adjustment to revalue inventory holdings (reported by various companies using potentially different accounting methods) to replacement cost. These adjustments are based on surveys that report the accounting valuation used in an industry and from information on how long goods are held in inventories. See Ribarsky (2004).

¹¹There is a great deal of heterogeneity in industry-specific trends of sales and inventory-to-sales ratios. Bils and Kahn (2000) argue that scale effects are small for inventory-to-sales ratio dynamics. Our work below shows that heterogeneity across inventory-to-sales ratios is unlikely to be an important factor for the dynamics of the aggregate inventory-to-sales ratio.

are defined in % deviations from respective HP trends. Inventory investment is defined in % points-of-output-fraction deviations from its HP trend. Unless otherwise noted, inventory-to-sales ratios, output, sales, etc. are in real units.¹²

We divide the data analysis into four parts. We first document the countercyclicality and other facts about the dynamics of the aggregate inventory-to-sales ratio over the business cycle. Second, we show that the countercyclicality of the aggregate inventory-to-sales ratio does not stem from compositional shifts at the disaggregate level. We then report the output and sales elasticities of inventory-sales ratios at the aggregate and industry levels. Finally, we show that inventory-sales ratios decline in response to (identified) exogenous monetary expansions, and that the facts documented for business cycles in general characterize those episodes as well.

A. Inventories and sales over the business cycle

This section describes the behavior of inventories over the business cycle¹³. Table 1 reports several unconditional moments for real inventory-to-sales ratios in the data. On average firms carry inventory stocks of about 1.4 months of sales according to the monthly data and 0.31 years of sales in the annual data (that is, 3.7 months of sales). Standard deviations of real inventory-to-sales ratios, 2.2% in the monthly and 3.9% in the annual data, are comparable to those of the aggregate output and sales. The most prominent feature of the aggregate inventory-to-sales ratios is their countercyclicality. The correlation of real inventory-to-sales ratios with output in the monthly data is -0.83 for Manufacturing and Trade, and -0.49 for Retail. In the annual data that correlation is -0.52. Correlations with sales are typically even more negative. Figure 1 reinforces this finding by plotting the aggregate (detrended) inventory-to-sales ratio and output in the Manufacturing and Trade sector. The negative correlation is evident.

B. Compositional bias

There is a considerable amount of heterogeneity in the level of stock relative to sales among firms. For example, in 1996 25th, 50th and 75th-percentile levels of inventory-to-sales ratios across 4-digit NBER industries are 0.23, 0.31 and 0.42 respectively. One concern

¹²Detailed descriptions of the data and definitions are available upon request.

¹³Ramey and West (1999) and Bils and Kahn (2000) are two earlier papers that also document some of these facts.

is therefore that the behavior of inventory-to-sales ratios may reflect a compositional shift. Given that the aggregate inventory-to-sales ratio is equal to the average of industry inventory-to-sales ratios (weighted by each industry’s sales), it may be the case that the drop of the aggregate ratio is simply evidence that low-inventory-to-sales industries expand more during booms.¹⁴

To assess this hypothesis, we divide the annual data into three equally sized bins according to the level of inventory-to-sales ratios: low, medium, and high. Table 2 reports, for each bin, the average inventory-to-sales ratio and the bin’s sales share. The lower bin accounts for half the sales in our sample and industries in this bin have an average inventory-to-sales ratio of 0.18. The medium and high bins have a mean inventory-to-sales ratios of 0.31 and 0.55, respectively, and account for a quarter of total sales each.

For compositional effects to be important, there must be large differences in how inventory-to-sales ratios fluctuate over the business cycle. Table 2 demonstrates that these differences, if any, are small. Specifically, correlations of bin-specific inventory-to-sales ratios with real GDP are very similar, -0.41, -0.33 and -0.47, respectively.

An alternative way of gauging the extent of compositional effects is to compare a common industry-level time effect, or a fixed-weight average inventory-to-sales ratio, to the aggregate ratio. Specifically we run the following panel regression:

$$\ln IS_{it} = \alpha \mathbf{D}_t + \beta \mathbf{D}_i + \varepsilon_{it} ,$$

where $\ln IS_{it}$ is the log inventory-to-sales ratio for industry i , \mathbf{D}_t is the vector of year dummies, \mathbf{D}_i are 4-digit industry dummies, and ε_{it} is the residual representing industry-level disturbances. Coefficients for year dummies, α_{year} , represent the fixed-weight average inventory-to-sales ratio. If compositional effects are important, the time series for the average and aggregate inventory-to-sales ratios should differ substantially. Figure 2 shows that this is not the case. Indeed the two series move together over the cycle, except for the 1994-1997 period when the aggregate ratio was somewhat smaller than the average (as in Figure 1, the two series are HP-filtered). We conclude that compositional shifts across sectors with high

¹⁴Solon, Barsky, Parker (1994) find such a compositional bias for aggregate real wages that are acyclical in the aggregate time series because the aggregate statistics are constructed in a way that gives more weight to low-skill workers during expansions than during recessions.

and low inventory-to-sales ratios do not contribute importantly to countercyclicality of the aggregate inventory-to-sales ratio.

C. Elasticities of inventory-to-sales ratios

To gauge the sensitivity of inventory-to-sales ratio movements to fluctuations in output and sales, we estimate their respective elasticities at the aggregate and industry levels. Table 3 reports aggregate elasticities.

In the monthly data, the output elasticity of inventory-to-sales ratio in Manufacturing and Trade is -0.77, and sales elasticity is -0.86. For Retail, elasticities are somewhat smaller (in absolute value), -0.49 and -0.77 respectively. In the NBER annual data, output and sales elasticities are -0.42 and -0.60 respectively.¹⁵

Elasticities are similar when we focus separately on inventories at different stages of disaggregation (raw materials, work-in-progress, and finished goods). Around 77% of all Manufacturing and Trade inventories are finished goods inventories, and 23% are raw materials and work-in-progress inventories. We compute elasticities of NIPA inventory-to-sales ratios by the stage of fabrication with respect to total Manufacturing sales (each time series is log HP-detrended). Respective sales elasticities are -1.01 for raw materials, -0.87 for work-in-progress, and -1.16 for finished goods inventories (-1.01 for total Manufacturing).

Iachoviello, Schiantarelli and Schuh (2007), hereafter ISS, argue that it is important to distinguish between output inventories (which they define as retail inventories), and the rest - which they refer to as input inventories. They find that the ratio of output inventories to consumption purchases is acyclical (the quarterly correlation with aggregate output is 0.10), whereas the ratio of input inventories to output is countercyclical (correlation with output is -0.89). Moreover, they report that the output inventory-to-consumption ratio is more volatile: between 2.5 and 4.6 times more volatile than input inventory-to-output ratio.

We show below that our results are not inconsistent with theirs. To see this, we first compute the elasticity of retail inventories with respect to real consumption expenditures by regressing real retail inventory-to-consumption ratios on respective real personal consumption of durables, non-durables and durables plus nondurables¹⁶ in the NIPA data for 1990:01-

¹⁵Measured elasticities are not sensitive to our detrending method (HP-filter). For example, if the trend is linear (in logs), output (sales) elasticities in NBER data change from -0.42 (-0.60) to -0.54 (-0.69).

¹⁶We argue that consumption of durable plus nondurable goods is a reasonable proxy of ISS's consumption.

2008:04. The respective elasticities are -1.02, -0.38, and -0.86 (or, at the quarterly frequency, -0.88, -0.42, and -0.62). ISS report elasticities of input inventory-to-consumption with respect to aggregate output, so we regress our three NIPA ratios on real GDP: respective elasticities are -0.50, 0.63, and 0.32. Hence, the apparent discrepancy in results is accounted for by the fact that we focus on how inventory-sales ratios within a given sector covary with that sector's sales or output, whereas ISS focus on how similar inventory-to-sales ratios covary with *aggregate* output.

Finally, at the industry level, across-time elasticities turn out to be of the same magnitude. Table 4 shows that across-time sales elasticity of inventory-to-sales ratio at the industry level is around -0.6. When we compare the across-time sales elasticity of inventory-to-sales ratio to its cross-section counterpart, we see that the cross-section elasticity is much smaller than the one across time: -0.14. These estimates of industry-level across-time and cross-sectional elasticities are not sensitive to controlling for various measures of industry inflation, marginal cost, markup, and firm concentration. Together, these correlations suggest that the negative correlation between inventory-sales ratios and sales is a cyclical phenomenon and not accounted for by technological factors (e.g. scale economies that allow larger firms to hold fewer inventories). This result is reminiscent of that of Bils and Kahn (2000), who find that for six large manufacturing industries, the inventory-to-sales ratio is insensitive to industry size.

D. Evidence from identified monetary shocks

Since the focus of this paper is on the real effects of monetary shocks, we extend our data analysis to document the behavior of inventory-to-sales ratios in the wake of monetary expansions and contractions. A necessary step here is to identify these monetary disturbances. We use the monthly NIPA data which has the advantage of higher-frequency observations.

We employ two available measures of monetary shocks: due to Romer and Romer (2004) and Christiano, Eichenbaum and Evans (1996). Both measures represent innovations to the federal fund's rate (in RR's case, intended federal funds rate). The Romer-Romer (RR) measure is based on narrative records of FOMC meetings and Federal Reserve's internal forecasts. The Christiano-Eichenbaum-Evans (CEE) measure corresponds to the innovation to the federal funds rate in their VAR.

Responses of output, inventory investment and inventory-to-sales ratios for each sector

are obtained by estimating the following OLS regression:

$$y_t = \alpha_0 + \alpha_1 t + \sum_{s=0}^{36} \beta_s ffs_{t-s} + \gamma y_{t-37} + \varepsilon_t , \quad (1)$$

where y_t is the dependent variable, α_0 , $\alpha_1 t$ are a full set of monthly dummies, ffs_t is a measure of monetary shocks, and ε_t is the zero-mean normally distributed error term, which is assumed to be AR(2):

$$\varepsilon_t = \rho_1 \varepsilon_{t-1} + \rho_2 \varepsilon_{t-2} .$$

Estimation of (1) yields a number of facts about impulse responses after monetary shocks.^{17,18}

Figure 3 reports impulse responses to CEE shocks for output, sales, inventory investment and inventory-to-sales ratio in Manufacturing and Trade. The solid lines in the figure denote point estimates of the different response functions, and the dashed lines report respective 95% confidence intervals. Output responses are positive for about half a year after the shock, reflecting the fact that monetary shocks affect the economy with a lag. Output responses are negative between around 6 and 30 months, with the trough at around 24 months after the shock. Sales responses are very similar to output responses. The response of inventory investment is small and statistically almost indistinguishable from zero. Finally, the response of inventory-to-sales ratios is negative for the first 5 months after the shock and then positive up until 2.5 years after the shock. For Retail, inventory-to-sales ratios are somewhat quicker to respond and less persistent – they are around zero in the first few months after the shock, then positive coming back to zero after 18 months after the shock (versus around 30 months for the Manufacturing and Trade sector)¹⁹.

We conclude that inventory-to-sales ratio is countercyclical even conditional on monetary disturbances, thus corroborating the results of Jung and Yung (2005). To quantify the extent of the ratio's comovement with output and sales, we estimate its elasticities *condi-*

¹⁷The results are provided for estimates based on CEE measure of monetary shocks. The results based on the RR measure are very similar.

¹⁸Bils, Klenow, Kryvtsov (2003) employ a similar method for estimating impulse responses of consumption and prices across BEA consumption sectors to CEE monetary shocks.

¹⁹As a robustness check, we have estimated our own measure of monetary policy shocks and have obtained similar results.

tional on monetary shock. Specifically, we regress the fitted inventory-to-sales ratio on fitted output (sales) from regression (1). For Manufacturing and Trade, conditional output and sales elasticities are only slightly lower, -0.64 and -0.71, than the unconditional elasticities. For the Retail sector, elasticities conditional on monetary shocks are slightly larger than the unconditional elasticities: -0.49 for output elasticity and -0.57 for sales elasticity.

3. Model

The economy consists of a continuum of final goods firms, indexed by i , that each sells a differentiated good produced using intermediate goods as inputs; a continuum of competitive intermediate goods firms; and a representative household that derives utility from the consumption of final goods, trades a complete set of state-contingent securities and supplies labor to the intermediate goods sector. Given that the intermediate goods sector is competitive, our use of the terms “retail” and “wholesale” is a matter of notational convention, as one can reinterpret the problem we present below as that of a firm which produces and retails the good simultaneously. In this sense, our use of inventory data from the Manufacturing and Trade Sector, and use of both “input” and “output” inventories is not inconsistent with the model. In particular, the fact that inventory accumulation reflects the behavior of cost is clearly not specific to a particular sector.

Finally, we assume that consumers, unlike firms, cannot store the good. Their problem of choosing how much to buy of each of the differentiated goods is thus static.

Consumers

Consumers have preferences over a continuum of consumption goods and leisure and maximize

$$\begin{aligned} & \max_{c_t(i), n_t, \mathbf{b}_{t+1}} E_0 \sum_{t=0}^{\infty} \beta^t U(c_t, n_t) \\ \text{s.t. } & \int_0^1 p_t(i) c_t(i) di + \mathbf{q}_t \cdot \mathbf{b}_{t+1} \leq W_t n_t + b^t + \Pi_t , \\ & c_t = \left(\int_0^1 v_t(i)^{\frac{1}{\theta}} c_t(i)^{\frac{\theta-1}{\theta}} di \right)^{\frac{\theta}{\theta-1}} , \\ & c_t(i) \leq z_t(i) . \end{aligned}$$

Here \mathbf{b}_{t+1} is a vector of state-contingent Arrow-Debreu securities that the consumer buys and \mathbf{q}_t is a vector of security prices, b^t is the quantity of the respective state's bonds the agent has purchased at $t - 1$, Π_t is firm profits, $W_t n_t$ is labor income, $c_t(i)$ is consumption of the variety i , and $p_t(i)$ is that variety's price. c_t is the CES aggregator over different varieties. In this economy the consumer will occasionally be turned down by stores with little inventory available for sales. We let $z_t(i)$ be the stock of inventories that firm i has available for sale. The consumer cannot buy more than $z_t(i)$ units.

It follows that the consumer's demand for each of the final goods is:

$$c_t(i) = v_t(i) \left(\frac{p_t(i)}{P_t} \right)^{-\theta} c_t \text{ if } v_t(i) \left(\frac{p_t(i)}{P_t} \right)^{-\theta} c_t < z_t(i) ,$$

$$c_t(i) = z_t(i) \text{ otherwise,}$$

where $P_t = \left[\int_0^1 v_t(i) [p_t(i) + \tilde{\mu}_t(i)]^{1-\theta} di \right]^{\frac{1}{1-\theta}}$ and $\tilde{\mu}_t(i)$ is the product of the multipliers on the consumer's budget constraint and that on the constraint that purchases do not exceed the stock of inventories available for sale. Since $\tilde{\mu}_t(i) > 0$ for a positive fraction of goods in equilibrium, P_t is higher than $\hat{P}_t = \left[\int_0^1 v_t(i) p_t(i)^{1-\theta} di \right]^{\frac{1}{1-\theta}}$, the usual expression for the aggregate price index in economies without stockouts. As a result, demand for unconstrained goods, $c_t(i)$, is greater than that in an economy without stockouts: $c_t(i) > v_t(i) \left(\frac{p_t(i)}{\hat{P}_t} \right)^{-\theta} c_t$. The consumer thus directs expenditures towards the unconstrained goods, whenever some goods are in limited supply. Also notice that here P_t is not equal to the aggregate price index, usually defined as $\int_0^1 p_t(i) \frac{c_t(i)}{C_t} di$.

Finally, the Euler equation and intra-temporal consumption leisure choice take the usual form:

$$-\frac{U_{n,t}}{U_{c,t}} = \frac{W_t}{P_t} ,$$

and

$$q(s^{t+1}|s^t) = \beta \pi(s^{t+1}|s^t) \frac{U_c(s^{t+1}) P(s^{t+1})}{U_c(s^t) P(s^t)} .$$

Retail Firms

Retailers buy a storable good from a perfectly competitive intermediate-goods industry

at a (nominal) price ω_t . One unit of the intermediate good produces $a_t(i)$ units of the final good. The productivity of the firm, $a_t(i)$, is idiosyncratic and follows a Markov process to be described below. The retailer sells the goods in the final goods market and faces a demand curve for its product given by:

$$c_t(i) = v_t(i) \left(\frac{p_t(i)}{P_t} \right)^{-\theta} c_t ,$$

as well as the constraint that it does not sell more than the stock of inventories, $z_t(i)$, it has available.

We next discuss the frictions we impose on the retailer's ability to change prices and adjust inventories. Retailers face two fixed adjustment costs: of adjusting prices, κ_p , and of ordering new goods from the intermediate firms, κ_s . A third restriction we impose is that price and inventory decisions are made prior to the realization of that period's demand $v(i)$. We assume for simplicity that $v(i)$ is serially uncorrelated. Finally, we assume an irreversibility: once produced, final goods cannot be returned.

To economize on the ordering costs and to insure against the possibility of high demand $v(i)$, firms carry non-zero inventories of the final good from one period to another. In particular, let $s_t(i)$ be the retailer's beginning-of-period inventory level, and $z_t(i)$ be the stock of inventory available after new goods were ordered (if any). The firm's expected sales, given the price it chooses, p , and the inventory stock, z , is equal to

$$\mathcal{R}(p, z) = \int_0^{\infty} \min \left(v \left(\frac{p}{P} \right)^{-\theta} c, z \right) d\Phi(v) ,$$

where $\Phi(v)$ is the distribution of preference shocks, assumed to be log-normal with variance σ^2 . Let $v^* = \frac{z}{\left(\frac{p}{P}\right)^{-\theta} c}$ denote the lowest level of the taste parameter for which the firm stocks out. Then we can write:

$$\mathcal{R}(p, z) = \left(\frac{p}{P} \right)^{-\theta} c \exp \left(\frac{\sigma^2}{2} \right) \Phi \left(\log(v^*) - \sigma^2 \right) + z \left(1 - \Phi \left(\log(v^*) \right) \right) .$$

We next write the firm's dynamic program. Let $V^{a,a}(p_{-1}, s_{-1}, a)$ be the firm's value of adjusting prices and inventories, given the beginning-of-period stock of inventories s_{-1} , and

inherited price p_{-1} , as well as productivity a . $V^{n,a}$ value of not adjusting prices and adjusting inventories, and $V^{a,n}, V^{n,n}$ similarly defined. The firm's continuation value is the envelope of these four options: $V = \max(V^{a,a}, V^{a,n}, V^{n,a}, V^{n,n})$. Then

$$\begin{aligned}
V^{a,a}(p_{-1}, s_{-1}, a) &= \max_{p, z \geq s_{-1}} \left[\frac{U_c}{P} [p\mathcal{R}(p, s_{-1} + i) - a\omega(z - s_{-1}) - (\kappa_p + \kappa_s)W] + \beta EV(p, s'_{-1}, a') \right], \\
V^{n,a}(p_{-1}, s_{-1}, a) &= \max_{z \geq s_{-1}} \left[\frac{U_c}{P} [p_{-1}\mathcal{R}(p_{-1}, s_{-1} + i) - a\omega(z - s_{-1}) - \kappa_s W] + \beta EV(p, s'_{-1}, a') \right], \\
V^{a,n}(p_{-1}, s_{-1}, a) &= \max_p \left[\frac{U_c}{P} [p\mathcal{R}(p, s_{-1}) - \kappa_p W] + \beta EV(p_{-1}, s'_{-1}, a') \right], \\
V^{n,n}(p_{-1}, s_{-1}, a) &= \left[\frac{U_c}{P} p\mathcal{R}(p_{-1}, s_{-1}) + \beta EV(p_{-1}, s'_{-1}, a') \right],
\end{aligned}$$

where $(z - s_{-1})$ is the amount ordered.

The law of motion for inventories is

$$s'_{-1} = \left(z - \min \left(v \left(\frac{p}{P} \right)^{-\theta} c, z \right) \right) (1 - \delta)$$

if the firm orders, and

$$s'_{-1} = \left(s_{-1} - \min \left(v \left(\frac{p}{P} \right)^{-\theta} c, s_{-1} \right) \right) (1 - \delta)$$

otherwise. Here δ is the rate at which the firm's existing inventories depreciate.

Intermediate good firms (wholesalers)

We assume that a continuum of perfectly competitive intermediate good firms produce the wholesale good using a production technology that is linear in labor

$$y = l.$$

They sell these goods to retailers at price ω and thus earn profits equal to

$$\pi = \omega y - Wl ,$$

where W is the nominal wage rate.

Equilibrium

In equilibrium free entry drives intermediate's profits to zero and thus $\omega = W$. We impose a quantity-theory equation $M = \int_0^1 p(i)c(i)di$ rather than derive the demand for money explicitly in order to avoid introducing the labor-leisure etc. distortions associated with other specifications (cash-in-advance or money-in-the-utility). We study an economy in which M is constant and then compute the response to a one-time unanticipated change in M . The equilibrium is then a path of prices, $P, p(i), W, \omega$, as well as quantities, $c, c(i), y(i), s(i), z(i), l, n$, such that firms and consumers optimize and markets clear.

A. Economy with no adjustment frictions

To build intuition for the workings of our economy, we first consider the case when the two adjustment costs are equal to 0 and there is no irreversibility constraint. We assume however that the firm cannot return its excess inventories after learning its preference shock: it can return back to the goods market only at the beginning of the next period. Finally, we assume away variability in productivity shocks, a , for now. In the absence of frictions, and in the absence of $z \geq s_{-1}$ constraint, the firm's value is linear in s_{-1} (with slope equal to the replacement cost, ω), and the choice of p and z reduces to:

$$\max_{p_t, z_t} \left[p_t - \frac{(1-\delta)}{1+r_t} \omega_{t+1} \right] \mathcal{R}(p_t, z_t) - \left[\omega_t - \frac{(1-\delta)}{1+r_t} \omega_{t+1} \right] z_t ,$$

where $\frac{1}{1+r_t} = \beta \frac{U_{c,t+1}/P_{t+1}}{U_{c,t+1}/P_t}$ is the rate at which the firm discounts the future.

To understand this expression, recall that the only difference now between our problem and the standard problem of a monopolistically competitive firm is that prices and inventories must be decided prior to the realization of demand. Thus, the choice of price is determined by solving the standard problem of a monopolist, where $\mathcal{R}(p, z)$ is the (expected) demand

function and $\frac{(1-\delta)}{1+r_t}\omega_{t+1}$ (as opposed to ω_t) is the marginal valuation of the goods the firm sells. The difference stems from the fact that the savings from lowering sales by one unit today only accrue next period when the firm returns to the goods market. The choice of z is also straightforward: a higher z increases expected sales, but the firm loses $\left[\omega_t - \frac{(1-\delta)}{1+r_t}\omega_{t+1}\right]$ in inventory carrying costs.

The firm's optimal price is then a markup over its shadow valuation of inventories:

$$p = \frac{\varepsilon_t}{\varepsilon_t - 1} \frac{(1 - \delta)}{1 + r_t} \omega_{t+1} ,$$

where $\varepsilon_t = -\frac{\mathcal{R}_p(p_t, z_t)p_t}{\mathcal{R}(p_t, z_t)}$ is the elasticity of (expected) sales. Moreover, the choice of z_t satisfies

$$1 - \Phi \left(\frac{z_t}{\left(\frac{p_t}{P_t}\right)^{-\theta} c_t} \right) = \frac{1 - \frac{(1-\delta)}{1+r_t} \frac{\omega_{t+1}}{\omega_t}}{\frac{p_t}{\omega_t} - \frac{(1-\delta)}{1+r_t} \frac{\omega_{t+1}}{\omega_t}} .$$

Notice that the left-hand side of this expression is the probability that the firm stocks out. As in Bils and Kahn (2000), the firm chooses a higher stock of inventories (a lower stockout probability) the higher the markup $\frac{p_t}{\omega_t}$, the higher the growth rate of costs, and the lower the cost of holding inventories, $\frac{1-\delta}{1+r_t}$.

To further build intuition for the mechanism at play and illustrate what is driving the elasticity of inventories with respect to markups and costs in the full version of the model, we find it useful to log-linearize the above expression:

$$\Gamma \hat{q}_t = \left[(1 - \Phi) \mu \hat{\mu}_t + \beta (1 - \delta) \Phi \hat{\Delta}_t \right]$$

where $q_t = \frac{z_t}{\left(\frac{p_t}{P_t}\right)^{-\theta} c_t}$ is the (beginning-of-period) inventory to (expected) sales ratio, $\mu_t = \frac{p_t}{\omega_t}$ is the markup, $\Delta_t = \frac{\omega_{t+1}}{(1+r_t)\omega_t}$ is the change in costs, $\Gamma = \Phi'q[\mu - \beta(1 - \delta)]$ is a constant, variables with a hat denote log-deviations from steady-state, and $1 - \Phi$ is the steady state probability of a stockout. This expression says that the inventory-to-sales ratio is much more sensitive to changes in costs than to changes in markups. This is because in the calibration of the model we will use, as in the data, stockout probabilities are low, i.e., Φ is close to 1

and so are depreciation rates and markups. Clearly, as the rate of depreciation increases, the sensitivity of stock of inventories to costs decreases, a point we return to below when we study the full model.

B. Parametrization

We next study our full model's quantitative implications. The period is a month, and thus the discount factor is equal to $.96^{\frac{1}{12}}$. We assume preferences of the form $U(c_t, n) = \log c_t - \psi_t n_t$. We think of ψ_t as preference shocks that in the various experiments will be allowed to vary with the monetary shock in order to generate stickiness in nominal wages and thus real rigidities. We assume $\log(a)$ takes one of three possible values: $\{-\bar{a}, 0, \bar{a}\}$. The

transition matrix is $\Pr(a = a_i | a_{-1} = a_j)$, and it is equal to
$$\begin{bmatrix} 1 - \rho & \rho & 0 \\ \frac{\rho}{2} & 1 - \rho & \frac{\rho}{2} \\ 0 & \rho & 1 - \rho \end{bmatrix}$$
. The

firm's log productivity is on average equal to 0 and shocks arrive with a Poisson-like process with probability ρ . The discrete arrival of shocks allows the model to generate the large price changes observed in the data without a large selection effect²⁰; Midrigan (2008) provides empirical evidence for this assumption. The upper bound on technology shocks, \bar{a} , is chosen to match an average size of price changes of 10%, as documented by Bilal and Klenow (2005) and Klenow and Kryvtsov (2008). We choose ρ to match the frequency of price changes of 5 months given that our model has no motive for temporary price changes. A recent paper by Kehoe and Midrigan (2008) shows that a model with a motive for temporary price changes (sales) in which prices change on average every 3 weeks (as in the data on prices in retail stores), generates real effects of money similar to those of a model without a motive for sales in which 'regular' prices change every 5 months. The intuition as to why real effects of money are larger is that temporary price changes most often return to their pre-change level and, as a result, do not respond to low frequency variations in monetary policy.

We set κ_p , the fixed cost of changing prices, to 0.7% of steady-state revenue, as measured directly in grocery stores in a study by Levy et. al, 1997. Given the infrequent arrival of idiosyncratic shocks to productivity, the level of κ_p has little effect on the frequency of price changes and only affects how the fraction of price changes responds to a monetary shock.

²⁰Midrigan (2008), Gertler-Leahy (2007).

The fixed ordering cost, κ_S , together with the volatility of demand, σ , are chosen so that the model generates a monthly inventory-to-sales ratio of 1.4 months, as in the U.S. retail trade sector, and a correlation between changes in prices and quantities of -0.2 at the monthly frequency. The last statistic was reported by Burstein and Hellwig (2007) for regular price changes at the monthly frequency using Dominick’s scanner data. The depreciation rate is set to 2.5%, a number in the range of estimates of inventory-carrying costs reported by Richardson (1995) and Stock and Lambert (2001). We think of δ as not only capturing physical depreciation of goods, but also obsolescence, shrinkage as well as additional inventory carrying costs other than foregone interest, taxes, insurance or rent costs. Finally, the elasticity of demand, θ , is set equal to 3, in the mid-range of estimates in retail markets.

Table 5 reports the parameter values used. Table 6 reports the moments targeted. In addition to the 4 moments we use in calibration to pin down $\rho, \bar{a}, \sigma, \kappa_s$ (the rest of the parameters are assigned), we do well at matching additional features of the data. In particular, the model predicts that firms stockout 5% of the time and that 56% of firms order inventories in a given week. These numbers are in the range of those reported in earlier work by Bils (2004) and Aguirregabiria (2005), who place the frequency of stockouts at 5%-8%;²¹ Aguirregabiria (1999), who observes a 79% frequency of orders in the data on chain-wide orders of a Spanish supermarket, and Alessandria et. al (2008), who use data from Hall-Rust (1999) on the orders of a large steel wholesaler and find a frequency of orders for domestic goods of 0.33 to 0.50 per month.

C. Decision rules

Before we proceed with experiments, we briefly discuss optimal firm policy rules in our environment. We refer the reader to Aguirregabiria (1999) for a more detailed analysis of how fixed costs of ordering and price changes affect a retailer’s optimal price and ordering decisions, and proofs (that rely on the K-concavity of the value functions) of the uniqueness of optimal policy rules.

²¹This range for the frequency of stockouts is consistent with a large marketing literature on the incidence and costs of stockouts. For example, in the most comprehensive worldwide study of stockouts, Gruen et al. (2002) report a range of 5%-10%. They also find that retailers are likely to lose almost one-half of the intended purchases when a consumer confronts an out-of-stock: around 40% of the time consumer does not buy at all from the store, and remaining 60% of time consumers substitute the purchase (within or across brands, or across time). Hence after taking into account the possibility of within-retailer switching and backlogged orders, an estimate of the frequency of stockouts is 2.5%.

Notice from the Bellman equations characterizing the firm's dynamic program, that the optimality conditions governing the choice of inventories and prices are (to simplify notation, we write these assuming away both uncertainty about a and the irreversibility constraint):

$$p : \frac{U_c}{P} [R(p, z) + pR_p(p, z)] + \beta \int_0^\infty V_1(p, s'(v)) d\Phi(v) = -\theta\beta(1 - \delta) \int_0^{v^*} V_2(p, s'(v)) d\Phi(v) ,$$

$$z : \beta(1 - \delta) \int_0^{v^*} V_2(p, s'(v)) d\Phi(v) = \frac{U_c}{P} [\omega - p(1 - \Phi(v^*))] ,$$

where $v^* = \frac{z}{\left(\frac{p}{P}\right)^{-\theta} C}$ is the cutoff level of the taste shock for which the firm does not stockout.

The firm thus orders inventories so as to equalize the marginal valuation of an additional unit of inventories next period to the marginal cost of acquiring it (which with stockouts is lower than the replacement cost, $\frac{U_c}{P}\omega$, as an additional unit of inventories allows the firm to avoid a stockout). Similarly, the firm chooses the price to equalize marginal revenue (which here includes the savings of the price adjustment cost, as measured by $V_1(p)$) to the marginal valuation of inventories next period. Notice here again the sense in which inventory behavior can be informative about the wages (wholesale prices) that the firm faces. The stickier wages are (the lower $\frac{U_c}{P}\omega$ in times of monetary expansion), the more incentive the firm has to increase inventories s' , so as to lower their marginal valuation (since inventory frictions imply that V_2 is (K)-concave), and to keep prices low.

Figure 4 plots the regions of the (p_{-1}, s_{-1}) -state space in which the firm finds it optimal to adjust both inventory and price (heavy-shaded regions in the upper and lower left corners), adjust only inventory (medium-shaded region to the left), or adjust only price (light-shaded regions to the right). We fix the productivity, a , to its mean here. The region where inaction along both margins is optimal is the unshaded center region. We also plot the optimal price (solid line) and order decisions (we plot z , the inventory available for sale, using dotted lines) conditional on the firm adjusting both prices and inventories. The figure corresponds to the medium level of retailer's productivity (higher/lower levels of productivity would shift the graph to the south-east/north-west corners of the figure).

Given the fixed costs, the adjustment regions are of the (S, s) type: the firm adjusts

its price or orders inventories when its old price or current inventories are too far out of line. Conditional on ordering, the return point for inventories is independent of (p_{-1}, s_{-1}) as both of them are reset (the order decision is drawn for a firm with a stock s_{-1} sufficiently low so that the irreversibility constraint does not bind). Finally, notice that the firm's optimal price, conditional on adjustment, is approximately equal to its frictionless optimum up to a point at which the irreversibility constraint on inventories binds and the firm's stock s_{-1} exceeds the return point. As s_{-1} increases above this point, the firm finds it optimal to run down its inventories by decreasing its price so as to sell more rapidly and avoid paying the depreciation cost.²² Not reported here is the optimal price of the firm, conditional on not ordering new inventories. This price is strictly decreasing in the amount of inventories available (and in particular, above the price conditional on ordering in Figure 4, except in the region to the right of the vertical line in which the two are equal) as a firm that does not order has higher shadow valuation of inventories.

Given that our focus is on the behavior of the inventory-to-sales ratio, we report in Figure 5 the (expected) inventory-to-sales ratio given the optimal inventory decisions in the p_{-1} space, again for $a = 0$. In particular, we report $\frac{z-R(p,z)}{R(p,z)}$ conditional on ordering (left panel) and not ordering (right panel), as a function of the firm's past price and given the firm's optimal price adjustment decision. Notice in the left panel of Figure 5 that when the past price is sufficiently out of line, the firm adjusts it and orders an amount that is independent of the past price. If the firm's past price is close to its optimum, the firm leaves it unchanged, but chooses an inventory level that, although decreasing in the past price, results in an expected inventory-to-sales ratio that increases in the firm's price. As in the simpler example we presented above, a higher price makes an additional unit of sales more valuable as the gains from avoiding a stockout are larger. In the right panel of the Figure we plot the same inventory-sales ratio conditional on the firm not ordering. In this case increases in the firm's price increase the inventory-sales ratio to an even greater extent by reducing sales without affecting inventories. This figure thus illustrates that in response to a decrease in markups, induced by a monetary expansion, firms would, holding all else constant, hold less inventories relative to sales, even conditional on ordering inventories.

²²Alessandria, Kaboski, Midrigan (2008) study how this feature of the model can account for the slow pass-through of imported goods prices at the retail level following a large devaluation.

4. Experiments

We consider a one-time 2% increase in the stock of money in period 1. This increase is announced at date 1 before firms make their price/inventory decisions. We derive decision rules for firms that take into account this increase in the stock of money, and the transition path for aggregate prices and quantities. This path is in turn solved for by using a shooting algorithm in which we require that the transition path implied by firms' decision rules is consistent with the path used to derive those decision rules.

A. No real rigidities

In the benchmark economy the disutility from work, ψ_t , is constant and equal to ψ . The real wage in this model is thus equal to

$$\frac{W_t}{P_t} = -\frac{U_{n,t}}{U_{c,t}} = \psi C_t ,$$

and responds one-for one to changes in aggregate consumption. Moreover, the price of wholesale goods faced by firms, expressed in units of marginal utility of consumption, $\frac{\omega_t U_{c,t}}{P_t} = \frac{W_t U_{c,t}}{P_t}$ (recall that this is the relevant price that affects ordering decisions) is irresponsive to the monetary shock: absent the drop in markups (and in the absence of the fixed costs of ordering which generate inaction for a subset of firms), the inventory-to-sales ratio would not respond to the monetary shock as the real cost of purchasing goods from the wholesale is constant during the transition.

Figure 6 plots impulse responses in this benchmark economy. Here the inventory-sales ratio drops, mainly because of the unexpected increase in consumption, but also because new orders drop below their steady-state level due to the drop in markups. Notice that inventories oscillate: the initial drop in the fraction of firms ordering increases the mass of firms that find it optimal to order next period, etc. These echo effects can be smoothed out by introducing heterogeneity in the size of ordering costs. Table 7 reports the average deviation of the real wage $\frac{W}{P}$ from its steady state level, relative to the deviation of consumption, C , from its steady-state value in the first 5 periods after the shock, as well as a similar measure for the elasticity of inventory-to-sales to sales. As seen above, the real wage moves one for one with consumption, whereas the price of purchasing inputs from wholesalers (normalized by U_c) is constant. This leads to a drop of inventory-to-sales ratio that is 2.66 larger than the increase

in sales. Clearly, this elasticity is too high (in absolute value) relative to the data.

B. Real rigidities

We next assume that the marginal disutility from work, ψ_t , changes in response to the monetary expansion. In particular, we choose the path for ψ_t so that the nominal wage rate follows

$$W_t = W_t^{*(1-\lambda)} W_{t-1}^\lambda, \text{ where}$$

$$W_t^* = \psi P_t C_t$$

is the wage rate absent shocks to the marginal disutility from work. We interpret positive values of λ as features of the labor (or intermediate good's market) that imply real rigidities that are stronger than in our Benchmark setup. Larger values of λ correspond to stronger real rigidities.

Figure 7 reports impulse responses computed assuming $\lambda = \frac{1}{2}$. Notice that in this case real wages increase much less than consumption, C , and as a result, the price of intermediate goods, $\frac{\omega_t U_{c,t}}{P_t}$, drops initially and then rises. Firms thus find it optimal to purchase inventories immediately, rather than wait. This is seen in the lower panels of the Figure which show that orders increase by 20%, whereas inventories increase on impact by 13%. Table 7 shows that the elasticity of real wages to consumption is 0.53, while that of inventory-to-sales relative to sales increases to 4.56, much higher than in the data. To conserve space, we do not report the model's implications for other macroeconomic variables (output, hours etc.). Briefly, given that output includes inventory investment, the counterfactual inventory response also implies counterfactually large increases in real output (in this calibration output is an order of magnitude more volatile than consumption).

We next ask: what degree of real rigidities, as parameterized by the value of λ , is necessary for the model to account for the elasticity of the inventory-sales ratio to output in the data of -0.8? A search over λ yields a value of $\lambda = 0.15$, thus not too far from the benchmark case with no real rigidities. As Table 7 indicates, in this case the elasticity of real wages to consumption is equal to 0.92 and that of the inventory-sales ratio with respect to sales is equal to -0.79, as in the data.

As in the example above with no adjustment costs, these results are in large part driven by the fact that inventories in this economy are highly sensitive to variation in costs (as captured by $\frac{U_{c,t}}{P_t}\omega_t$). Inventories are sensitive to costs because in the parametrization we considered above the inventory carrying cost (the rate of depreciation) is fairly low, 2.5% per month on average. Firms thus find it optimal to raise the stock of inventories in response to a decline in the cost of acquiring them as little of the extra purchase will be lost to depreciation. To see this, recall that the optimality condition for inventory investment is:

$$\beta(1 - \delta) \int_0^{v^*} V_2(p, s'(v)) d\Phi(v) = \frac{U_c}{P} [\omega - p(1 - \Phi(v^*))]$$

In models with no frictions, V is linear in the stock of inventories, s , and the shadow value of inventories is equal to their replacement cost, $V_2 = \frac{U_c}{P}\omega$, as the firm can buy and sell inventories freely at the market price, ω . The economy we study allows for several sources of concavity in V : a fixed cost of ordering, a stockout probability, and an irreversibility. However, what really matters is the curvature of V at the return point (the inventory level that satisfies the first order condition) at which the stock of inventories is high and the first two frictions are muted (as the probability of a stockout or re-ordering again next period is low). Moreover, the irreversibility is quantitatively important only if the rate of depreciation is sufficiently high. As a result, the firm's valuation of inventories, $V(s)$ is close to linear, and small changes in the cost of inventories, $\frac{U_c}{P}\omega$, require large changes in the stock of inventories to satisfy the optimality condition. As above, changes in markups, as captured by p , play little role as the stockout probability, $1 - \Phi$, is close to 0, in the model and in the data.

This argument is reminiscent of two findings in recent work in related contexts. House (2007) shows that the sensitivity of investment in capital to its price in real business cycle models is high because of low capital depreciation rates. He argues that this feature of the model explains the Thomas (2002) and Khan and Thomas (2007) "irrelevance" result that an economy with fixed costs of adjusting capital is virtually indistinguishable from an economy without fixed costs. In the inventory literature, Jung and Yun (2005) study a sticky price model with inventories in which real rigidities are due to partial indexation of prices to inflation and the use of intermediate inputs as a factor of production. They show that one needs to assume high (in excess of 50% per quarter) depreciation rates in order to allow the

model's impulse responses of inventories to a monetary shock to be in line with those in the data. Our results above suggest that this sensitivity of inventories to aggregate shocks in the presence of low rates of depreciation is not specific to the exact reason that makes it optimal for firms in the model to hold inventories.

Our finding that an elasticity of real wages (marginal cost) to output of 0.92 is necessary to account for the behavior of inventories in the data stems from the assumption on preferences (log-utility) that we have made. More generally, given the lack of the curvature in V implied by our parametrization of the inventory frictions, what the model requires in order to prevent large fluctuations of inventories in response to business cycle shocks is the cost of inventories, $\frac{U_c}{P}\omega$, that is roughly constant over the cycle, or that the real marginal cost be roughly proportional to

$$\frac{\omega}{P} \sim \frac{1}{U_c}.$$

In other words, the elasticity of real marginal cost to consumption (which is typically assigned a value of one or higher) in order to account for the joint behavior of consumption and interest rates.

$$\frac{d \log \frac{\omega}{P}}{d \log c} \sim -\frac{U_{cc}}{U_c}c.$$

5. Extensions

In this section, we consider several extensions to the analysis above in order to gauge the robustness of our results.

A. No fixed costs of ordering

We first show that the sensitivity of inventories to costs is not driven to a great extent by the fixed costs of ordering. To see this, we reproduce the experiments above for an economy without fixed costs of ordering. We parameterize most parameters in this economy as we do above; the only difference is that we set $\kappa_s = 0$ and choose the volatility of taste shocks, $\sigma_v = 0.56$, somewhat higher than earlier, in order to match an inventory-to-sales ratio of 1.40, as in the data. Table 8 reports the results of the experiments. Clearly, without fixed costs of ordering the sensitivity of inventories to the monetary shock under various degrees of real rigidities is reduced. The reason is that now that firms order each period there are no variations in inventory investment along the extensive margin (fraction of firms ordering).

Nevertheless, we still find that an economy with an elasticity of real wages to consumption of one-half generates a procyclical inventory-to-sales ratio (the elasticity is 0.67, smaller than the 4.56 in the economy with fixed ordering costs). A search over the value of λ that allows the model to exactly replicate the elasticity of the inventory-sales ratio of -0.8 in the data implies an elasticity of real marginal cost to consumption of 0.77, somewhat smaller than the earlier 0.92.

B. Elasticity of inventories to costs in the micro-data

We next ask, does the model predict a counterfactually large sensitivity of inventories to costs relative to what we observe in the micro data? To do so, we employ a dataset of prices, inventories, sales, and costs in a Spanish supermarket dataset used by Aguirregabiria (1999) in his microeconomic study of markups and inventories in retail firms. We refer the reader to the original study for a description of the data. Here it suffices to say that we use a panel of monthly observations inventories and wholesale costs for 534 products (brands) sold by the supermarket chain for a period of 29 months from January 1990 to May 1992. Products are mostly non-perishable foodstuffs, e.g. beans, chocolate, canned fish, olive oil, dried fruits, and household goods, e.g., laundry detergents, soap, batteries. To gauge the sensitivity of inventories to costs, we estimate the following regression:

$$\log(s_{it}) = \alpha_i + \beta \log(c_{it}) + \varepsilon_{it}$$

where s_{it} is the end-of-period stock of inventories for product i in period t , and c_{it} is the wholesale price at which the retailer purchases the good from a supplier and α_i is a product-specific fixed effect. Table 9 reports the results of this regression in the data, as well as in the model with fixed costs we study above. The elasticity of inventories to costs in the model is roughly three times greater than in the data.

There are two potential interpretations of this result. One interpretation is that relationships between wholesalers are far more complex than what we have assumed here. Recall that we have assumed that quantities are demand-determined and that the wholesaler stands ready to sell whatever amount of the good the retailer demands. Nakamura (2008) cites evidence to support this interpretation and argues that wholesale prices are substantially more difficult to interpret than retail prices. Indeed, a typical pattern observed in microeconomic

data on wholesale and retail prices is that high-frequency variation in wholesale prices is typically highly correlated with high-frequency variation in retail prices. For example, Goldberg and Hellerstein (2008) and Eichenbaum, Jaimovich and Rebelo (2008) report that the retail price closely follows high-frequency movements in the wholesale price: a temporary price cut at the wholesale level is most likely also accompanied by a temporary price cut by the retailer, and vice versa: when the wholesaler raises his price back, the retailer follows. Such evidence is difficult to reconcile with the assumption that quantities are demand-determined. If this were indeed the case, the retailer would find it optimal to buy when the wholesale price is low, and gradually raise its retail price as the stock of inventory is depleted: the wholesaler's decision to raise its price after a temporary sale should have little consequence in this case on the retailer's price.

A second interpretation is that our model indeed understates the degree of curvature of the firm's shadow valuation of inventories. For example, higher rates of depreciation than what we have assumed may make it costlier for firms to raise the amount of inventories it carries in response to favorable cost shocks. This would also be able to reconcile the findings above that wholesale and retail prices are highly correlated at high frequencies with the assumption that quantities are demand-determined. Hence, we ask next: can increasing the rate of depreciation (inventory-carrying cost) account for the low sensitivity of inventories to costs in the data?

To conserve space, we briefly summarize the results of our experiments here. We have found that with depreciation rates of roughly 25% per month the model can indeed account for the elasticity of inventories to costs in the data at the micro-level. Clearly, this implies a lower elasticity at the aggregate level as well. However, we also find that with such high rates of depreciation it is difficult to account for the relatively high inventory holdings we observe in the data (0.70 in grocery stores, 1.30 at the retail level and 1.40 for Manufacturing and Trade as a whole). In the model, one needs excessively large volatility of idiosyncratic shocks to induce firms to hold inventories, which in turn generates counterfactual quantity implications. For example, with a $\delta = 0.25$ the model implies that firms stockout 30% of the time (5-8% in the data), most price changes are temporary price increases due to stockouts (in the data most price changes are temporary price cuts), as well as that price changes are much larger than in the data (10%). Thus, simply raising the rate of depreciation in the model allows one to account for one set of microeconomic facts only at the expense of others.

Jung and Yun (2005) suggest an alternative solution that may in principle reduce the sensitivity of inventories to costs: penalizing deviations of the firm's inventory-sales ratio from a given fixed target. We argue that modifications of the model along these lines would again produce counterfactual microeconomic implications. In particular, in Table 10 we report the volatility of the inventory-to-sales ratio in the Spanish supermarket data and compare it to that in the model. We measure volatility using the ratio of the interquartile range to the median, a statistic similar to the coefficient of variation, but less susceptible to outliers. As the Table reports, our model implies a volatility of the inventory-to-sales ratio, $\frac{s_{it}}{q_{it}}$, slightly lower than that in the data (the ratio of the interquartile range to the median is 1.24 in the model vs. 1.35 in the data). Thus introducing additional features of the model to reduce the unconditional volatility of the inventory-to-sales ratio would again produce counterfactual microeconomic implications.

C. Calvo time-dependent pricing and adjustment costs

In a companion paper²³ we extend the analysis above to a Calvo-type time-dependent setting with inventories. We show there that the results above are robust to this extension. In particular, we find again that inventories are highly sensitive to cyclical fluctuations in the cost of acquiring and holding them. This is true both in a simple Calvo model with a stockout-avoidance motive for holding inventories, as well as in a richer Smets-Wouters (2007) - type model. Thus even modest degrees of real rigidities (elasticities of real marginal cost to output lower than the inverse intertemporal elasticity of substitution) predict a counterfactually high increase in the inventory-to-sales ratio in response to monetary expansions. We also consider extensions in which we allow for adjustment costs on inventory investment (or more generally factors of production). These do allow the model to simultaneously match the behavior of real wages and other time-series in the data. These modifications imply, however, that a firm's shadow valuation of inventories increases sharply during booms despite the sluggishness of factor prices. As a result, the model's implications for the behavior of inflation resemble those of a model with little real rigidities and no adjustment costs.

²³Kryvtsov and Midrigan (2008).

6. Conclusions

As pointed out by Bilts and Kahn, the behavior of inventories is informative about how markups and real marginal costs vary with the cycle. We show here that a model which combines the stockout-avoidance and (s,S) motives for holding inventories, calibrated to match salient features of the microeconomic data, requires an elasticity of real marginal cost with respect to output of slightly less than the inverse intertemporal elasticity of substitution, in order to account for the behavior of inventories in the aftermath of monetary policy shocks. A greater degree of real rigidities implies a counterfactually high increase in the inventory-to-sales ratio in response to monetary expansions. In contrast, in the data this ratio persistently declines during booms. These results are robust to a number of variations of the model and are mainly driven by the relatively low inventory carrying costs (rates of depreciation) that are observed in the data (and are necessary to reconcile the model with the relatively large average inventory-to-sales ratios). We conclude that the data on inventories pose a challenge for models in which real rigidities take the form of low sensitivity of real marginal cost to output.

To us, the most promising approach to resolving the challenge posed by the data on inventories for models of real rigidities are frictions in the asset market that disconnect the real interest rate implied by consumer's pricing kernel from the opportunity cost of funds of the firms. In particular, Blinder and Maccini (1991) review evidence that suggests that inventories are insensitive to real interest rates²⁴. Allowing for frictions that limit the firms' ability to transform inventories into interest-paying balances in the asset market, as in the segmented asset market model of Alvarez, Atkeson and Kehoe (2002) or in the model of regime-switching and learning by Maccini, Moore and Schaller (2004), may be promising in tackling this challenge. Other potentially useful extensions include a) allowing additional sources of countercyclical markups (other than nominal price rigidities) that would decrease the benefits of carrying inventories during booms, and b) additional frictions on the firms' ability to purchase and carry inventories (e.g., non-linear rates of depreciation, capacity constraints) that reduce the sensitivity of inventories to costs. Exploring alternative mechanisms

²⁴Although, as our model suggests, this evidence must be interpreted with caution, as it is the interest rates times the expected change in the cost of purchasing inventories, that determines the cost of holding inventories. In our Benchmark model, inventories are also fairly insensitive to variations in interest rates because costs of purchasing inventories rise sufficiently in booms to offset the drop in the real interest rate.

that can address the disconnect between inventory data and models of real rigidities is an interesting avenue for future research.

References

- [1] Alvarez, Fernando, Andrew Atkeson and Patrick Kehoe, 2002, “Money, Interest Rates, and Exchange Rates with Endogenously Segmented Markets,” *Journal of Political Economy*, 110: 73-112.
- [2] Aguirregabiria, Victor, 1999. “The Dynamics of Markups and Inventories in Retailing Firms,” *The Review of Economic Studies*, 66(2), 275-308.
- [3] Aguirregabiria, Victor, 2005. “Retail Stockouts and Manufacturer Brand Competition,” mimeo.
- [4] Alessandria, George, Joseph Kaboski, and Virgiliu Midrigan, 2008. “Inventories, Lumpy Trade, and Large Devaluations,” Federal Reserve Bank of Philadelphia Working Paper 08-3.
- [5] Ball, Laurence and David Romer, 1990. “Real Rigidities and the Non-Neutrality of Money,” *The Review of Economic Studies*, 57(2), 183-203.
- [6] Barro, Robert, 1977, “Long-Term Contracting, Sticky Prices, and Monetary Policy,” *Journal of Monetary Economics*, 3: 305-316
- [7] Bils, Mark, 2004. “Studying Price Markups from Stockout Behavior,” mimeo.
- [8] Bils, Mark and James A. Kahn, 2000. “What Inventory Behavior Tells Us About Business Cycles,” *American Economic Review*, 90(3), 458-481.
- [9] Bils, Mark and Peter J. Klenow, 2004. “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112, 947-985.
- [10] Bils, Mark, and Klenow, Peter J. and Oleksiy Kryvtsov, 2003. “Sticky Prices and Monetary Policy Shocks,” *Federal Reserve Bank of Minneapolis Quarterly Review*, Winter, 2-9.

- [11] Blinder, Alan and Louis Maccini, 1991, "Taking Stock: A Critical Assessment of Recent Research on Inventories," *Journal of Economic Perspectives*, 5(1): 73-96.
- [12] Burstein, Ariel and Christian Hellwig, 2007. "Prices and Market Shares in a Menu Cost Model," mimeo.
- [13] Caplin, Andrew, 1985. "The Variability of Aggregate Demand with (S,s) Inventory Policies," *Econometrica*, 53(6), 1395-1409.
- [14] Chang, Yongsung, and Hornstein, Andreas and Pierre-Daniel Sarte, 2006. "Understanding How Employment Responds to Productivity Shocks in a Model with Inventories," Federal Reserve Bank of Richmond Working Paper 2006-06.
- [15] Chari, V. V. and Patrick J. Kehoe and Ellen R. McGrattan, 2000. "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?," *Econometrica*, 68(5), 1151-1180.
- [16] Christiano, Lawrence J. and Eichenbaum, Martin and Charles L. Evans, 1999. "Monetary policy shocks: What have we learned and to what end?," *Handbook of Macroeconomics*, in: J. B. Taylor & M. Woodford (ed.), *Handbook of Macroeconomics*, edition 1, volume 1, chapter 2, pages 65-148 Elsevier.
- [17] Christiano, Lawrence J. and Eichenbaum, Martin and Charles L. Evans, 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, University of Chicago Press, 113(1), 1-45.
- [18] Dotsey, Michael and Robert J. King, 2005. "Implications of state-dependent pricing for dynamic macroeconomic models," *Journal of Monetary Economics*, 52(1), 213-242.
- [19] Dotsey, Michael and Robert G. King, 2006. "Pricing, Production, and Persistence," *Journal of the European Economic Association*, 4(5), 893-928.
- [20] Eichenbaum, Martin, Nir Jaimovich and Sergio Rebelo, 2008. "Reference Prices and Nominal Rigidities," mimeo.
- [21] Eichenbaum, Martin and Jonas Fisher, 2004. "Evaluating the Calvo Model of Sticky Prices," mimeo.

- [22] Erceg, Christofer, Dale Henderson and Andrew Levin, 2000, “Optimal Monetary Policy with Staggered Wage and Price Contracts,” *Journal of Monetary Economics*, 46: 281-313.
- [23] Fisher, Jonas and Andreas Hornstein, 2000. “(S, s) Inventory Policies in General Equilibrium,” *The Review of Economic Studies*, 67(1), 117-145.
- [24] Gertler, Mark and Leahy, John, 2006. “A Phillips Curve with an Ss Foundation,” NBER Working Paper 11971.
- [25] Gruen, Thomas W and Corstein, Daniel S. and Sundar Bharadwaj, 2002. “Retail Out-of-Stocks: A Worldwide Examination of Extent, Causes and Consumer Responses,” the Grocery Manufacturers of America.
- [26] Hall, Robert, 2005, “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 95(1): 50-65.
- [27] Hall, George and John Rust, 2000. “An Empirical Model of Inventory Investment by Durable Commodity Intermediaries,” *Carnegie-Rochester Conference Series on Public Policy*, 52, 171-214.
- [28] Hall, George and John Rust, 2002. “Econometric Methods for Endogenously Sampled Time Series: The Case of Commodity Price Speculation in the Steel Market,” mimeo.
- [29] House, Christopher, L., 2008. “Fixed Costs and Long-Lived Investments.” *NBER Working Paper*, No. W14402.
- [30] Iacoviello, Matteo, Fabio Schiantarelli and Scott Schuh, 2007. “Input and Output Inventories in General Equilibrium,” Federal Reserve Bank of Boston Working Paper No. 07-16.
- [31] Jung Yongseung and Tack Yun, 2005, “Monetary Policy Shocks, Inventory Dynamics, and Price Setting Behavior ” mimeo
- [32] Kehoe, Patrick J. and Virgiliu Midrigan, 2007. “Sales and the real effects of monetary policy,” Working Papers 652, Federal Reserve Bank of Minneapolis.
- [33] Khan, Aubhik and Julia Thomas, 2007a. “Inventories and the business cycle: An equilibrium analysis of (S,s) policies.” *American Economic Review*, 97(4), 1165-88.

- [34] Khan, Aubhik and Julia Thomas, 2007b. “Explaining Inventories: A Business Cycle Assessment of the Stockout Avoidance and (S,s) Motives,” *Macroeconomic Dynamics*, 11(5), 638-64.
- [35] Khan, Aubhik and Julia Thomas, 2008. “Inventories and State-Dependent Pricing,”
- [36] Kimball, Miles, 1995. “The Quantitative Analytics of the Basic Neomonetarist Model,” *Journal of Money, Credit and Banking*, 27(4), 1241-1277.
- [37] Klenow, Peter J. and Oleksiy Kryvtsov, 2008. “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” forthcoming in *The Quarterly Journal of Economics*.
- [38] Klenow, Peter J. and Jon Willis, 2006. “Real Rigidities and Nominal Price Changes,” mimeo.
- [39] Kryvtsov, Oleksiy and Virgiliu Midrigan, 2008, “Inventories, Markups and Real Rigidities in New Keynesian Business Cycle Models,” mimeo.
- [40] Levy, Daniel, and Bergen, Mark, and Dutta, Shantanu, and Robert Venable, 1997. “The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains,” *The Quarterly Journal of Economics*, 112(3), 791-825.
- [41] Maccini, Louis J., and Bartholomew J. Moore, and Huntley Schaller, 2004. “The Interest Rate, Learning, and Inventory Investment,” *The American Economic Review*, 94(5), 1303-1327.
- [42] Nakamura, Emi and Jón, Steinsson, 2007. “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model,” mimeo.
- [43] Nakamura, Emi and Jón, Steinsson, 2008. “Five Facts About Prices: A Reevaluation of Menu Cost Models,” forthcoming in *The Quarterly Journal of Economics*.
- [44] Nakamura, Emi, 2008, “Pass-Through in Retail and Wholesale,” *American Economic Review*, 98(2), 430-437.
- [45] Ramey, Valerie A and Daniel J. Vine, 2004. “Why Do Real and Nominal Inventory-Sales Ratios Have Different Trends?,” *Journal of Money, Credit and Banking*, 36(5), 959-63.

- [46] Ramey, Valerie A. and Kenneth D. West, 1999. "Inventories," Handbook of Macroeconomics, in: J. B. Taylor & M. Woodford (ed.), Handbook of Macroeconomics, edition 1, volume 1, chapter 13, pages 863-923
- [47] Ribarsky, Jennifer, 2004, "Real Inventories, Sales, and Inventory-Sales Ratios for Manufacturing and Trade, 1997:I-2003:III," *Survey of Current Business*, February 1, 2004.
- [48] Richardson, H. 1995. "Control Your Costs then Cut Them" *Transportation and Distribution*, 94-96.
- [49] Romer, Christina D. and David H. Romer, 2004. "A New Measure of Monetary Shocks: Derivation and Implications," *The American Economic Review*, 94(4), 1055-1084.
- [50] Smets, Frank and Rafael Wouters, 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *The American Economic Review*, 97(3), 586-606.
- [51] Stigler George and James Kindhal, 1970, "The Behavior of Industrial Prices," Columbia University Press.
- [52] Solon, Gary and Barsky, Robert and Jonathan A. Parker, 1994. "Measuring the Cyclicity of Real Wages: How Important is Composition Bias," *The Quarterly Journal of Economics*, 109(1), 1-25.
- [53] Stock, James R. and Douglas M. Lambert, 2001. "Strategic Logistics Management," McGraw-Hill, 4th edition.
- [54] Wen, Yi, 2008. "Input and Output Inventory Dynamics," Federal Reserve Bank of St. Louis Working Paper 2008-008A.
- [55] Woodford, Michael, 2003, "Interest and Prices: Foundations of a Theory o Monetary Policy," Princeton University Press.

Table 1: Moments for Aggregate Inventory-to-Sales ratio

	Mean	Std	Serr corr	Correlation with output sales	
<i>NIPA monthly</i>					
Manufacturing and Trade	1.41	2.19	0.88	-0.83	-0.83
Retail	1.31	2.08	0.72	-0.49	-0.61
<i>NBER annual</i>					
Manufacturing	0.31	3.90	0.31	-0.52	-0.66

Note: Data are from the BEA National Income and Product Accounts monthly data from January 1967 to December 1997 and the NBER Manufacturing Productivity Database from 1957 to 1996. Sales are defined as real final sales to domestic purchasers in NIPA data and real value of industry shipments for NBER data. Output is the sum of final sales and the change in the end-of-period real inventory stock. Inventory-to-sales ratio is defined as the ratio of the end-of-period inventory stock to final sales in that period. All data are HP filtered. Output, sales and inventory-to-sales ratio are defined in % deviations from respective HP trends. Inventory investment is defined in % points-of-output-fraction deviations from its HP trend.

Table 2: Inventory-to-sales ratio by low, medium and high bins

	Level of Inventory-to-Sales Ratio			
	Low	Medium	High	Aggregate
	<i>Mean</i>			
Inventory-to-sales ratio	0.18	0.31	0.55	0.31
Sales share	0.47	0.27	0.26	1
	<i>Correlation with real GDP</i>			
Inventory-to-sales ratio	-0.41	-0.33	-0.47	-0.47

Note: The NBER Manufacturing Productivity Database from 1957 to 1996. The panel of inventory-to-sales ratios is divided into 3 bins corresponding to low, medium or high ratio. Means are un-weighted. Sales share is the fraction of total sales for inventory-to-sales in the bin to total sales in the panel.

Table 3: Elasticities of aggregate inventory-to-sales ratio

	output elasticity	sales elasticity
<i>NIPA monthly</i>		
Manufacturing and Trade	-0.77	-0.86
Retail	-0.49	-0.70
<i>NBER annual</i>		
Manufacturing	-0.42	-0.60

Note: Data are from BEA National Income and Product Accounts monthly data from January 1967 to December 1997 and the NBER Manufacturing Productivity Database from 1957 to 1996. Elasticities are regression coefficients with log inventory-to-sales as dependent variable and log output (or log sales) as independent variable, in addition to fixed time effects.

Table 4: Sales elasticities of industry-level inventory-to-sales ratio

	Across time	Cross-section
sales elasticity	-0.57	-0.12

Note: The NBER Manufacturing Productivity Database from 1957 to 1996. Across time elasticity is a panel regression coefficient with log industry inventory-to-sales as dependent variable and log industry sales as independent variable. Cross-section elasticity is a regression coefficient with mean log industry inventory-to-sales as dependent variable and mean log industry sales as independent variable. All regressions include fixed time and industry effects.

Table 5: Parameter values

κ_p	κ_s	\bar{a}	ρ	θ	δ	σ
0.007	0.019	0.09	0.21	3	0.025	0.47

Table 6: Moments in data and Model

	Data	Model
<i>Used in calibration</i>		
Frequency Δp	0.20	0.20
Mean $ \Delta p $ if adjust	0.10	0.10
Inventory-sales ratio	1.40	1.40
Correlation $\Delta \log p, \Delta y$	-0.23	-0.22
<i>Additional moments</i>		
Fraction stockouts	0.05-0.08	0.05
Frequency orders	0.33-0.79	0.56

Table 7: Elasticities of real wages and I/S to sales

		mean $\frac{\widehat{W/P}}{\widehat{C}}$	mean $\frac{\widehat{I/S}}{\widehat{S}}$
Benchmark	$\lambda = 0$	1	-2.66
Large real rigidities	$\lambda = \frac{1}{2}$	0.53	4.56
IS-consistent real rigidities	$\lambda = 0.15$	0.92	-0.79

Table 8: Economy without adjustment costs

	$\frac{\widehat{W/P}}{\widehat{C}}$	$\frac{\widehat{I/S}}{\widehat{S}}$ no cost
Benchmark	1	-2.29
Smaller real rigidities	0.92	-1.29
Large real rigidities	0.53	0.67
Inventory-consistent	0.77	-0.79

Table 9: Inventories and Costs in a Spanish Supermarket

	Data	Model
Dependent	$\log(s_{it})$	$\log(s_{it})$
$\log(c_{it})$	-2.38	-7.41
# obs.	14325	
R^2	0.02	

Table 10: Inventory-to-Sales ratio in a Spanish supermarket

	Data	Model
	$\frac{s_{it}}{q_{it}}$	$\frac{s_{it}}{q_{it}}$
median	0.65	1.44
$\frac{iqr}{median}$	1.35	1.24
# obs.	14182	

Figure 1: Inventory-Sales Ratio and Output
in US Manufacturing




Figure 2: Role of compositional effects




Figure 3: Response to +1% innovation to Fed Funds Rate
 Manufacturing and Trade. Jan '67 - Dec '96







Figure 4: Inaction regions and decision rules




Figure 5a: Expected I/S cond. on ordering & optimal price decision




Figure 5b: Expected I/S cond. on not ordering & optimal price decision




Figure 6: Benchmark







Figure 7: Large RR

