

Inventories, Markups, and Real Rigidities in Menu Cost Models

OLEKSIY KRYVTSOV

Bank of Canada

and

VIRGILIU MIDRIGAN

Federal Reserve Bank of Minneapolis and New York University

First version received August 2010; final version accepted May 2011 (Eds.)

A growing consensus in New Keynesian macroeconomics is that nominal cost rigidities, rather than countercyclical markups, account for the bulk of the real effects of monetary policy shocks. We revisit these conclusions using theory and data on inventories. We study an economy with nominal rigidities in which goods are storable. Our theory predicts that if costs of production are sticky and markups do not vary much in response to, say, expansionary monetary policy, firms react by excessively accumulating inventories in anticipation of future cost increases. In contrast, if the data inventories are fairly constant over the cycle and in response to changes in monetary policy. We show that costs must increase and markups must decline sufficiently in times of a monetary expansion in order to reduce firm's incentive to hold inventories and thus bring the model's inventory predictions in line with the data. Versions of the model consistent with the dynamics of inventories in the data imply that countercyclical markups account for a sizable fraction of the response of real variables to monetary shocks.

Key words: Inventories, markups, costs

JEL Codes: E31, F12.

1. INTRODUCTION

A widely held view in macroeconomics is that changes in monetary policy affect real economic activity because prices are sticky. What gives rise to price stickiness in the aggregate, and the extent to which prices are sticky is, however, a matter of considerable debate.¹

Since prices are equal to a markup times marginal costs, price stickiness in the aggregate can arise via one of two channels. One channel is countercyclical variation in markups, due to menu costs of price adjustment that prevent firms from changing their prices or other imperfections in the product market that make it optimal for firms to lower markups during booms.² A second channel, often referred to as real rigidities, is stickiness in costs. Frictions in the labour markets that give rise to wage rigidities, as well as the firms' ability to flexibly vary the workweek

1. See, for example, Chari *et al.* (2000) and Woodford (2003).

2. See Golosov and Lucas (2007) and Gertler and Leahy (2008) who study models with menu costs of price adjustment. Rotemberg and Woodford (1999) survey several alternative mechanisms that generate countercyclical markups.

of capital and labour, may imply that costs of production respond only gradually to monetary shocks.³

How strong each of these channels are has important implications for the strength of the monetary transmission mechanism, the role of nominal shocks in accounting for business cycle fluctuations, as well as the conduct of monetary and fiscal policy.⁴ Our goal in this article is to use theory and data in order to measure the extent to which markups and costs vary in response to monetary policy shocks.

Our approach is to study data on inventories through the lens of a New Keynesian model in which we introduce demand uncertainty and thus a stockout-avoidance motive for holding inventories. We focus on inventories because theory predicts a tight relationship between markups, costs, and inventories, as forcefully argued by Bils and Kahn (2000). The model predicts that when markups decline, firms reduce their stock of inventories relative to sales since inventories are less valuable when markups and profits are lower. Similarly, if costs are sticky after, say, an expansionary monetary shock, firms take advantage of the temporarily low costs by substituting intertemporally and building up a larger stock of inventories to run down in future periods.

A salient feature of the data is that the stock of inventories reacts much less to monetary policy shocks than sales do. The inventory–sales ratio is thus countercyclical and declines after an expansionary monetary policy shock. For our model to reproduce this fact, two conditions must be satisfied. First, costs of production must increase immediately after an expansionary monetary shock in order for firms not to substitute intertemporally by investing in inventories. Second, markups must decline to reduce the firm’s incentive to build up their stock of inventories in response to the lower interest rates that accompany episodes of monetary expansions. Overall, for the model to reproduce the behaviour of inventories in the data, countercyclical markups, rather than cost rigidities, must account for the bulk of the real effects of monetary policy shocks. Hence, as Bils and Kahn (2000) do, albeit using a different methodology and for monetary-driven business cycles, we find that markups are strongly countercyclical.

Our results stand in sharp contrast to a number of findings in existing work. A growing consensus in New Keynesian macroeconomics is that sticky nominal costs, rather than variable markups, account for the bulk of the response of real activity to monetary policy shocks. Studies of micro-price data find that input costs change infrequently and respond gradually to nominal shocks, while consumer prices tend to respond quickly to changes in costs.⁵ Moreover, Christiano *et al.* (2005) find that cost rigidities, as opposed to countercyclical markups, are necessary to account for salient facts of key U.S. macroeconomic time-series. These observations led researchers to conclude that wage rigidities and stickiness in input costs, rather than countercyclical markups, must be the dominant source of monetary non-neutrality.

Our article revisits these conclusions. Although input costs are indeed sticky in the data, the relationship between inputs costs and marginal costs is highly sensitive to what one assumes about the production technology, as Rotemberg and Woodford (1999) forcefully illustrate. In contrast, theory predicts a very robust relationship between the behaviour of inventories and that of marginal costs, which we document and exploit in this article.

3. See, for example, Christiano *et al.* (2005) and Dotsey and King (2006).

4. See Woodford (2003) and Galí (2008), who show how optimal monetary policy varies depending on whether the frictions are in the goods or labour markets. Woodford (2003) and Dotsey and King (2006) show how the size of the real effects from monetary shocks increases as one increases the degree of cost rigidities. More recently, Hall (2009) argues that countercyclical markups greatly amplify the effect of changes in government spending on output.

5. Bils and Klenow (2004); Klenow and Kryvtsov (2008); Goldberg and Hellerstein (2012); Eichenbaum *et al.* (2011). See also Nekarda and Ramey (2010) who argue that markups are, in fact, procyclical.

We begin our analysis by reviewing several well-known facts about inventories.⁶ In the data, inventories are procyclical but much less volatile than sales. The aggregate U.S. stock of inventories increases by about 0.16% for every 1% increase in sales during a business cycle expansion. We reach a similar conclusion when conditioning fluctuations on identified measures of monetary policy shocks. In response to an expansionary monetary policy shock, the stock of inventories increases by about 0.34% for every 1% increase in sales. Hence, the aggregate stock of inventories is relatively sticky and the aggregate inventory–sales ratio is countercyclical.

We then turn to the model. Our baseline model is characterized by price and wage rigidities, decreasing returns to scale, as well as convex adjustment costs that limit the firms' ability to rapidly change their stock of capital. These features imply that marginal costs are very responsive to monetary policy shocks: decreasing returns to labour make it costly for firms to change production very much, preventing them from varying their stock of inventories. Moreover, since prices are sticky, markups are countercyclical. These two features of the model, strongly procyclical marginal costs and countercyclical markups, imply that inventories are much less volatile than sales, as in the data. Importantly, countercyclical markups account for the bulk of the real effects of monetary policy shocks in our baseline model.

We then demonstrate that eliminating these two key ingredients of the baseline model—countercyclical markups and volatile marginal costs—visibly worsens the model's ability to account for the inventory facts. When we eliminate the decreasing returns to labour, marginal costs increase much more gradually in response to expansionary monetary shocks. Since the cost of carrying inventories is low, both in the model and in the data, firms find it optimal to invest in inventories in anticipation of future increases in production costs. The model thus predicts that inventories are much more volatile than in the data. Similarly, when we eliminate price rigidities and thus the countercyclical variation in markups, firms find it optimal to take advantage of the lower interest rates that accompany monetary expansions and excessively build up their stock of inventories. Absent a decline in markups to reduce the profits from holding inventories, the model fails to account for the inventory data.

We have studied a number of extensions of our model and have found that these results are robust to the exact model of inventories: the (S,s) model with fixed ordering costs as opposed to a stockout-avoidance model, details about the monetary policy rule, the rate at which inventories depreciate, the degree of demand uncertainty, as well as allowing inventories to be held at multiple stages of production. In all of these experiments, we found that markups must account for the bulk of the real effects of monetary shocks for the model to be able to account for the behaviour of inventories in the data. Our results thus stand in sharp contrast to the findings of Christiano *et al.* (2005), who estimate parameter values that imply essentially no role for markups in accounting for the real effects of monetary policy shocks.

Our work is related to a number of recent papers that study the behaviour of inventories, costs, and markups over the business cycle. Our starting point is the observation of Bills and Kahn (2000) that inventories are closely linked to markups and costs. The main difference between our work and that of Bills and Kahn is that they use data on input prices directly, together with a partial equilibrium model of inventories, in order to measure marginal costs. They find that the growth rate of marginal costs is acyclical, and hence the intertemporal substitution motive is weak in the data. They therefore conclude that markups must be countercyclical for the model to account for the countercyclical inventory–sales ratio in the data.

Khan and Thomas (2007) have recently argued that a countercyclical inventory–sales ratio is not necessarily evidence of countercyclical markups. They study the dynamics of inventories

6. See, for example, Ramey and West (1999) and Bills and Kahn (2000).

in a general equilibrium model driven by productivity shocks. Their model accounts well for the behaviour of inventories in the data, despite the fact that markups are constant. These authors show that general equilibrium considerations, and in particular capital accumulation, are critical to this result. Diminishing returns to labour reduce the response of marginal costs to a productivity shock and hence the incentive for inventory accumulation.

As Khan and Thomas (2007) do, we explicitly study the dynamics of inventories in a general equilibrium setting and find an important role for diminishing returns to variable factors in accounting for the inventory facts. While their focus is on productivity shocks, ours is on monetary shocks in an economy with nominal rigidities. We find that in our economy, countercyclical markups play an important role: absent markup variation, the model's predictions are grossly at odds with the data. The difference in our results stems from the special nature of monetary shocks in driving fluctuations in output. Unlike productivity shocks, monetary shocks affect real activity only if nominal prices are sticky and do not react immediately to changes in monetary policy. Since prices are equal to a markup times costs, monetary shocks can affect output only if either markups vary or if nominal costs are sticky. Hence, if markups are constant, monetary policy shocks can generate real effects only if nominal costs are sticky. Cost stickiness gives rise, however, to strong variability in inventories due to intertemporal substitution in production, which is at odds with the data on inventories.⁷

Finally, our article is closely related to the work of Klenow and Willis (2006) and Burstein and Hellwig (2007), who also measure the strength of real rigidities⁸ using theory and micro-price data. These researchers focus on an alternative type of real rigidity, in the form of strategic complementarities in price setting, and find weak evidence of such complementarities.

2. DATA

In this section, we review several salient facts regarding the cyclical behaviour of inventories. These facts are well known from earlier work.⁹ We discuss them briefly for completeness, as they are central to our quantitative analysis below.

We use data from the Bureau of Economic Analysis National Income and Product Accounts (NIPA) on monthly final sales and inventories for the U.S. Manufacturing and Trade sectors from January 1967 to December 2009.¹⁰ These two sectors of the economy account for most (85%) of the U.S. inventory stock; the rest of the stock is in mining, utilities, and construction.

All series are real. Our measure of sales is real final domestic sales. We define production as the sum of final sales and the change in the end-of-period inventory stock. We construct the inventory–sales ratio as the ratio of the end-of-period inventory stock to final sales in that period. When reporting unconditional business cycle moments, we detrend all series using a Hodrick–Prescott filter with a smoothing parameter equal to 14,400. We also use a measure of identified monetary policy shocks to report statistics conditional on identified exogenous monetary policy shocks.

Figure 1A presents the time series of sales and the inventory–sales ratio for the manufacturing and trade sectors. The figure shows that the two series are strongly negatively correlated and are

7. See Jung and Yun (2005), Chang *et al.* (2006), and Wen (2011), who also study the business cycle predictions of inventory models.

8. See Ball and Romer (1990).

9. See Ramey and West (1999) and Bills and Kahn (2000).

10. The 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).

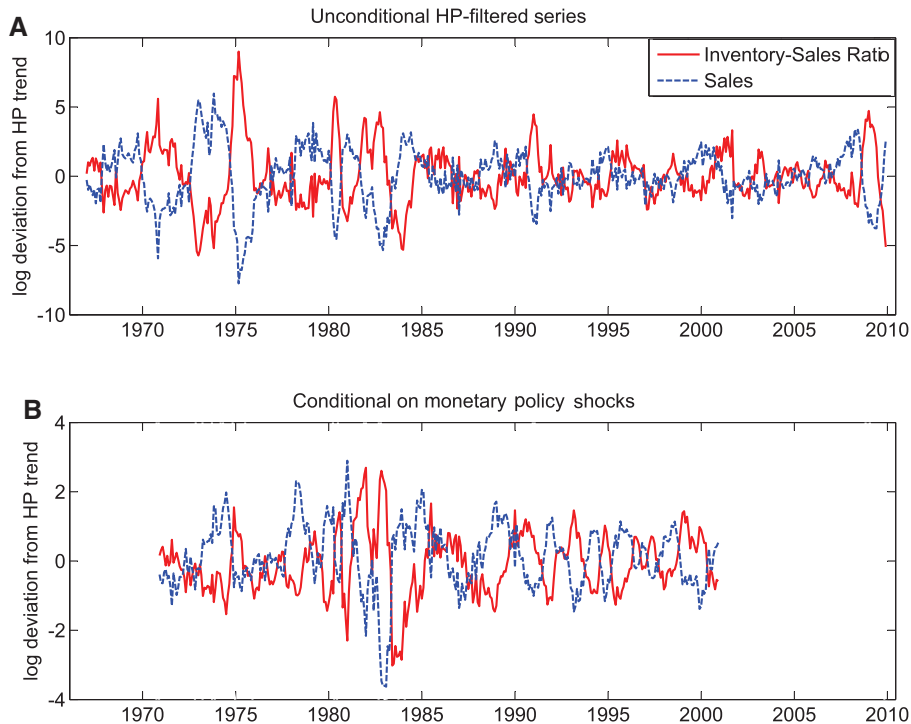


FIGURE 1

Inventories and sales, U.S. manufacturing and trade.

(A) Unconditional HP-filtered series and (B) Conditional on monetary policy shocks

almost equally volatile. Recessions are associated with a decline in sales and a similarly sized increase in the inventory–sales ratio. Likewise, expansions are associated with an increase in sales and a decline in the inventory–sales ratio of a similar magnitude.

Table 1 quantifies what is evident in the figure. The column labelled “Unconditional” reports unconditional statistics for these series. We focus on the series for the entire manufacturing and trade sector and later briefly discuss the statistics for the retail sector in isolation.

Notice in the first column of Table 1 that the correlation between the inventory–sales ratio and sales in manufacturing and trade is equal to -0.82 . The standard deviation of the inventory–sales ratio is about as large as the standard deviation of sales. Consequently, the elasticity of the inventory–sales ratio with respect to sales is equal to -0.84 .¹¹ In other words, the inventory–sales ratio declines by about 0.84% for every 1% increase in sales. The stock of inventories is thus fairly constant over the cycle, increasing by only 0.16% ($= -0.84 + 1$) for every 1% increase in sales. Note also that the inventory–sales ratio is very persistent: its autocorrelation is equal to 0.87.

The fact that the *stock* of inventories is fairly constant over time may seem to contradict the well-known fact that inventory investment is strongly procyclical and accounts for a sizable proportion of the volatility of GDP.¹² There is, in fact, no contradiction, since inventory *investment*

11. This elasticity is defined as the product of the correlation and the ratio of the standard deviations, or equivalently, as the slope coefficient in a regression of the log inventory–sales ratio on log sales.

12. See, for example, Ramey and West (1999).

TABLE 1
Inventory facts, U.S. NIPA, January 1967–December 2009

	Unconditional		Conditional on monetary shocks	
	Manufacturing and Trade	Retail	Manufacturing and Trade	Retail
$\rho(IS_t, S_t)$	−0.82	−0.65	−0.71	−0.57
$\sigma(IS_t)/\sigma(S_t)$	1.03	1.17	0.93	1.28
elast. IS_t w.r.t. S_t	−0.84	−0.76	−0.66	−0.73
elast. I_t w.r.t. S_t	0.16	0.24	0.34	0.27
$\rho(IS_t, IS_{t-1})$	0.87	0.72	0.88	0.74
$\rho(Y_t, S_t)$	0.98	0.89	0.98	0.84
$\rho(Y_t)/\sigma(S_t)$	1.12	1.14	1.11	1.17
$\rho(Y_t, \Delta I_t)$	0.55	0.47	0.63	0.50
$\rho(\Delta Y_t)/\sigma(Y_t)$	0.23	0.46	0.20	0.51

Notes: All series are real, monthly. IS_t , ΔI_t , S_t , Y_t denote real inventory–sales ratio, inventory investment, and final sales, respectively. The column labelled “Unconditional” reports statistics for HP (14 400)-filtered data. The column labelled “Conditional on monetary shocks” reports statistics computed using data projected on current and 36 lags of Christiano *et al.* (1999) measures of monetary policy shocks estimated using a VAR for 1960:01–2000:12.

is small relative to the entire stock of inventories: the average monthly inventory investment is equal to 0.22% of the inventory stock in the manufacturing and trade sector.

We next report the facts on inventory investment. One way to do so is by exploiting the following accounting identity:

$$Y_t = S_t + \Delta I_t,$$

where Y_t is production, S_t is sales and ΔI_t is inventory investment. To measure the volatility of inventory investment, we compare the standard deviation of production to that of sales (both expressed as log deviations from an HP trend). Notice in Table 1 that production and sales are strongly correlated and that production is 1.12 times more volatile than sales. Alternatively, the standard deviation of inventory investment (expressed, as is typical in the inventory literature, as a fraction of production, $\Delta I_t/Y_t$) is equal to 0.23 times the standard deviation of production. Also note that inventory investment is also strongly procyclical: its correlation with production is equal to 0.55. These last two statistics jointly imply that inventory investment (expressed as a fraction of production) increases by 0.13% ($= 0.55 \times 0.23$) for every 1% increase in production.

The other columns of Table 1 present several additional robustness checks. We note that the facts above also characterize the behaviour of firms in the retail sector: the elasticity of inventories to sales is equal to 0.24 and production is 1.14 as volatile as sales.¹³ These facts also hold when we condition on measures of monetary policy shocks. To see this, we project the data series on current and 36 lags of Christiano *et al.* (1999) measures of monetary policy shocks and recompute these statistics.¹⁴ We plot the resulting series in Figure 1B. Although monetary shocks account for a smaller fraction of the business cycle (the standard deviation of these series is about one-third as large when conditioning on measures of monetary shocks), the main patterns we documented above are evident now as well. As Table 1 shows, the elasticity of inventories to sales is now

13. Iacoviello *et al.* (2007) report that the inventory–sales ratio in retail is acyclical, as they find a low correlation between the inventory–sales ratio in retail and aggregate GDP. Their results are consistent with ours, since at the monthly frequency, aggregate GDP is imperfectly correlated with retail sales (the correlation of 0.10). These results thus reinforce our conclusion that the stock of inventories is fairly constant over the cycle.

14. Our results are robust to using an alternative measure of monetary policy shocks, due to Romer and Romer (2004).

equal to 0.34 (0.27 for retail) and production is 1.11 times more volatile than sales (1.17 in retail). Thus, in response to an expansionary monetary policy shock, both sales and inventory investment increase, but inventory investment increases much less than sales, and so the inventory–sales ratio declines.

The evidence in this section is robust to the detrending method, the level of aggregation and the stage of fabrication of inventories. We have also excluded the last five years of the sample to ensure that our results are not driven by the large recent recession and have found similar results. See our Supplementary Appendix for more details.

3. MODEL

We study a monetary economy populated by a large number of infinitely lived households and three types of firms: producers of intermediate goods, distributors, and final good firms. Intermediate good producers and final good firms are perfectly competitive. Distributors are monopolistically competitive, have sticky prices, and hold inventories.

In each period the commodities are differentiated varieties of labour services, a final labour service, money, intermediate goods, a continuum of differentiated varieties of goods sold by distributors, and a final good. The final good is used for consumption and investment.

In each period t , this economy experiences one of infinitely many events s_t . We denote by $s^t = (s_0, \dots, s_t)$ the history (or state) of events up to and including period t . The probability density, as of period 0, of any particular history s^t is $\pi(s^t)$. The initial realization s_0 is given.

The shocks in this economy are aggregate shocks to the money supply and idiosyncratic demand shocks. We describe the idiosyncratic shocks below. We assume that the supply of money follows a random-walk process of the form

$$\log M(s^t) = \log M(s^{t-1}) + g_m(s^t), \quad (1)$$

where $g_m(s^t)$ is money growth, a normally distributed i.i.d. random variable with mean 0 and standard deviation σ_m .

3.1. Households

Households consume, trade money and bonds, and work. We assume frictions in the labour market in the form of sticky wages. Households are organized in monopolistically competitive unions, indexed by j . Each union supplies a differentiated variety of labour services, $l_j(s^t)$, that aggregates into a final labour service, $l(s^t)$, according to

$$l(s^t) = \left(\int_0^1 l_j(s^t)^{\frac{\vartheta-1}{\vartheta}} dj \right)^{\frac{\vartheta}{\vartheta-1}}, \quad (2)$$

where ϑ is the elasticity of substitution across different types of labour services. Each union chooses its wage, $W_j(s^t)$, and faces demand for its labour services given by

$$l_j(s^t) = \left(\frac{W_j(s^t)}{W(s^t)} \right)^{-\vartheta} l(s^t), \quad (3)$$

where $l(s^t)$ is the amount of final labour hired by firms, and $W(s^t)$ is the aggregate wage rate:

$$W(s^t) = \left(\int W_j(s^t)^{1-\vartheta} dj \right)^{\frac{1}{1-\vartheta}}. \quad (4)$$

We assume that the markets for state-contingent money claims are complete. We represent the asset structure by having complete, state-contingent, one-period nominal bonds. Let $B_j(s^{t+1})$ denote the consumer's holdings of such a bond purchased in period t and state s^t with payoffs contingent on a particular state s^{t+1} at date $t+1$. One unit of this bond pays one unit of money at date $t+1$ if the particular state s^{t+1} occurs and 0 otherwise. Let $Q(s^{t+1}|s^t)$ denote the price of this bond in period t and state s^t . Clearly, $Q(s^{t+1}|s^t) = Q(s^{t+1})/Q(s^t)$, where $Q(s^t)$ is the date 0 price of a security that pays one unit of money if history s^t is realized.

The problem of union j is to choose its member's money holdings, $M_j(s^t)$, consumption, $c_j(s^t)$, state-contingent bonds, $B_j(s^{t+1})$, as well as a wage, $W_j(s^t)$, to maximize the household's utility:

$$\sum_{t=0}^{\infty} \int_{s^t} \beta^t \pi(s^t) [u(c_j(s^t)) - v(l_j(s^t))] ds^t \quad (5)$$

subject to the budget constraint

$$\begin{aligned} M_j(s^t) + \int_{s^{t+1}} Q(s^{t+1}|s^t) B_j(s^{t+1}) ds^{t+1} \\ \leq M_j(s^{t-1}) - P(s^{t-1}) c_j(s^{t-1}) + W_j(s^t) l_j(s^t) + \Pi_j(s^t) + B_j(s^t), \end{aligned} \quad (6)$$

a cash-in-advance constraint,

$$P(s^t) c_j(s^t) \leq M_j(s^t), \quad (7)$$

and subject to the demand for labour given by (3) as well as the frictions on wage setting we describe later. We assume that utility is separable between consumption and leisure.

Here, $P(s^t)$ is the price of the final good and $\Pi_j(s^t)$ are firm dividends. The budget constraint says that the household's beginning-of-period balances are equal to unspent money from the previous period, $M_j(s^{t-1}) - P(s^{t-1}) c_j(s^{t-1})$, labour income, dividends, as well as returns from bond holdings. The household divides these balances into money holdings and purchases of state-contingent bonds.

We assume Calvo frictions on wage setting. The probability that any given union is allowed to reset its wage at date t is constant and equal to $1 - \lambda_w$. A measure λ_w of the unions leave their nominal wages unchanged. We choose the initial bond holdings of unions so that each union has the same present discounted value of income. Even though unions differ in the wages they set and hence the amount of labour they supply, the presence of a complete set of securities and the separability between consumption and leisure implies that they make identical consumption and savings choices in equilibrium. Since these decision rules are well understood, we simply drop the j subscript and note that the bond prices satisfy

$$Q(s^{t+1}|s^t) = \beta \pi(s^{t+1}|s^t) \frac{u_c(c(s^{t+1}))}{u_c(c(s^t))} \frac{P(s^t)}{P(s^{t+1})}, \quad (8)$$

where $\pi(s^{t+1}|s^t) = \pi(s^{t+1})/\pi(s^t)$ is the conditional probability of s^{t+1} given s^t . Similarly, the date 0 prices satisfy

$$Q(s^t) = \beta^t \pi(s^t) \frac{u_c(c(s^t))}{P(s^t)}. \quad (9)$$

3.2. *Final good firms*

The final good sector consists of a unit mass of identical and perfectly competitive firms. The final good is produced by combining the goods sold by distributors (we refer to these goods as *varieties*) according to

$$q(s^t) = \left(\int_0^1 v_i(s^t)^{\frac{1}{\theta}} q_i(s^t)^{\frac{\theta-1}{\theta}} di \right)^{\frac{\theta}{\theta-1}}, \tag{10}$$

where $q_i(s^t)$ is the amount of variety i purchased by a final good firm, $v_i(s^t)$ is a variety-specific shock and θ is the elasticity of substitution across varieties. We assume that $v_i(s^t)$ is an i.i.d. lognormal random variable. We adopt the i.i.d. assumption for simplicity, as it allows us to characterize the firm’s inventory decisions in closed-form, although, as we show below, this assumption counterfactually implies a negative correlation between sales and inventory investment at the firm level. The data, in contrast, shows little correlation between sales and inventory investment, a feature that the (S,s) model we study in the Robustness section can account for.

At the beginning of period t , distributors have $z_i(s^t)$ units of the good available for sale. We describe how the distributors choose $z_i(s^t)$ below. Given the price and inventory adjustment frictions we assume, distributors will occasionally be unable to meet all demand and will thus stock out. We assume, in case of a stockout, a rationing rule under which all final good firms are able to purchase an equal share of that distributor’s goods. Since the mass of final good firms is equal to 1, $z_i(s^t)$ is both the amount of goods the distributor has available for sale, as well as the amount of goods that any particular final good firm can purchase.

The problem of a firm in the final good sector is therefore

$$\max_{\{q_i(s^t)\}} P(s^t)q(s^t) - \int_0^1 P_i(s^t)q_i(s^t)di, \tag{11}$$

subject to the constraint

$$q_i(s^t) \leq z_i(s^t) \quad \forall i \tag{12}$$

and the final good production technology (10). Cost minimization by the final good firms implies the following demand for each variety:

$$q_i(s^t) = v_i(s^t) \left(\frac{P_i(s^t) + \mu_i(s^t)}{P(s^t)} \right)^{-\theta} q(s^t), \tag{13}$$

where $\mu_i(s^t)$ is the multiplier on (12). Notice here that the shocks, $v_i(s^t)$, act as *demand* shocks for a distributor. Perfect competition implies that the price of the final good, $P(s^t)$, is equal to

$$P(s^t) = \left[\int_0^1 v_i(s^t) [P_i(s^t) + \mu_i(s^t)]^{1-\theta} di \right]^{\frac{1}{1-\theta}}. \tag{14}$$

Also note that if the inventory constraint binds, then $\mu_i(s^t)$ satisfies

$$P_i(s^t) + \mu_i(s^t) = \left(\frac{z_i(s^t)}{v_i(s^t)P(s^t)^{\theta}q(s^t)} \right)^{\frac{1}{\theta}}. \tag{15}$$

The left-hand side of this expression is the price that a distributor that stocks out would have chosen absent the price adjustment frictions. Since such a distributor faces inelastic demand, it

would like to increase its price to the point at which final good firms demand exactly all of its available goods. Together with the inventory frictions we describe later, price adjustment frictions give rise to stockouts in the equilibrium of this economy, since they prevent distributors from increasing their prices.

3.3. Intermediate good firms

There is a continuum of intermediate good firms that sell a homogeneous good to distributors. Any such firm owns its capital stock, hires labour and sells the good to the distributors. Intermediate good firms augment their capital stock by purchasing investment goods from final good producers at price $P(s^t)$. We assume that investment is subject to a convex cost of installing new capital.

The production function takes the Cobb–Douglas form

$$y(s^t) = \left(l(s^t)^\alpha k(s^{t-1})^{1-\alpha} \right)^\gamma, \quad (16)$$

where $y(s^t)$ is output, $k(s^{t-1})$ is the amount of capital a producer owns at date t , $l(s^t)$ is the amount of labour it hires, and γ is the degree of returns to scale.

Let $\Omega(s^t)$ be the price of the intermediate good. Recall that the price of the final good, used for investment, is $P(s^t)$ and the wage rate is $W(s^t)$. The intermediate good producer solves

$$\max_{y(s^t), x(s^t), l(s^t)} \sum_{t=0}^{\infty} \int_{s^t} \mathcal{Q}(s^t) [\Omega(s^t)y(s^t) - P(s^t)(x(s^t) + \phi(s^t)) - W(s^t)l(s^t)] ds^t \quad (17)$$

subject to

$$k(s^t) = (1 - \delta)k(s^{t-1}) + x(s^t)$$

and the production function (16). Here, δ is the depreciation rate of capital and $\phi(s^t)$ is the adjustment cost:

$$\phi(s^t) = \frac{\xi}{2} \left(\frac{x(s^t)}{k(s^{t-1})} - \delta \right)^2 k(s^{t-1}). \quad (18)$$

Letting $R(s^t) = (1 - \alpha)\gamma\Omega(s^t)y(s^t)/k(s^{t-1})$ denote the marginal product of capital, perfect competition and optimization by intermediate good firms imply the following relationship between the price of intermediate inputs, $\Omega(s^t)$, the wage rate and the firm's stock of capital:

$$\Omega(s^t) = \frac{1}{\gamma} (1 - \alpha)^{-(1-\alpha)} \alpha^{-\alpha} \left(W(s^t)^\alpha R(s^t)^{1-\alpha} \right) y(s^t)^{\frac{1}{\gamma}-1}. \quad (19)$$

3.4. Distributors

Distributors purchase goods from intermediate good firms at a price $\Omega(s^t)$, convert these to distributor-specific varieties and sell these varieties to final good firms. A key assumption we make is that the distributor chooses how much to order prior to learning the value of $v_i(s^t)$, its demand shock (but after learning the realization of all other shocks, including the monetary shocks). This assumption introduces a precautionary motive for holding inventories, the stockout-avoidance motive.

Distributors face two frictions. First, they must choose how much to order, $y_i(s^t)$, and the price to set, $P_i(s^t)$, prior to learning their demand shock, $v_i(s^t)$. Second, they change prices

infrequently, in a Calvo fashion. A fraction $1 - \lambda_p$ of randomly chosen distributors are allowed to reset their nominal prices in any given period; the remaining λ_p of distributors leave their prices unchanged.

Let $n_i(s^{t-1})$ denote the stock of inventories distributor i has at the beginning of date t . If the firm orders $y_i(s^t)$ additional units, the amount it has available for sale is equal to

$$z_i(s^t) = n_i(s^{t-1}) + y_i(s^t). \tag{20}$$

Given a price, $P_i(s^t)$, and an amount of goods available for sale, $z_i(s^t)$, the firm's sales are

$$q_i(s^t) = \min \left(v_i(s^t) \left(\frac{P_i(s^t)}{P(s^t)} \right)^{-\theta} q(s^t), z_i(s^t) \right). \tag{21}$$

The firm's problem is therefore to choose $P_i(s^t)$ and $z_i(s^t)$ so as to maximize its objective given by

$$\sum_{t=0}^{\infty} \int_{s^t} Q(s^t) (P_i(s^t) q_i(s^t) - \Omega(s^t) y_i(s^t)) ds^t, \tag{22}$$

where $n_i(s_0)$ is given. The constraints are the demand function in equation (21), the restriction that $z_i(s^t)$ and $P_i(s^t)$ are not measurable with respect to $v_i(s^t)$ (but measurable with respect to the money growth shocks), the constraint that $P_i(s^t) = P_i(s^{t-1})$ in the absence of a price adjustment opportunity, as well as the law of motion for inventories:

$$n_i(s^t) = (1 - \delta_z) (z_i(s^t) - q_i(s^t)), \tag{23}$$

where δ_z is the rate at which inventories depreciate.

3.5. Equilibrium

Consider now this economy's market-clearing conditions and the definition of the *equilibrium*. The market-clearing condition for labour states that the amount of the final labour service supplied by households is equal to the amount hired by the intermediate good firms:

$$\left(\int_0^1 l_j(s^t)^{\frac{\vartheta-1}{\vartheta}} dj \right)^{\frac{\vartheta}{\vartheta-1}} = l(s^t). \tag{24}$$

Similarly, the market-clearing condition for the final good states that consumption and investment sum up to the total amount of the final good produced:

$$c(s^t) + x(s^t) + \phi(s^t) = q(s^t) \tag{25}$$

The market-clearing condition for intermediate inputs is

$$y(s^t) = \int_0^1 y_i(s^t) di. \tag{26}$$

We can rewrite this further as

$$y(s^t) = \int_0^1 q_i(s^t) di + \frac{1}{1 - \delta_z} \int_0^1 (n_i(s^t) - (1 - \delta_z) n_i(s^{t-1})) di, \tag{27}$$

which says that total production is equal to sales plus inventory investment.

An *equilibrium* for this economy is a collection of allocations for households $c(s^t)$, $M(s^t)$, $B(s^{t+1})$, $l_j(s^t)$, and wages $W_j(s^t)$; prices and allocations for firms $P_i(s^t)$, $q_i(s^t)$, $y_i(s^t)$, $z_i(s^t)$, $n_i(s^t)$, $y(s^t)$, $l(s^t)$, $x(s^t)$, $k(s^t)$; and aggregate prices $W(s^t)$, $P(s^t)$, $\Omega(s^t)$, and $Q(s^t)$, all of which satisfy consumer and firm maximization and the market-clearing conditions.

We solve for the equilibrium numerically, using first-order perturbation methods that are standard in the New Keynesian literature. We show, in the Supplementary Appendix, that perturbation methods are very accurate in our context: global projection methods yield nearly identical results.

3.6. Decision rules

To build intuition for our results, we next discuss the decision rules in this economy and the determinants of the inventory–sales ratio. Combining equations (22) and (23), the firm’s problem of how many goods to make available for sale, $z_i(s^t)$, reduces to

$$\max_{z_i(s^t)} \left[P_i(s^t) - (1 - \delta_z) \int_{s^{t+1}} \frac{Q(s^{t+1})}{Q(s^t)} \Omega(s^{t+1}) ds^{t+1} \right] D(P_i(s^t), z_i(s^t), s^t) \quad (28)$$

$$- \left[\Omega(s^t) - (1 - \delta_z) \int_{s^{t+1}} \frac{Q(s^{t+1})}{Q(s^t)} \Omega(s^{t+1}) ds^{t+1} \right] z_i(s^t),$$

where

$$D(P_i(s^t), z_i(s^t), s^t) = \int \min \left(v \left(\frac{P_i(s^t)}{P(s^t)} \right)^{-\theta} q(s^t), z_i(s^t) \right) dF(v) \quad (29)$$

are the firm’s expected sales and $F(v)$ is the (lognormal) distribution of demand shocks, v .

Equation (28) shows that on the one hand, a higher $z_i(s^t)$ increases expected sales by reducing the probability of a stockout (the first term of the equation). On the other hand, the firm loses $\Omega(s^t) - (1 - \delta_z) \int_{s^{t+1}} \frac{Q(s^{t+1})}{Q(s^t)} \Omega(s^{t+1})$ in inventory carrying costs: the difference between the cost of ordering in the current period and the discounted cost of ordering in the next period (the second term of the equation). The solution to the problem in (28) is

$$1 - F(v_i^*(s^t)) = \frac{1 - r^I(s^t)}{b_i(s^t) - r^I(s^t)}, \quad (30)$$

where

$$v_i^*(s^t) = z_i(s^t) / \left[\left(\frac{P_i(s^t)}{P(s^t)} \right)^{-\theta} q(s^t) \right] \quad (31)$$

is the amount of goods available for sale scaled by a term that captures demand, while

$$r^I(s^t) = (1 - \delta_z) \int_{s^{t+1}} \frac{Q(s^{t+1})}{Q(s^t)} \frac{\Omega(s^{t+1})}{\Omega(s^t)} ds^{t+1} \quad (32)$$

is the return to holding inventories and

$$b_i(s^t) = P_i(s^t) / \Omega(s^t) \quad (33)$$

is the markup.

The left-hand side of equation (30) is the probability of a stockout: a higher inventory stock, relative to demand, lowers this probability. The right-hand side of this expression is decreasing in the return to holding inventories and the firm's markup. As in Bils and Kahn (2000), a higher return to holding inventories makes it optimal for firms to increase the amount of goods available for sale and lower the probability of a stockout. Similarly, a higher markup makes stockouts especially costly, since the profit lost by failing to make a sale is greater. Higher markups thus lead firms to lower the probability of a stockout by increasing the amount offered for sale.

We can gain some insight about how the amount of goods available for sale varies over time by log linearizing equation (30) around the steady state:

$$\hat{v}_i^*(s^t) = \frac{1}{(\bar{b} - \beta(1 - \delta_z))\bar{F}^t} \left[(1 - \bar{F})\bar{b}(\hat{P}_i(s^t) - \hat{\Omega}(s^t)) + \beta(1 - \delta_z)\bar{F}\hat{r}_i(s^t) \right],$$

where hats denote log deviations from the steady state, \bar{b} is the steady-state markup, and $1 - \bar{F}$ is the steady-state stockout probability. Notice that when the steady-state cost of carrying inventories, δ_z , is lower, the amount of goods available for sale is more sensitive to fluctuations in the return to holding inventories. Intuitively, if the cost of carrying inventories is low, firms are able to intertemporally substitute orders to take advantage of temporarily lower input prices.¹⁵

Also note that the return to holding inventories, $r^I(s^t)$, can be written up to a first-order approximation as

$$r^I(s^t) \approx (1 - \delta_z) \frac{\int_{s^{t+1}} \Omega(s^{t+1}) / \Omega(s^t) ds^{t+1}}{1 + i(s^t)}, \tag{34}$$

where $i(s^t)$ is the nominal risk-free interest rate:

$$i(s^t) = \left[\int_{s^{t+1}} \frac{Q(s^{t+1})}{Q(s^t)} ds^{t+1} \right]^{-1} - 1.$$

Equation (34) states that the return to holding inventories increases with the expected change in costs, $\Omega(s^{t+1}) / \Omega(s^t)$, and decreases with the nominal interest rate. Hence, the model's implications for how inventories react to monetary policy shocks depend on the expected change in costs, as well as the response of interest rates and markups.

So far we have discussed the model's implications for the amount of goods the firm makes available for sales, $v_i^*(s^t)$. This object, on its own, is not useful to evaluate the model empirically because we do not directly observe it in the data. Notice, however, that the model predicts a monotone relationship between the expected end-of-period inventory–sales ratio and $v_i^*(s^t)$. In particular, using equations (21) and (23) and integrating over the distribution of demand shocks, we have that the expected end-of-period inventory–sales ratio is equal to

$$IS_i(s^t) = \frac{v_i^*(s^t)F(v_i^*(s^t)) - e^{-\sigma_v^2/2}F(v_i^*(s^t)e^{-\sigma_v^2})}{e^{-\sigma_v^2/2}F(v_i^*(s^t)e^{-\sigma_v^2}) + v_i^*(s^t)(1 - F(v_i^*(s^t)))}, \tag{35}$$

and increasing in $v_i^*(s^t)$ for a lognormal distribution F .

15. See House (2008), who makes a similar argument in the context of a model with investment.

TABLE 2
Parameterization

Calibrated parameters			Targets	Data	Model
δ_z	Inventory depreciation	0.011	I/S ratio	1.4	1.4
σ_v	s.d. demand shocks	0.626	Frequency stockouts	0.05	0.05
ξ	Investm. adj. cost	46.2 (0.05% of x)	$\sigma(x_t)/\sigma(c_t)$	4	4
Assigned parameters					
	period	1 month			
θ	elast. subst. goods	5			
ϑ	elast. subst. labour	5			
$1 - \lambda_w$	freq. wage changes	1/12			
$1 - \lambda_p$	Freq. price changes	1/12			
α	Labor share	2/3			
δ	Capital depreciation	0.01			
γ	Returns to scale	0.9			
β	Discount factor	$0.96^{1/12}$			
σ	1/IES	2			
χ	Inverse Frisch elasticity	0.40			

We also briefly discuss how firms choose prices in this economy. Absent price rigidities, when $\lambda_p = 0$, the firm's price is simply a markup over its shadow valuation of inventories,

$$P_i(s^t) = \frac{\varepsilon_i(s^t)}{\varepsilon_i(s^t) - 1} (1 - \delta_z) \int_{s^{t+1}} \frac{Q(s^{t+1})}{Q(s^t)} \Omega(s^{t+1}), \quad (36)$$

where the markup depends on price elasticity of expected sales, $\varepsilon_i(s^t)$. This elasticity is equal to θ (the elasticity of substitution across varieties) times the share of sales in the states in which the firm does not stock out. Markups thus decrease with the distributor's inventory stock, since a greater stock of inventories lowers the probability of a stockout and raises the demand elasticity. With price rigidities, the firm's price, conditional on the firm being able to reset it, is simply a weighted average of future frictionless prices in (36).

4. QUANTITATIVE RESULTS

We next show that our model accounts well for the dynamics of inventories, production and sales in the data. We then trace this result to our assumptions on pricing and technology that imply that markups are countercyclical. In the aftermath of expansionary monetary policy shocks, costs of production rise quickly while prices do not, so that firms find it costly and less valuable to build up their stock of inventories. As a result, the inventory–sales ratio falls. We then show that versions of the model that imply that costs increase much more gradually and that markups are less countercyclical predict, counterfactually, that the inventory–sales ratio increases sharply after monetary policy expansions.

4.1. *Parameterization*

Table 2 reports the parameter values we used in our quantitative analysis. We set the length of the period equal to one month and therefore choose a discount factor of $\beta = 0.96^{1/12}$. We assume preferences of the form $u(c) - v(l) = c^{1-\sigma}/(1-\sigma) - l^{1+\chi}/(1+\chi)$. We set $\sigma = 2$, a commonly

used value in the business cycle literature. As we show below, our baseline parameterization requires an intertemporal elasticity of substitution, $1/\sigma$, less than unity in order to account for the drop in nominal interest rates after a monetary policy shock. We set $\chi = 0.4$, implying a Frisch elasticity of labour supply of 2.5, consistent with the estimates of Rogerson and Wallenius (2009).

We assume that the growth rate of the money supply is serially uncorrelated. We choose the standard deviation of money growth, σ_m , equal to 0.23% so that the model matches the standard deviation of the exogenous component (identified using the Christiano *et al.* (1999) measure of monetary policy shocks) of the U.S. monthly growth rate of the money supply.

We set the rate at which capital depreciates, δ , equal to 0.01 and the share of capital in production equal to $\alpha = 1/3$, standard choices in existing work. We follow Khan and Thomas (2008) and set the value of the span-of-control parameter, γ , equal to 0.9, in the range of values used in models of firm dynamics. We set the elasticity of substitution across intermediate goods (θ) equal to 5, implying a 25% markup. This elasticity is somewhat higher than estimates of price elasticities reported in IO studies—around 3, but lower than elasticities that are consistent with estimates of markups from production function estimates (*e.g.* Basu and Fernald (1997))—around 10. There is little agreement about the value of ϑ , the elasticity of substitution across varieties of labour. Christiano *et al.* (2005) set $\vartheta = 21$, implying a 5% markup, while Smets and Wouters (2007) set $\vartheta = 3$, implying a 50% markup. We set $\vartheta = 5$, implying a markup of 25%, in the middle of the values used in these two studies.

We choose the size of the capital adjustment cost, ξ , to ensure that the model reproduces the fact that the standard deviation of investment is four times greater than the standard deviation of consumption. The value of ξ is equal to 46.2, implying that resources consumed by capital adjustment costs account for about 0.05% of the total amount of investment in simulations of our economy. The size of these adjustment costs is somewhat smaller than that used by Chari, Kehoe, and McGrattan (2000), who find that adjustment costs range between 0.09% and 0.40% of the total investment.

We assume that wages and prices change on average once every 12 months ($\lambda_w = \lambda_p = 1 - 1/12$), consistent with what is typically assumed in existing studies. This degree of price stickiness is somewhat higher than what Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008) report, but is consistent with the findings of Kehoe and Midrigan (2010).

The parameters that are specific to our inventory model are the rate at which inventories depreciate, δ_z , and the volatility of demand shocks, σ_v . These two parameters jointly determine the steady-state inventory–sales ratio and the frequency of stockouts. We thus choose values for these parameters so that the model can reproduce the 1.4 ratio of inventories to monthly sales in the U.S. manufacturing and trade sector and the 5% frequency of stockouts that Bills (2004) reports using micro-price data from the Bureau of Labor Statistics (BLS). As Table 2 shows, matching an inventory–sales ratio of 1.4 and a frequency of stockouts of 5% requires a volatility of demand shocks of $\sigma_v = 0.63$ and a depreciation rate of 1.1%.

4.2. Aggregate implications

Figure 2 reports the impulse responses of nominal and real variables to a 1% expansionary shock to the money supply in our model with capital, decreasing returns and sticky prices. We refer to this parameterization as our *baseline*. Figure 2A shows that both wages and prices respond gradually to the increase in the money supply: a year after the shock, prices and wages increase by about one-half of a percent. In contrast, the marginal cost of production, $\Omega(s^t)$, is quite flexible and responds essentially one-for-one to the monetary shock, despite the wage stickiness. As equation (19) makes clear, the marginal cost rises because the marginal product of capital increases sharply after the shock, as well as because of the decreasing returns to scale in the

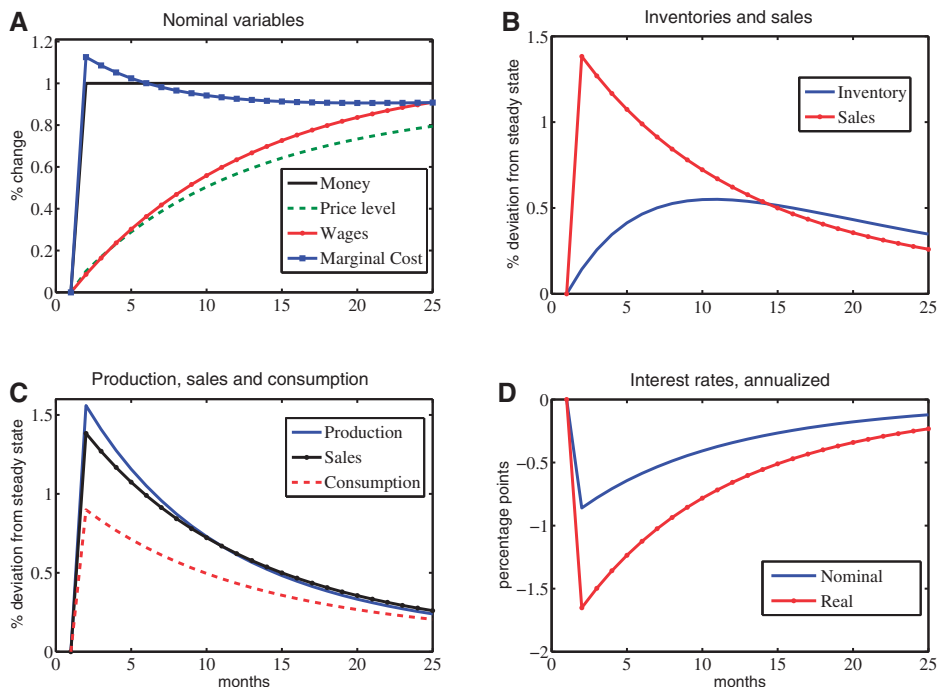


FIGURE 2

Impulse responses to monetary shock. Baseline model.

(A) Nominal variables, (B) Inventories and sales, (C) Production, sales and consumption, and (D) Interest rates, annualized

production of intermediate goods. Since the marginal cost is fairly flexible, but the price level is sticky, the average level of markups, as measured by $P(s^t)/\Omega(s^t)$, declines, thus reducing the distributor's incentives to hold inventories.

Figure 2B shows the responses of the stock of inventories and sales. We report, as in the data, the response of *real* sales and the aggregate end-of-period real inventory stock defined as $I(s^t) = \int_0^1 [z_i(s^t) - q_i(s^t)] di$. Notice that sales increase by about 1.4% and gradually decline, while the stock of inventories increases gradually and reaches its peak of about 0.6% about a year after the shock. Since inventories increase much more gradually than sales do, the inventory–sales ratio declines.

Figure 2C shows the response of production, sales, and aggregate consumption. Recall that production is equal to the sum of all orders made by distributors and is therefore equal to total sales plus inventory investment. Since production increases by about 1.6% in response to the monetary shock and sales by only 1.4%, the excess production contributes to an increase in inventory investment. This increase in inventory investment is small, however, relative to the total stock of inventories in the economy, and hence the increase in the inventory stock is gradual.

Finally, Figure 2D shows that both nominal and real interest rates decline in response to the monetary policy shock, as in the data.¹⁶ The nominal interest rate declines by less than the real rate does, because of expected inflation, but nevertheless declines because of the large drop

16. See the Supplementary Appendix for evidence on how consumption and interest rates respond to monetary shocks in the data.

TABLE 3
Business cycle predictions of the model

	(1) Data	(2) Baseline	(3) Labour only	(4) Flexible prices
<i>Impulse response of consumption to monetary shock</i>				
Average response	0.55	0.45	0.60	0.25
Maximum response	1.01	0.90	1.02	0.33
Half-life, months	16.20	10.5	13.1	26.2
Markup contribution		0.89	0.36	0.08
<i>Inventory statistics</i>				
$\rho(IS_t, S_t)$	-0.71	-0.91	0.99	0.52
$\rho(IS_t)/\sigma(S_t)$	0.93	0.82	3.17	1.15
Elast. IS_t to S_t	-0.66	-0.75	3.15	0.60
Elast. I_t to S_t	0.34	0.25	4.15	1.60
$\rho(IS_t, IS_{t-1})$	0.88	0.74	0.82	0.94
$\rho(Y_t)/\sigma(S_t)$	1.11	1.10	3.82	1.47
$\rho(Y_t, \Delta I_t)$	0.63	0.89	0.97	0.85
$\rho(\Delta I_t)/\sigma(Y_t)$	0.20	0.10	0.88	0.41
Elast. ΔI_t to Y_t	0.12	0.09	0.85	0.35
<i>Investment statistics</i>				
$\sigma(x_t)/\sigma(c_t)$	4	4	-	4

Notes: All variables are HP-filtered with smoothing parameter 14,400. The average consumption response is computed for first 24 months after shock.

in the real interest rates induced by the monetary policy expansion and price rigidities. One key assumption that we have made that ensures this drop in nominal interest rates is that the intertemporal elasticity of substitution is less than 1, so that $\sigma > 1$. To see this, note that the binding cash-in-advance constraint implies that the nominal interest rate in our setup is equal to

$$1 + i(s^t) = \left[\beta \int \frac{1}{g_m(s^{t+1})} \left(\frac{c(s^{t+1})}{c(s^t)} \right)^{1-\sigma} \pi(s^{t+1}|s^t) ds^{t+1} \right]^{-1}. \tag{37}$$

Since money growth, $g_m(s^t)$, is serially uncorrelated and consumption is expected to decline towards the steady state, the nominal interest rate declines as long as $\sigma > 1$.

In Table 3, we report the business cycle properties of the model and compare them to the data. We report two sets of statistics. The first set are measures of the real effects of money which summarize the impulse response of aggregate consumption, $c(s^t)$, to a monetary shock. Alternatively, since the cash-in-advance constraint, $\log c(s^t) = \log M(s^t) - \log P(s^t)$, binds in our model, these statistics summarize the degree of aggregate price stickiness: the extent to which prices trail the change in the money supply.

Column 2 of Table 3 shows that the average response of consumption in the first two years after the shock is equal to 0.45% in our baseline parameterization. The maximum consumption response is equal to 0.90%. The half-life of the consumption response, our measure of the persistence, is equal to 10.5 months. Comparison with the ‘‘Data’’ column shows that the model predicts somewhat smaller real effects of monetary policy shocks. In the data, the average consumption response to identified monetary shocks is about 0.55% and consumption is more persistent, with a half-life of about 16 months.¹⁷ Our results are thus consistent with those of Chari *et al.* (2000),

17. We compute these impulse responses in the data using the Christiano *et al.* (1999) measures of monetary policy shocks. See the Supplementary Appendix for details.

who find that sticky price models are unable, absent additional sources of rigidities, to account for the persistence of output and consumption in the data.

We next ask: what is the role of markup variation in accounting for the real effects of monetary shocks? As noted above, the effect of an increase in the money stock on consumption in our model is determined solely by the degree of aggregate price stickiness. The aggregate price level is sticky for two reasons: costs of production are sticky because of wage rigidity, and markups decline because of price rigidity. To decompose the role of these two effects, note that the cash-in-advance constraint implies that

$$\Delta \ln(c(s^t)) = \underbrace{\Delta [\ln(M(s^t)) - \ln(\Omega(s^t))]}_{\text{cost rigidity term}} + \underbrace{\Delta [\ln(\Omega(s^t)) - \ln(P(s^t))]}_{\text{markup term}}. \quad (38)$$

The response of consumption to a monetary shock is thus equal to the sum of two terms. The first term captures the extent to which costs, $\Omega(s^t)$, decline relative to the money stock. The second term captures the extent to which prices decline relative to costs. Table 3 reports the average response of the markup term relative to the average response of consumption. This ratio is equal to the area between the price and cost impulse responses relative to the area between the money and price responses in Figure 2 and is our measure of the fraction of the real effects accounted for by countercyclical markups. As the row labelled “markup contribution” in Table 3 shows, markups account for almost 90% of the response of consumption to the monetary shock.¹⁸

The second set of statistics we report are those that characterize the behaviour of inventories, sales, and production. To compute these statistics, we HP-filter simulated time series for these variables, as we have done in the data, with a smoothing parameter of 14,400. We then contrast the model’s predictions with those in the data for the manufacturing and trade sector. We focus on the statistics that are conditional on identified monetary policy shocks, although recall from Table 1 that the unconditional statistics are very similar.

The model does a very good job in accounting for the behaviour of inventories in the data. It predicts a countercyclical inventory–sales ratio—the correlation with sales is -0.91 in the model versus -0.71 in the data. The inventory–sales ratio is about 0.82 times more volatile than sales in the model (0.93 in the data). Consequently, the elasticity of the inventory–sales ratio to sales is equal to -0.75 in the model (-0.66 in the data). The elasticity of inventories with respect to sales is equal to 0.25, so that a 1% increase in sales is associated with an increase in the inventory stock of only 0.25%.

Also note that the inventory–sales ratio is somewhat less persistent than in the data. Its autocorrelation is equal to 0.74 in the model, compared to 0.88 in the data. This lack of persistence is partly due to the lack of persistence in sales in the model (the autocorrelation of sales is equal to 0.80 in the model and 0.89 in the data).

What accounts for the model’s implication that the inventory–sales ratio decreases in periods with expansionary monetary shocks? Recall that the inventory–sales ratio in this economy increases with the expected change in costs and the markup, and depends negatively on the nominal interest rate. Since costs increase immediately in response to an unanticipated increase in the money supply, the expected change in costs is small. If anything, as Figure 2 shows, costs are expected to decline in the first few periods after the shock as the economy accumulates more capital. The decline in nominal interest rates counteracts this effect, and the return to holding inventories is essentially constant. The drop in markups thus makes it optimal for firms to lower

18. A variance decomposition based on equation (38) using simulated data from the model yields a very similar result to that using impulse responses.

their stock of inventories relative to sales, since lower markups reduce the gains from holding inventories.

Consider next the model's predictions for the relative volatility of orders, sales and inventory investment. Note that production is 1.10 times more volatile than sales in the model, thus almost as much as in the data (1.11). As in the data, inventory investment is procyclical: the correlation between production and inventory investment is equal to 0.89 in the model and 0.63 in the data. Finally, note that inventory investment is about half as volatile in the model as in the data. The standard deviation of inventory investment (expressed as a fraction of output) is equal to 0.10 of the standard deviation of output in the model and 0.20 in the data. Together, these last two statistics imply that the elasticity of inventory investment to output is about 0.09 in the model and 0.12 in the data.

Why does the model produce too little volatility in inventory investment relative to the data? It turns out that the model fails to account for the *high-frequency* variability of inventory investment, although it does account well for the variability of inventory investment at *business cycle* frequencies. To see this, we also report statistics from the data that are filtered using the Baxter–King (1999) bandpass filter which isolates fluctuations in the data at frequencies between one and eight years and find that the resulting series for inventory investment is 0.12 times as volatile as output. When we apply an identical filter to the model, the standard deviation of inventory investment is equal to 0.10, thus not too different from the data. Hence, the failure of the model to account for the variability of inventory investment in the data is due to its failure to account for the volatile high-frequency component, documented by Hornstein (1998) and Wen (2005).

4.3. *The role of countercyclical markups*

We argue next that our model's ability to account for the dynamics of inventories in the data is largely accounted for by the presence of countercyclical markups. To see this, we next study two economies in which markups are much less responsive to monetary policy shocks.

First, consider what happens when we eliminate the decreasing returns to labour by eliminating capital ($\alpha = 1$) and the decreasing returns to scale ($\gamma = 1$). Figure 3 reports the impulse responses to a monetary shock in this economy, which we refer to as the *labour only* economy. Since labour is the only factor of production, the marginal cost is equal to the wage and responds gradually to the monetary shock. Even though prices are sticky, the markup declines much less than in our baseline parameterization, since costs do not increase much relative to prices.

Notice in Figure 3B that in this economy the stock of inventories increases by more than 4% after an expansionary monetary shock and that production increases by about 7%. Table 3 shows that the inventory–sales ratio is now strongly procyclical, with an elasticity of inventories to sales of about 4.2, a lot larger than in the data. Moreover, production is very volatile as well, about four times more volatile than sales. Because costs are sticky, markups decline much less now and account for only about one-third of the increase in consumption.

Why are inventories so volatile in this version of the model? The return to holding inventories in equation (34) sharply increases after an expansionary monetary shock, since the nominal interest rate declines and the cost of production is expected to increase. The higher return to holding inventories, combined with only a minor drop in markups, makes it optimal for firms to sharply increase their stock of inventories.

We next study the role of price rigidities by reverting to our original economy with decreasing returns to labour but assuming that prices are flexible. Since prices are flexible, markups barely change in response to the monetary shock: less than one-tenth of the increase in consumption is due to the decline in markups. Although inventories are now somewhat less volatile than in

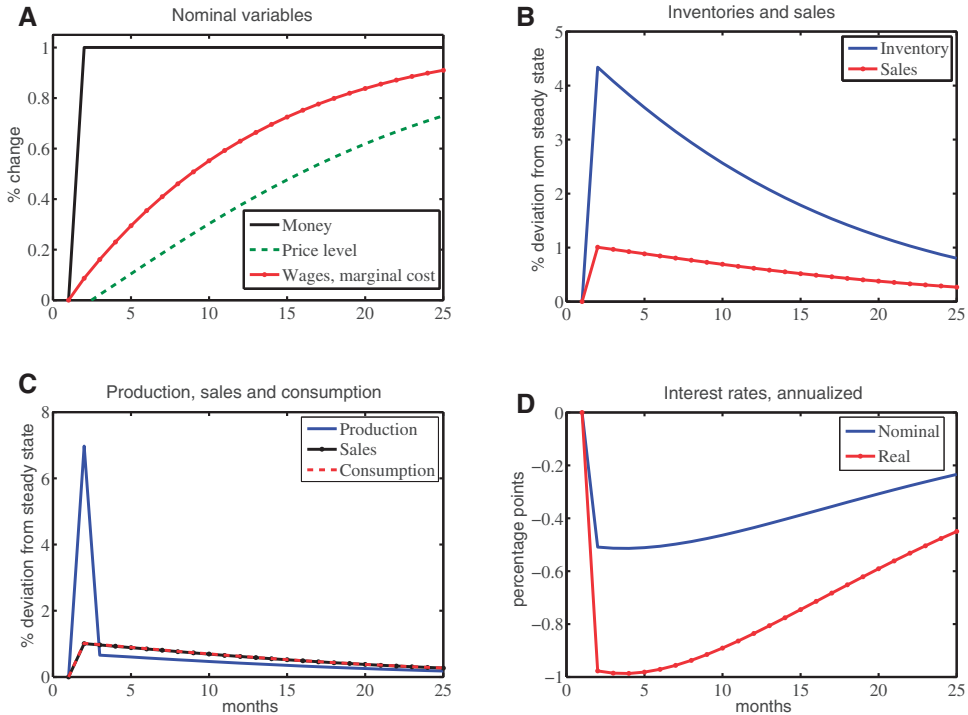


FIGURE 3

Impulse responses to monetary shock. Labor only.

(A) Nominal variables, (B) Inventories and sales, (C) Production, sales and consumption, and (D) Interest rates, annualized

the economy with labour only, they are nevertheless much more volatile than in the data. The inventory–sales ratio is procyclical: the stock of inventories increases by about 1.6% for every 1% increase in sales. Moreover, production is now about 47% more volatile than sales, much more so than in the data. Overall, this model once again cannot reproduce the key features of the inventory data.

Thus, contrary to what Khan and Thomas (2007) find for *productivity* shocks, countercyclical markups must play an important role for our model to account for the response of inventories to *monetary* shocks. This difference stems from the special nature of monetary shocks. Unlike productivity shocks, which affect output even in the absence of adjustment in capital or labour inputs, monetary shocks can affect output only if either markups adjust or if nominal costs are sticky. If costs are too sticky, firms intertemporally substitute production in anticipation of future cost changes, a feature that is at odds with the data. Hence, the model requires countercyclical markups to generate real effects from monetary shocks and at the same time be consistent with the inventory data.

4.4. Role of inventories for fluctuations

How are the model's predictions influenced by the presence of inventory accumulation? We answer this question by comparing our model's predictions to those of a standard New Keynesian

model without inventories, the limiting case of our model when σ_v goes to 0. We leave all parameters, including the size of the capital adjustment costs, unchanged.

We find that the model's predictions for the volatility of consumption and investment are essentially identical to those of the model with inventories. The average response of consumption to a monetary shock is identical to that in the model with inventories (0.45%), while investment is 4.01 times more volatile than consumption (four times more volatile in the model with inventories). Finally, output is about 8% more volatile in the model with inventories than in the model without inventories. Hence, as Khan and Thomas (2007) find in the context of an (S, s) model with productivity shocks, our model also implies that inventories play little role in amplifying business cycle fluctuations.

5. ROBUSTNESS

We now discuss several additional experiments that we have conducted. We first show that all of our results hold in an alternative popular model of inventories, the (S, s) model with fixed ordering costs. We then briefly report on our results from several experiments in which we modify our assumptions on technology, the process for monetary policy, the rate at which inventories depreciate, the nature of capital adjustment costs, as well as the stage of fabrication at which inventories are held. To conserve on space, we very briefly discuss these extensions and refer the reader to the Supplementary Appendix for a more in-depth discussion.

5.1. (S, s) inventory models

We next study an economy in which distributors face a fixed cost of holding inventories. Such a fixed cost makes it optimal for distributors to order infrequently and hold inventories. We show that our results are robust to this modification.

Each period a distributor faces a cost $\kappa_i(s^t)$ of ordering inventories. As in Khan and Thomas (2007), this fixed cost is an i.i.d. random variable drawn from a distribution that we specify below. The distributor's problem now becomes

$$\max_{P_i(s^t), y_i(s^t), \varphi_i(s^t)} \sum_{t=0}^{\infty} \int_{s^t} Q(s^t) (P_i(s^t) q_i(s^t) - \Omega(s^t) y_i(s^t) - W(s^t) \kappa_i(s^t) \varphi_i(s^t)) ds^t, \quad (39)$$

subject to the demand function in equation (21), the Calvo pricing restrictions, and the law of motion for inventories in equations (20) and (23). Here $\varphi_i(s^t) = 1$ if $y_i(s^t) \neq 0$ and $\varphi_i(s^t) = 0$ otherwise.

5.1.1. Parameterization and micro implications. We study two versions of the economy with fixed costs. In the first set of experiments, we assume that κ_i is uniformly distributed on the interval $[0, \bar{\kappa}]$. In the second experiment, we assume that κ_i is equal to 0 with probability λ_z and equal to a prohibitively large number with probability $1 - \lambda_z$. We refer to the second economy as the *Calvo orders* economy since the probability of ordering is constant and independent of the inventory stock.

Table 4 reports the parameter values that we have used for these two versions of the model. We reduce the standard deviation of demand shocks 3-fold, to $\sigma_v = 0.20$, in order to show that our

TABLE 4
Parameterization of (S,s) economies

		Data	Stockout-avoidance	Uniform order costs	Calvo orders
<i>Parameters</i>					
δ_z	Inventory depreciation		0.011	0.011	0.011
σ_v	s.d. demand shocks		0.626	0.20	0.20
ξ	Investm. adj. cost		46.2	46.2	46.2
κ	Upper bound on ordering cost		–	0.045	–
λ_z	Probability of ordering		–	–	0.73
<i>Targets</i>					
	I/S ratio	1.4	1.4	1.4	1.4
<i>Additional moments</i>					
	Frequency stockouts	0.05	0.05	0.045	0.068
	Frequency of orders	0.79	1	0.38	0.73
	s.d. ln(sales)	0.55	0.60	0.20	0.29
	s.d. Δ ln(sales)	0.77	0.84	0.28	0.40
	Autocor. sales	0.03	0.00	0.00	–0.03
	Corr(inv invest., sales)	0.06	–0.64	–0.11	–0.02

Note: The fixed ordering cost is expressed as a fraction of the distributor's mean order (conditional on ordering) in the ergodic steady state.

original results are not driven by the volatility of demand shocks.¹⁹ We leave all other parameters unchanged.

We choose, in the model with uniform ordering costs, an upper bound on ordering costs, $\bar{\kappa}$, equal to 4.5% of the distributor's steady-state orders, in order to generate an inventory–sales ratio of 1.4. We choose the constant hazard of ordering in the Calvo model equal to 0.73 so that the model also matches this ratio. We note that absent the fixed costs, the model would produce an inventory–sales ratio equal to 0.4 when the standard deviation of demand shocks is equal to 0.20. Hence, the stockout-avoidance motive accounts for about 30% of the stock of inventories held by distributors; the rest is accounted for by the fixed ordering costs.

Table 4 also reports several statistics from the data that we have not explicitly targeted in our calibration. Notice that both versions of the (S, s) model predict a frequency of stockouts that is not too different from the 5% in the data: 4.5% in the model with uniform ordering costs and 6.8% in the Calvo model. With Calvo ordering, stockouts occur more frequently because distributors cannot choose to order whenever the stock of inventories falls close to zero.

The next set of statistics we report concern the variability of sales, frequency of orders and the correlation of inventory investment with sales at the micro level. For a sense of how these statistics compare with the data, we use the Spanish supermarket data from Aguirregabiria (1999). This data set is a panel of monthly observations on inventories, sales and orders for 534 products (mostly non-perishable foods and household supplies) sold by the supermarket chain for a period of 29 months from 1990 to 1992. Although it is difficult to draw broader conclusions about the aggregate economy from a particular supermarket, we are not aware of other data sets that contain product-level information on inventories, sales and orders at the monthly frequency.

This particular supermarket orders individual products with a frequency of 0.79 per month. Sales of individual good are very volatile: the monthly standard deviation of log sales is equal to 0.55, while that of changes in log sales is equal to 0.77. There is very little persistence in sales

19. We do not eliminate demand shocks altogether because demand shocks smooth out the kink in the distributor's profit function and facilitate computations.

TABLE 5
Business cycle predictions of (S, s) economies

	Data		Model					
			(1) Uniform ordering costs			(2) Calvo ordering		
	Baseline	Labor only	Flex prices	Baseline	Labor only	Flex prices		
<i>Impulse response of consumption to monetary shock</i>								
Average response	0.55	0.43	0.56	0.23	0.42	0.57	0.25	
Maximum response	1.01	0.80	0.97	0.33	0.82	0.92	0.34	
Half-life, months	16.2	12.6	11.9	20.6	10.4	15.1	20.7	
Markup contribution		0.90	0.38	0.03	0.87	0.38	0.08	
<i>Inventory statistics</i>								
$\rho(I/S_t, S_t)$	-0.71	-0.89	0.99	0.45	-0.88	0.99	0.68	
$\sigma(IS_t)/\sigma(S_t)$	0.93	0.71	2.02	2.09	0.75	4.10	1.28	
Elast. I/S_t to S_t	-0.66	-0.63	2.00	0.94	-0.66	4.06	0.87	
Elast. I_t to S_t	0.34	0.37	3.00	1.94	0.34	5.06	1.87	
$\rho(IS_t, IS_{t-1})$	0.88	0.81	0.82	0.97	0.77	0.90	0.96	
$\sigma(Y_t)/\sigma(S_t)$	1.11	1.08	1.76	1.50	1.09	1.15	1.39	
$\rho(Y_t, \Delta I_t)$	0.63	0.61	0.83	0.84	0.89	0.50	0.83	
$\sigma(\Delta I_t)/\sigma(Y_t)$	0.20	0.12	0.71	0.45	0.10	0.54	0.37	
Elast. ΔI_t to Y_t	0.12	0.07	0.59	0.38	0.09	0.27	0.31	
<i>Investment statistics</i>								
$\sigma(x_t)/\sigma(c_t)$	4	3.95	-	3.98	3.95	-	3.93	

from one month to another: the autocorrelation of sales is equal to 0.03. Finally, the correlation between inventory investment and sales is very weak and equal to 0.06.

As Table 4 shows, the stockout-avoidance model produces similar variability of sales as in the data: the standard deviation of sales is equal to 0.60 compared to 0.55 in the data. That model cannot reproduce, however, the correlation between inventory investment and sales in the data. This correlation is strongly negative in the model (-0.64) and is therefore much at odds with the weak correlation of 0.06 in the data. Intuitively, since orders must be chosen prior to the realization of the demand shock, an unexpected positive demand shock increases sales and also leads to a reduction in the distributor's stock of inventories.

The (S, s) models do much better than the stockout-avoidance model along several dimensions. First, in both models distributors order infrequently: 0.38 of the months in the model with uniform costs and 0.73 of the months in the Calvo model (0.79 in the data). Even though Calvo firms order twice as frequently, the timing of those orders is exogenous and uncertain, leading them to hold inventories as a precaution against the possibility of not being able to order again soon. Second, the correlation between inventory investment and sales is much closer to the data: -0.11 in the model with uniform ordering costs and -0.02 in the Calvo model, although still negative. Intuitively, periods in which distributors order a new batch of inventories are periods with high inventory investment. Whether the distributor orders or not is largely determined by the size of the fixed cost it draws and the initial stock inventories, not by the amount it sells in that period.

5.1.2. Aggregate implications. Table 5 shows that the (S, s) model with capital, decreasing returns and sticky prices has very similar aggregate implications as our original model. For example, the elasticity of inventories to sales is equal to 0.37 in the model with uniform costs and 0.34 in the Calvo model, in line with the data (0.34), and only slightly greater than in the original model (0.25). Similarly, the elasticity of inventory investment to output is equal to 0.07 in the model with uniform costs and 0.09 in the Calvo model, very close to the 0.09 in the

original model. The model's implications for how consumption responds to monetary shocks and for the relative variability of investment and consumption are also very similar.

Assuming constant returns to labour visibly raises the variability of inventories in all models, although more so in the Calvo model in which firms order more frequently and are better able to take advantage of the temporarily low production costs. The elasticity of inventories to sales is equal to about 3 in the model with uniform ordering costs and about 5 in the Calvo model, compared to an elasticity of about 4 in the model without fixed costs of ordering.

Finally, eliminating price rigidities also produces very similar implications to those in the stockout-avoidance model. Since markups move little, firms invest strongly in inventories after an expansionary monetary shock. The elasticity of inventories to sales is equal to 1.94 in the model with uniform costs and 1.87 in the model with Calvo ordering, very close to the value of 1.60 in the stockout-avoidance model. The implications for all other variables are also very similar.

To summarize, our results are robust to the exact way in which we introduce a motive for holding inventories: both the stockout-avoidance and (S, s) models require that markups decrease in response to monetary shocks and costs increase rapidly in order to account for the inventory data. The reason for this result is that in both classes of models, the strength of the intertemporal substitution effect is primarily determined by a single parameter: the cost of carrying inventories, δ_z . As long as this cost is low, firms react strongly to changes in the return to carrying inventories regardless of the underlying reason for inventory accumulation.

5.2. *Extensions of the original model*

We consider next several perturbations of the assumptions on technology in our original stockout-avoidance model. All of these experiments are described in more detail in the Supplementary Appendix.

Our baseline model's results are robust to eliminating the decreasing returns at the firm level by setting $\gamma = 1$. The elasticity of inventories to sales is now equal to 0.49, thus only slightly greater than the elasticity of 0.34 in the data. In contrast, assuming increasing returns to scale, by setting $\gamma = 1.25$, implies that marginal costs increase only gradually in response to monetary shocks and the model predicts that the elasticity of inventories to sales is equal to 1.19, much greater than in the data.

Another extension we have considered is to introduce variable capital utilization in our baseline model. Such a modification is a popular approach to reducing the variability of marginal costs in the New Keynesian literature. We find, however, that such an extension once again implies that inventories are too volatile compared to the data, since marginal costs are not sufficiently responsive to monetary shocks.

We have also studied the role of wage rigidities. When wages are flexible, marginal costs overshoot after a, say, expansionary monetary shock, since wages increase sharply due to the greater disutility from work. Firms, anticipating future cost declines, sell out of their current stock of inventories and inventory investment decreases.

5.3. *Alternative specifications of monetary policy*

We have also studied economies with several alternative specifications of monetary policy. We first assumed that money growth is no longer i.i.d., but rather has a serial correlation of 0.61, a number often used in earlier work. We find that inventories become slightly countercyclical now (the elasticity with respect to sales is equal to -0.17), since nominal interest rates increase after a monetary expansion due to the higher expected inflation. The rise in nominal interest rates is

not enough to offset the expected increases in costs in the models with labour as the only factor of production or with flexible prices: both models continue to predict very volatile responses of inventories, though somewhat less volatile than in the economy with i.i.d. money growth.

Since nominal interest rates decline after an expansionary monetary shock in the data, we have also considered an extension of the model with persistent money growth in which consumers have habit persistence in preferences. Habit persistence implies that nominal interest rates fall even in response to persistent expansionary money shocks and the model's implications are now much closer to those of the model with i.i.d. money growth. Our baseline model with decreasing returns and sticky prices now produces an elasticity of inventories to sales of 0.18, very similar to the 0.25 value in the model with i.i.d. money growth.

We have also studied an economy in which monetary policy follows a Taylor rule. When we estimate the parameters of this rule using historical U.S. data, we find that the model's inventory implications are very similar to those in the original model with i.i.d. money growth.

5.4. *Inventory depreciation rate*

In our original experiments, we have set the rate at which inventories depreciate, δ_z , equal to 1.1%. The logistics literature reports inventory carrying costs that are somewhat higher, in the range of 1.5% to 3.5%.²⁰ We ask next whether our results are robust to assuming a greater rate of inventory depreciation.

We first increase δ_z to 2.5%, a number in the mid-range of those reported by Richardson (1995). The elasticity of inventories to sales is now equal to 0.10 in the baseline model with decreasing returns, thus slightly lower than the 0.25 in the original parameterization. In contrast, the models with labour as the only factor of production or with flexible prices strongly overstate the volatility in the data. The elasticity of inventories to sales is equal to 2.75 in the model with labour only and 1.44 in the model with flexible prices, both much higher than in the data. Raising the depreciation rate even further, to 3.5%, does not change these counterfactual implications very much.

Our results are also robust to allowing the inventory carrying cost to be a convex, rather than linear, function of the inventory stock. For example, when we consider a quadratic specification of the inventory carrying cost such that the depreciation rate is 0 when the stock is 0 and increasing at a rate necessary to match the steady-state frequency of stockouts and the inventory–sales ratio, we find that our baseline model with decreasing returns and sticky prices produces an elasticity of inventories to sales of 0.15. In contrast, the model with labour only produces a much greater elasticity of 3.62, and the model with sticky prices produces an elasticity of 1.38. The cubic specification of the carrying cost function produces similar results.

5.5. *Lower inventory–sales ratio*

One concern about our original calibration of a 1.4 inventory–sales ratio is that the stock of inventories in the data reflects stocks of final goods as well as stocks of intermediate goods. Since in our model firms only hold inventories of final goods, the concern is that our choice of an inventory–sales ratio of 1.4 is too high relative to the data.²¹

20. Richardson (1995) reports annual inventory carrying costs (excluding the “cost of money,” which is already accounted for in our model) that range from 19% to 43%, implying monthly carrying costs around 1.5–3.5%.

21. Note, however, that the inventory–sales ratio is also equal to 1.4 in the retail sector, which only holds inventories of finished goods.

Here, we ask whether our results are robust to reducing the volatility of demand shocks to $\sigma_v = 0.356$ so that the model matches an inventory–sales ratio of 0.7, half of that in our original setup. We find that the model’s implications for the response of inventories to monetary shocks change little. In the baseline parameterization the elasticity of inventories to sales is equal to 0.37, only slightly greater than in the model with a higher inventory–sales ratio. As in the original experiment, the models with labour only and with flexible prices produce a much more volatile stock of inventories.

5.6. *Productivity shocks*

An economy driven solely by productivity shocks and in which wages and prices are flexible reproduces the key features of the inventory data very well. Similarly, when we introduce productivity shocks into a model with sticky prices and wages and monetary policy shocks, the model predicts an elasticity of inventory to sales of 0.33 and a relative variability of inventory investment of 0.27, both very close to the data. In contrast, the model with sticky prices and wages in which productivity shocks are the only source of aggregate uncertainty, predicts that markups are strongly procyclical and that inventories are much more volatile than in the data.

5.7. *Role of capital adjustment costs*

Absent capital adjustment costs, our baseline model predicts that inventories are countercyclical, since the large spike in investment after a monetary expansion leads to an increase in interest rates. Such a model predicts, however, a relative volatility of investment to consumption of 150, much greater than in the data. We show in the Supplementary Appendix that this high variability of investment absent adjustment costs is not specific to our models with inventories. Standard New Keynesian models have very similar counterfactual implications for investment and interest rates. Our Supplementary Appendix discusses several alternative approaches to reducing the variability of investment in the model, as well as some plant-level evidence on the nature of capital adjustment costs.

5.8. *Inventories at two stages of production*

Our analysis has focused exclusively on finished goods inventories held by distributors. Inventories are held, however, at all stages of production. We show in the Supplementary Appendix that inventories of intermediate inputs account for about two-thirds of all inventories in the Manufacturing sector. However, since the wholesale and retail sectors hold large stocks of inventories, the share of inventories of intermediate inputs in the total stock of inventories in the U.S. manufacturing and trade sectors is equal to only 23%.

We then study an economy in which distributors hold inventories of finished goods and intermediate good producers hold inventories of materials. This model does a good job at reproducing the inventory statistics in the data, suggesting that our results are robust to introducing inventories at multiple stages of production.

5.9. *Lower share of inventory goods in final goods production*

Our Benchmark model assumes that final goods are produced solely using inputs of intermediate goods that can be stored as inventories. This feature, commonly used in the New Keynesian literature, contrasts with the assumption made in the work of Wen (2011) and Khan and Thomas (2007), in which the share of goods held in inventory in the production of final goods is equal to

0.7 and 0.5, respectively. We show in the Supplementary Appendix that our results are robust to reducing the share of storable goods in the production of final goods. In particular, we assume that the technology for producing final goods is

$$c(s^t) + x(s^t) + \phi(s^t) = a(s^t) = q(s^t)^\varphi \left[l_F(s^t)^\alpha k_F(s^t)^{1-\alpha} \right]^{1-\varphi},$$

where $q(s^t)$ is, as earlier, a Dixit-Stiglitz aggregator over varieties of storable intermediate inputs, while k_F and l_F are the amounts of capital and labour used in the final goods sector. We set φ , the share of storable goods in the production of final goods, equal to 1/2.

Our Supplementary Appendix shows that our results change very little in this alternative parameterization of the model. Once again, the baseline model with strongly countercyclical markups accounts well for the dynamics of inventories in the data. In contrast, versions of the model with flexible prices or constant returns to labour predict that inventories are much more volatile than in the data.

6. CONCLUSIONS

We embed a motive for inventory accumulation in a standard New Keynesian model with price and wage rigidities. The model predicts a strong relationship between inventories and the dynamics of costs and markups. We use the theory, together with data on inventories, to evaluate the role of cost rigidities and markups in accounting for the real effects of monetary policy shocks. In the data inventories adjust slowly in response to monetary shocks and are much less volatile than sales. Our theory interprets this fact as implying that countercyclical markups account for a sizable fraction of the real effects of monetary shocks.

Acknowledgments. The views expressed herein are those of the authors and not necessarily those of the Bank of Canada, the Federal Reserve Bank of Minneapolis or the Federal Reserve System. We thank three anonymous referees and our editor, Francesco Caselli, for numerous valuable suggestions. We are also indebted to George Alessandria, Ariel Burstein, Mike Dotsey, Mark Gertler, Boyan Jovanovic, Huntley Schaller, Jon Willis, and seminar participants at a number of venues for comments and suggestions.

SUPPLEMENTARY DATA

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

- AGUIRREGABIRIA, V. (1999), "The Dynamics of Markups and Inventories in Retailing Firms", *Review of Economic Studies*, **66**, 275–308.
- BALL, L. and ROMER, D. (1990), "Real Rigidities and the Non-neutrality of Money", *Review of Economic Studies*, **57**, 183–203.
- BASU, S. and FERNALD, J. (1997), "Returns to Scale in U.S. Production: Estimates and Implications", *Journal of Political Economy*, **105**, 249–283.
- BAXTER, M. and KING, R.G. (1999), "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series", *Review of Economics and Statistics*, **81**, 575–593.
- BILS, M. (2004), "Studying Price Markups from Stockout Behavior", (mimeo).
- BILS, M. and KAHN, J.A. (2000), "What Inventory Behavior Tells Us About Business Cycles", *American Economic Review*, **90**, 458–481.
- BILS, M. and KLENOW, P.J. (2004), "Some Evidence on the Importance of Sticky Prices", *Journal of Political Economy*, **112**, 947–985.
- BURSTEIN, A. and HELLOWIG, C. (2007), "Prices and Market Shares in a Menu Cost Model", (NBER Working Paper 13455).
- CHANG, Y., HORNSTEIN, A. and SARTE, P.-D.G. (2006), "Understanding How Employment Responds to Productivity Shocks in a Model with Inventories", (Working Paper 06-6, Federal Reserve Bank of Richmond).

- CHARI, V.V., KEHOE, P.J. and McGRATTAN, E.R. (2000), "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?" *Econometrica*, **68**, 1151–1180.
- CHRISTIANO, L.J., EICHENBAUM, M. and EVANS, C.L. (1999), "Monetary Policy Shocks: What Have We Learned and to What End?", in J. B. Taylor and M. Woodford (ed.) *Handbook of Macroeconomics*, vol. 1, 65–148. (Elsevier).
- CHRISTIANO, L.J., EICHENBAUM, M. and EVANS, C.L. (2005), "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy", *Journal of Political Economy*, **113**, 1–45.
- DOTSEY, M. and KING, R.G. (2006), "Pricing, Production, and Persistence", *Journal of the European Economic Association*, **4**, 893–928.
- EICHENBAUM, M., Jaimovich, N. and REBELO, S. (2011), "Reference Prices, Costs, and Nominal Rigidities", *American Economic Review*, **101**, 234–262.
- GALÍ, J. (2008), *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*. (Princeton University Press, Princeton, New Jersey).
- GERTLER, M. and LEAHY, J. (2008), "A Phillips Curve with an Ss Foundation", *Journal of Political Economy*, **116**, 533–572.
- GOLDBERG, P.K. and HELLERSTEIN, R. (2012), "A Structural Approach to Identifying the Sources of Local Currency Price Stability", *Review of Economic Studies*, **20**, 1–36.
- GOLOSOV, M. and LUCAS R.E. Jr., (2007), "Menu Costs and Phillips Curves", *Journal of Political Economy*, **115**, 171–199.
- HALL, R. (2009), "By How Much Does GDP Rise If the Government Buys More Output?", *Brookings Papers on Economic Activity*, **2**, 183–231.
- HOUSE, C.L. (2008), "Fixed Costs and Long-Lived Investments", (Working Paper 14402, NBER).
- HORNSTEIN, A. (1998), "Inventory Investment and the Business Cycle", Federal Reserve Bank of Richmond *Economic Quarterly*, **84**, 49–71.
- IACOVIELLO, M., Schiantarelli, F. and Schuh, S. (2007), "Input and Output Inventories in General Equilibrium", (Working Paper No. 07-16, Federal Reserve Bank of Boston).
- JUNG, Y. and YUN, T. (2005), "Monetary Policy Shocks, Inventory Dynamics, and Price-Setting Behavior", (Working Paper 2006-02, Federal Reserve Bank of San Francisco).
- KEHOE, P.J. and MIDRIGAN, V. (2010), "Prices are Sticky After All", (Working Papers 16364, NBER).
- KHAN, A. and THOMAS, J. (2007), "Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies", *American Economic Review*, **97**, 1165–1188.
- KHAN, A. and THOMAS, J.K. (2008), "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics", *Econometrica*, **76**, 395–436.
- KLENOW, P.J. and KRYVTSOV, O. (2008), "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics*, **123**, 863–904.
- KLENOW, P.J. and WILLIS, J. (2006), "Real Rigidities and Nominal Price Changes", (mimeo).
- NAKAMURA, E. and STEINSSON, J. (2008), "Five Facts about Prices: A Reevaluation of Menu Cost Models", *Quarterly Journal of Economics*, **123**, 1415–1464.
- NEKARDA, C.J. and RAMEY, V.A. (2010), "The Cyclical Behavior of the Price-Cost Markup", (mimeo).
- RAMEY, V.A. and WEST, K.D. (1999), "Inventories", in Taylor, J.B. and Woodford, M. (eds). *Handbook of Macroeconomics*, vol. 1, 863–923, (Elsevier).
- RIBARSKY, J. (2004), "Real Inventories, Sales, and Inventory-Sales Ratios for Manufacturing and Trade, 1997:I–2003:III", *Survey of Current Business*, **84**, 178–191.
- RICHARDSON, H. (1995), "Control Your Costs Then Cut Them", *Transportation and Distribution*, **December**, 94–96.
- ROGERSON, R. and WALLENIUS, J. (2009), "Retirement in a Life Cycle Model of Labor Supply with Home Production", (Working Paper wp205, University of Michigan).
- ROMER, C.D. and ROMER, D.H. (2004), "A New Measure of Monetary Shocks: Derivation and Implications", *American Economic Review*, **94**, 1055–1084.
- ROTEMBERG, J.J. and WOODFORD, M. (1999), "The Cyclical Behavior of Prices and Costs", in Taylor, J.B. and Woodford, M. (eds). *Handbook of Macroeconomics*, vol. 1, 1051–1135, (Elsevier).
- SMETS, F. and WOUTERS, R. (2007), "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach", *American Economic Review*, **97**, 586–606.
- WEN, Y. (2005), "Understanding the Inventory Cycle", *Journal of Monetary Economics*, **52**, 1533–1555.
- WEN, Y. (2011), "Input and Output Inventory Dynamics", *American Economic Journal: Macroeconomics*, **3**, 181–212.
- WOODFORD, M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, (Princeton University Press, Princeton, New Jersey).