

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

Oleksiy Kryvtsov
Bank of Canada

Virgiliu Midrigan[†]
New York University

April 2008

Abstract

Real rigidities that limit the responsiveness of real marginal cost to output fluctuations have been argued critical to the quantitative performance of sticky price models. We argue, in the spirit of Bils and Kahn (2000), that the elasticity of real marginal cost to output can be inferred from the behavior of inventories over the cycle. We show that a model which combines the stockout-avoidance and (s, S) motives for holding inventories, calibrated to match salient features of the microeconomic data, requires an elasticity of real marginal cost with respect to output of slightly less than unity (0.92), in order to account for the observed decline in inventory-sales ratios in the aftermath of monetary policy shocks.

JEL classifications: E31, F12.

Keywords: Inventory, menu costs, real rigidities.

*The views expressed here are those of the authors and do not reflect those of the Bank of Canada

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

1. Introduction

Real economic activity responds to changes in the stance of monetary policy with long lags. Identified exogenous changes in monetary policy lead to output, consumption and investment changes that peak after at least one and a half years after the shock; these real effects dissipate only after at least three years¹. Models with nominal price rigidities can account for these long lags only in the presence of strong degrees of strategic complementarity (real rigidities) in price setting. These real rigidities make it optimal for those firms that change prices to only partially respond to the monetary disturbances and thus keep prices close to those of their competitors that have not had a chance to reset. Absent real rigidities these models predict that the effect of monetary disturbances die out completely two to three quarters after the shock. This number is the average duration of price stickiness in the US data according to recent micro-studies of prices, and thus the time it takes for all firms to reset their prices after the shock².

Two types of real rigidities have received widespread attention in recent work³. The first type emphasizes strategic complementarities in price setting arising from losses to a particular firm of having its price deviate from that of its competitors (increased concavity of a firm's profit function in the relative price space). These complementarities can arise either due to non-constant demand elasticities, as in economies with (smoothed) kinked demand curves suggested by Kimball (1995). Similarly, specific factors of production or any other source of upward-sloping marginal cost curves at the firm level, would also make it optimal for firms to keep prices close to those of their competitors.

A second type of real rigidities stresses slow responsiveness of (aggregate) real factor prices and thus real marginal cost in response to output fluctuations. Simple variants of sticky price models predict an immediate increase in real wages, rental rates etc. in times of output expansions. As a result adjusting firms find it optimal to raise prices sharply in times of

¹Christiano, Eichenbaum and Evans (2005), Romer and Romer (2004), Friedman (1968).

²Bils and Klenow (2005), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2007).

³Ball and Romer (1990).

monetary expansions and render the real effects of money small and short-lived. Allowing for variations of the model in which real marginal cost responds slowly to shocks can thus render the impulse responses in sticky price monetary models closer to the data. Sticky nominal wage, variable labor and capital utilization, use of intermediate factors of production (whose prices are sticky) are a few examples that have been stressed important in recent work⁴.

An important question in light of the role real rigidities play in models with nominal rigidities is: how large are real rigidities in the data? Are the degrees of real rigidities necessary to reconcile sticky price models with the data empirically plausible?

One can measure real rigidities emphasizing the curvature of the profit function (e.g., kinked demand curves) by recognizing that these models predict that the larger this curvature is, the less will a given firm allow its price deviate from that of its competitors. This is the strategy pursued by Klenow-Willis (2006) who argue that the degree of curvature implied by typical parameterizations of the Kimball (1995) formulation of kinked demand curves generates implausibly small price changes as firms are unwilling to respond to idiosyncratic, say, cost disturbances for fear of losing (not gaining) customers. Similarly, Dotsey and King (2005) show that the large curvature of the profit function implied by this type of real rigidities require implausibly large menu costs (5.5% of revenue) in order for one to be able to sustain the duration of price contracts observed in the data. In principle it may be possible to reconcile real rigidities emphasizing high curvature in the profit function with the observed frequency and magnitude of price changes by recognizing that the relevant price affecting, say, demand elasticities or the marginal cost, is the relative price of the firm rather than that of any given product a firm sells.⁵ Nevertheless, existing single-product sticky price models emphasizing this first type of real rigidities are at odds with important features of the micro-data. The conclusion that the micro-economic evidence for this first type of real rigidities is weak is reinforced by recent work by Burstein and Hellwig (2007) who find that the degree of decreasing returns inferred using scanner price data in grocery stores is too weak to amplify

⁴Dotsey and King (2002), Christiano, Eichenbaum and Evans (2002), Nakamura and Steinsson (2007).

⁵Point due to Miles Kimball.

the real effects of monetary policy considerably.

Real rigidities emphasizing sluggish adjustment of nominal factor prices thus appear critical to the success of sticky-price-based explanations of the observed real effects of monetary policy. Measuring how sticky factor prices are is difficult in practice. For one, compositional issue plague measurement as documented by Solon, Barsky and Parker (1994) for real wages. Moreover, average, rather than marginal factor prices are typically observed. Finally, the assumption that factor inputs are demand determined may be invalid in the data, either because of long-term relationships between buyers and suppliers of factors of production, or because suppliers may be unable/unwilling to supply all that the buyer demands.

In this paper we pursue an alternative approach to measuring how sticky factor prices are in the data. We propose to use the observed behavior of inventories, rather than (in addition to) factor prices to measure how the real marginal cost of producing goods varies with the cycle (in response to monetary shocks).

Our approach is inspired by the work of Bils and Khan (2000) who argue that the behaviour of inventories over the cycle can provide valuable information about the cyclicity of real marginal costs and markups. In particular, in a sticky price model with sticky factor prices (in our model wages) several forces primarily determine a firm's incentives to invest in inventories in times of say, a monetary expansion. Countercyclical markups arising from nominal price stickiness render the firms' losses from stockouts smaller and decrease the desired inventory-sales ratio. Sticky factor prices (and thus expected future increases in the cost of acquiring inventories) increase the desired inventory-sales ratio by making firms better off buying low today rather than high tomorrow and also by reducing the extent to which nominal price stickiness translates into countercyclical markups. Finally, to the extent to which the monetary expansion is unanticipated, the actual inventory-sales ratio may be reduced if firms order inventories infrequently, as in the data, because of economies of scale in ordering inventories. Our goal is to quantify the relative strength of each of these forces using a model of sticky prices generalized to allow for an inventory-holding motive. We use the model to ask: what elasticity of real factor prices to output can allow the model to reproduce

the dynamic response of the inventory-to-sales ratio to a monetary shock we observe in the data?

We start by documenting that monetary expansions are associated with increases in the stock of inventories and declines in the inventory-sales ratio: the elasticity of inventory-sales ratios to sales is equal to -0.8. This is similar to what of Bils-Khan (2000) find for business cycles in general: although inventories are procyclical, inventory-sales ratios are negatively correlated with output at business cycle frequencies. Moreover, we document, using detailed data from the NBER productivity database and the Bureau of Economic Analysis, that this decline is not simply a compositional artifact resulting from industries that have on average lower inventory-to-sales ratios expanding more during booms.

We then formulate and calibrate a model in which nominal prices change infrequently because of fixed costs of price adjustment firms face and in which firms hold inventories because of two frictions on the technology of ordering goods we assume. In particular, we assume that ordering a new batch of inventories entails payment of an ordering cost that is independent of the shipment's size. Furthermore, the firm is assumed to order before a taste shock that determines the consumers' demand for its goods is revealed. As a result of these two frictions the firm holds excess inventories in order to minimize the probability of a stockout and the ordering costs. We show that the model, though parsimoniously parameterized, accords well with the frequency of orders, stockouts, as well as the average inventory-to-sales ratio observed in the data.

We then study the response of our economy to monetary policy shocks and ask: what is the elasticity of real marginal cost with respect to output that the model requires to match the elasticity of inventory-sales ratio to sales of -0.8 we document in the data. We vary the responsiveness of real marginal cost to output in the model by introducing a wedge in the optimal labor-leisure choice that allows us to mimic a range of responses for the real wage without modeling the frictions that account for these responses explicitly. We show that the parameterization in which nominal wages increase one for one with the monetary shock (and, given our assumption on preferences, real marginal cost increases one-for-one with output)

predicts, counterfactually, a decline in both inventories, in addition to the inventory-sales ratio. In contrast, the stickier nominal wages are, the more inventories and inventory-sales ratios rise in response to the monetary expansion. We then show that very modest degrees of nominal wage stickiness are necessary in our model to reconcile the behavior of inventories in the wake of a monetary shock in the data. The implied elasticity of real marginal cost that is necessary to account for the response of the inventory-sales ratio in the data is equal to 0.92, only slightly below unity⁶. In contrast, the range of elasticities used in earlier work range from 3 (the baseline parameterization used by Chari, Kehoe, McGrattan (2002)), to 0.33 (Dotsey and King (2005)).

This paper proceeds as follows. Section 2 documents a number of facts about inventories and computes the response of inventory-to-sales ratio to measures of exogenous monetary policy shocks. It then shows that the decline in inventory-to-sales ratio is pervasive across all industries, and not simply a compositional artifact. Section 3 presents the model and discusses optimal decision rules in this environment. Section 4 computes impulse response to monetary policy shocks under several assumptions regarding the elasticity of real marginal cost with respect to output (i.e, stickiness of nominal wages) and asks what elasticity is consistent with the decline in the inventory-to-sales ratio observed in the data. Section 5 concludes.

2. Data

In this section we document several salient facts regarding the cyclical behavior of inventories. The ratio is central to assessing the dynamic properties of markups, marginal cost, and in turn, the size of real rigidities over the business cycle. We employ two datasets: the (annual) NBER Manufacturing Productivity Database with data on output, sales, inventories and price deflators in 459 4-digit manufacturing industries from 1957 to 1996; and the Bureau

⁶Khan and Thomas (2007) and Wen (2008) show that a real business cycle model with capital modified to allow for fixed cost and stockout-avoidance motives for holding inventories, respectively, predicts counter-cyclical inventory-to-sales ratios and procyclical inventory levels. Their results are consistent with ours: in their models, investment in capital during booms drives up the real marginal cost of acquiring inventories.

of Economic Analysis (NIPA) data on monthly final sales, inventories, and inventory-to-sales ratios for Manufacturing and Trade sectors from January 1967 to December 1997. These data sets complement each other as they allows us to study inventory behaviour for a range of business cycle frequencies - monthly and annual; and for a different levels of aggregation - industry-level and aggregate.

Sales are defined as real final sales to domestic purchases in NIPA data and as the real value of industry shipments in case of the NBER data. Output is the sum of final sales and the change in the end-of-period inventory stock. Inventory-to-sales ratio is defined as the ratio of the end-of-period inventory stock to final sales in that period. Shipments deflators are used for industry-level data and CPI (less food and energy) is used for aggregate data to deflate nominal variables. All data are HP filtered.⁷ Output, sales and inventory-sale ratios are defined in % deviations from respective HP trends. Inventory investment is defined in % points-of-output-fraction deviations from its HP trend. Unless otherwise noted, inventory-to-sales ratios, output, sales, etc. are in real units. Detailed description of the data is provided in the Appendix.

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

⁷There is a great deal of heterogeneity in industry-specific trends of sales and inventory-to-sales ratios. Bils and Kahn (2000) argue that scale effects are small for inventory-to-sales ratio dynamics. Our work below shows that heterogeneity across inventory-to-sales ratios is unlikely to be an important factor for the dynamics of the aggregate inventory-to-sales ratio.

A. Inventories and sales over the business cycle

This section describes the behavior of inventories over the cycle⁸. Table 1 reports several unconditional moments for real inventory-to-sales ratios in the data. On average firms carry inventory stocks of about 1.4 months of sales according to the monthly data and 0.23 years of sales in the annual data (that is, 2.8 months of sales). Standard deviations of real inventory-to-sales ratios, 2.2% in the monthly and 3.9% in the annual data, are comparable to those of the aggregate output and sales. The most prominent feature of the aggregate inventory-to-sales ratios is their countercyclicality. The correlation of real inventory-to-sales ratios with output in the monthly data is -0.83 for Manufacturing and Trade, and -0.49 for Retail. In the annual data that correlation is -0.52. Correlations with sales are typically even more negative.

Figure 1 plots the aggregate (detrended) inventory-to-sales ratio, real GDP as well as a measure of real interest rates.⁹ As documented by King and Watson (1994), the real interest rate is negatively correlated with current and future output. In our annual data the contemporaneous correlation of the real interest rate with output is -0.43, while with one-year-lead (lag) output is -0.37 (-0.26). Since real interest rates affect the opportunity cost of holding inventories, their low value in the wake of the output boom should increase inventory accumulation absent additional channels.

B. Compositional bias

There is a considerable amount of heterogeneity in the level of stock relative to sales among firms. For example, in 1996 25th, 50th and 75th-percentile levels of inventory-to-sales ratios across 4-digit NBER industries are 0.23, 0.31 and 0.42 respectively. One concern is therefore that the behavior of inventory-to-sales ratios may reflect a compositional shift. Given that the aggregate inventory-to-sales ratio is equal to the average of industry inventory-to-sales ratios (weighted by each industry's sales), it may be the case that the drop of the

⁸Ramey and West (1999) and Bils and Kahn (2000) are two earlier papers that document similar facts.

⁹Real interest rates are defined as 1-year treasury bill rates minus realization of the backward-looking 1-year inflation rate. See Horstein and Hellwig (1999) for a discussion of alternative definitions of the real interest rate.

aggregate ratio is simply evidence that low-inventory-to-sales industries expand more during booms.¹⁰

To assess this hypothesis we divide the annual data into 3 bins according to the level of inventory-to-sales ratios: low, medium, and high. Table 2 reports, for each bin, the average inventory-to-sales ratio and the average share of sales in total sales over the sample. Inventory-to-sales ratios from the lower third average about 0.18 and account for about a half of sales in the sample. Medium and high thirds have mean inventory-to-sales ratios of 0.31 and 0.55 respectively, and account for a quarter of total sales each.

For compositional effects to be important there must be large differences in how inventory-to-sales ratios or sales shares for each bin fluctuate over the business cycle. Table 2 demonstrates that these differences, if any, are small. Specifically correlations of bin-average inventory-to-sales ratios with real GDP are very similar, -0.41, -0.33 and -0.47 for low, medium and high bins respectively. Correlations of bin sales shares with real GDP are small and positive, reflecting the fact sales for each bin are higher in booms - correlation of annual sales shares with real GDP is 0.13. Sales-share-output correlation for low bin is only slightly higher than that for high bin, 0.12 vs 0.03. This shift of sales shares towards low inventory-to-sales ratios during booms is small - the resulting correlation of the aggregate inventory-to-sales ratio with real GDP is -0.47, which is close to bin correlations.

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

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

where $\ln IS_{it}$ is the log inventory-to-sales ratio for industry i , \mathbf{D}_t is the vector of year dummies,

¹⁰Solon, Barsky, Parker (1994) find similar composition bias for aggregate real wages that are acyclical in the aggregate time series because the aggregate statistics are constructed in a way that gives more weight to low-skill workers during expansions than during recessions.

D_i are 4-digit industry dummies, and ε_{it} is the residual representing industry-level disturbances. Coefficients for year dummies, α_{year} , represent the fixed-weight average inventory-to-sales ratio. If compositional effects are important, the time series for the average and aggregate inventory-to-sales ratios should differ substantially. Figure 2 shows that this is not the case. Indeed the two series move together over the cycle, except for the 1994-1997 period when the aggregate ratio was somewhat smaller than the average. We conclude that compositional shifts across sectors with high and low inventory-to-sales ratios do not contribute importantly to countercyclicality of the aggregate inventory-to-sales ratio.

C. Elasticities of inventory-to-sales ratios

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

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

At the industry level across-time elasticities turn out to be of the same magnitude. Table 4 shows that across-time sales elasticity of inventory-to-sales ratio at the industry level is around -0.6. When we compare the across-time sales elasticity of inventory-to-sales ratio to its cross-section counterpart, we see that the cross-section elasticity is much smaller than the one across time: -0.14. These estimates of industry-level across-time and cross-sectional elasticities are not sensitive to controlling for various measures of industry inflation, marginal cost, markup, and firm concentration. Together, these correlations suggests that the negative correlation between inventory-sales ratios and sales is a short-run, rather than long-run phenomenon. This result is reminiscent of that of Bils and Kahn (2000) who find for six large manufacturing industries that inventory-to-sales ratio dynamics is insensitive to the industry size. The across-time elasticities of -.60 to -.86 of the inventory-to-sales ratio

are thus mostly due to short-run fluctuations.

D. Evidence from identified monetary shocks

Since the focus of this paper is to use inventory behaviour to gauge the role of real rigidities in accounting for the real effects of monetary shocks, we extend our data analysis to document the behavior of inventory-to-sales ratios in the wake of monetary expansions and contractions. A necessary step here is to identify these monetary disturbances. The estimation is based on monthly NIPA which has more time-series variation.

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

Responses of output, inventory investment and inventory-to-sales ratios for each sector are obtained by estimating the following OLS regression:

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

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

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

Estimation of (1) yields a number of facts about impulse responses after monetary shocks.¹¹ First, the estimated federal funds rate increases sharply up to about 1.2 % points within one quarter after the shock and comes back to zero slowly, after 1 and a half years,

¹¹The results are provided for estimates based on CEE measure of monetary shocks. The results based for RR measure are very similar.

see Figure 3.

Figure 4 reports impulse responses to CEE shocks for output, sales, inventory investment and inventory-to-sales ratio in Manufacturing and Trade. Output responses are positive for about half a year after the shock, reflecting the fact that monetary shocks affect the economy with a lag. Output responses are negative between around 6 and 30 months, with the trough at around 24 months after the shock. Sales responses are very similar to output responses. The responses of inventory investment is small and statistically almost indistinguishable from zero. Finally, the response of inventory-to-sales ratios is negative for the first 5 months after the shock and then positive up until 2.5 years after the shock. For Retail, inventory-to-sales ratios are somewhat quicker to respond and less persistent – they are around zero in the first few months after the shock, then positive coming back to zero after 18 months after the shock (versus around 30 months for the Manufacturing and Trade sector).

We back the above estimation with our own simple VAR for the Manufacturing and Trade sector. The VAR includes the following variables: output, inflation (CPI less food and energy), sales, inventory-to-sales ratio, and federal funds rate, Choleski-ordered in that order. Figure 5 provides impulse responses to one-standard-deviation positive innovation to the federal funds rate. Responses are very similar to those estimated for CEE and RR monetary shocks. In particular, within half a year after the shock output and sales responses are positive and inventory-to-sales ratio is negative. After that for at least 2 years output and sales are negative, and inventory-to-sales ratio is positive, peaking at about 1 year after the shock. Responses for Retail are similar.

We conclude that inventory-to-sales ratio moves countercyclically over the part of business cycles caused by monetary disturbances. To quantify the extent of the ratio's comovement with output and sales we estimate its respective elasticities *conditional* on monetary shock. Specifically, we regress the fitted inventory-to-sales ratio on fitted output (sales) from regression (1). For Manufacturing and Trade conditional output and sales elasticities are slightly lower, -0.64 and -0.71, than the unconditional elasticities. For the Retail sector elas-

ticities conditional on monetary shocks are slightly larger than the unconditional elasticities: -0.49 for output elasticity and -0.57 for sales elasticity.

3. Model

The economy consists of a continuum of final goods firms, indexed by z , that each sell a differentiated good produced using intermediate goods as inputs; a continuum of competitive intermediate goods firms; and a representative household that derives utility from the consumption of final goods, trades a complete set of state-contingent securities and supplies labor to the intermediate goods sector.

Consumers

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

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

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

Letting γ_t denote the multiplier on the consumer's budget constraint, and $\mu_t(z)$ denote

the multiplier on the constraint that the consumer buys no more than the stock of inventories available for sale, the optimality conditions governing his choice of hours, bond and goods purchases are:

$$U_{c,t} c_t^{\frac{1}{\theta}} v_t(z)^{\frac{1}{\theta}} c_t(z)^{-\frac{1}{\theta}} = p_t(z) \gamma_t + \mu_t(z)$$

$$\mu_t(z)(s_t(z) - c_t(z)) = 0, \quad c_t(z) \leq s_t(z), \mu_t(z) \geq 0$$

$$q_t(s^{t+1}) = \beta \pi_t(s^{t+1}) \frac{\gamma(s^{t+1})}{\gamma_t(s^t)}$$

$$-U_n = \gamma_t W_t$$

Notice that the first FOC can be written as

$$\tilde{\gamma}_t v_t(z)^{\frac{1}{\theta}} c_t(z)^{-\frac{1}{\theta}} = p_t(z) + \tilde{\mu}_t(z)$$

or

$$c_t(z) = v_t(z) [p_t(z) + \tilde{\mu}_t(z)]^{-\theta} \tilde{\gamma}_t^\theta, \quad (2)$$

where $\tilde{\gamma}_t = U_{c,t} c_t^{\frac{1}{\theta}} \gamma_t^{-1}$ and $\mu_t(z) = \tilde{\mu}_t(z) \gamma_t^{-1}$.

If $s_t(z)$ is sufficiently large, the inventory-availability constraint is not binding and $c_t(z) = v_t(z) p_t(z)^{-\theta} \tilde{\gamma}_t^\theta$ and $\tilde{\mu}_t(z) = 0$. If $s_t(z)$ is too low, then $c_t(z) = s_t(z)$ and $\tilde{\mu}_t(z) > 0$. Notice that when $\tilde{\mu}_t(z) > 0$, the consumer's perceived price of good z is $p_t(z) + \tilde{\mu}_t(z)$, i.e., whatever price the firm would have had to choose (were it able to do so) to ensure consumers purchase exactly $s_t(z)$ units of the good.

To find $\tilde{\gamma}_t$, we make use of $c_t^{\frac{\theta-1}{\theta}} = \int_0^1 v(z)^{\frac{1}{\theta}} c(z)^{\frac{\theta-1}{\theta}} dz$ in 2:

$$c_t^{\frac{\theta-1}{\theta}} = \tilde{\gamma}_t^{1-\theta} \left[\int_0^1 v_t(z) [p_t(z) + \tilde{\mu}_t(z)]^{1-\theta} dz \right],$$

or

$$\tilde{\gamma}_t = P_t c_t^{\frac{1}{\theta}}$$

where

$$P_t = \left[\int_0^1 v_t(z) [p_t(z) + \tilde{\mu}_t(z)]^{1-\theta} dz \right]^{\frac{1}{1-\theta}}$$

is the consumer's perceived aggregate price index, i.e., a consumption-weighted average of the perceived prices $p_t(z) + \tilde{\mu}_t(z)$ the consumer is facing. We can write P_t as:

$$P_t = \left[\int_{\tilde{\gamma}_t^\theta v_t(z) p_t(z)^{-\theta} < s_t(z)} v_t(z) p_t(z)^{1-\theta} dz + \tilde{\gamma}_t^{1-\theta} \int_{\tilde{\gamma}_t^\theta v_t(z) p_t(z)^{-\theta} \geq s_t(z)} v_t(z)^{\frac{1}{\theta}} s_t(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{1}{1-\theta}}$$

and thus show that $\tilde{\gamma}_t$ satisfies:

$$\tilde{\gamma}_t = \left[\int_{\tilde{\gamma}_t^\theta v_t(z) p_t(z)^{-\theta} < s_t(z)} v_t(z) p_t(z)^{1-\theta} dz + \tilde{\gamma}_t^{1-\theta} \int_{\tilde{\gamma}_t^\theta v_t(z) p_t(z)^{-\theta} \geq s_t(z)} v_t(z)^{\frac{1}{\theta}} s_t(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{1}{1-\theta}} c_t^{\frac{1}{\theta}} \quad (3)$$

The consumer demand for each of the final goods is thus:

$$c_t(z) = v_t(z) \left(\frac{p_t(z)}{P_t} \right)^{-\theta} c_t \quad \text{if } v_t(z) \left(\frac{p_t(z)}{P_t} \right)^{-\theta} c_t < s_t(z)$$

$$c_t(z) = s_t(z) \quad \text{otherwise}$$

It is also clear (as $\tilde{\mu}_t(z) > 0$) that $\tilde{\gamma}_t > \hat{P}_t c_t^{\frac{1}{\theta}}$ where \hat{P}_t is the usually-defined aggregate price index: $\hat{P}_t = \left[\int_0^1 v_t(z) p_t(z)^{1-\theta} dz \right]^{\frac{1}{1-\theta}}$ and thus $c_t(z) > v_t(z) \left(\frac{p_t(z)}{\hat{P}_t} \right)^{-\theta} c_t$ which is the demand for unconstrained goods in an economy without stockouts. The consumer directs expenditures towards the unconstrained goods whenever some goods are in limited supply.

Given $\tilde{\gamma}_t$, we find $\gamma_t = \frac{U_{c,t} c_t^{\frac{1}{\theta}}}{\tilde{\gamma}_t} = \frac{U_{c,t}}{P_t}$ which can then be used to solve the consumer's

labor-leisure choice and price assets. Notice that $P_t > \hat{P}_t$ and thus the marginal utility of consumption lower than in the absence of stockouts. Finally,

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

and

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

Retail Firms

Retailers buy a storable good from a perfectly competitive intermediate-goods industry at (nominal) price ω_t . One unit of the intermediate good produces $a_t(z)$ units of the final good. $a_t(z)$, the productivity of the firm, is idiosyncratic to each firm and follows a Markov process. Retailers face two fixed adjustment costs, of adjusting prices, κ_p , and of ordering new goods from the intermediate firms, κ_s . We assume that price and inventory decisions are made prior to the realization of that period's demand $v(z)$. We assume for simplicity that $v(z)$ is serially uncorrelated.

To economize on the ordering costs and to ensure against the possibility of high demand $v(z)$, firms carry non-zero inventories of the final good from one period to another. In particular, let $s_t(z)$ be the retailer's beginning-of-period inventory level. The retailer solves the following dynamic program, where $V^{a,a}$ is the firm's value of concomitantly adjusting inventories and prices, $V^{n,a}$ the value of adjusting inventories, but not the price, etc. Given price p , inventory available for sale s^a and log-normally distributed demand $\log v \sim N(0, \sigma^2)$, the firm's expected sales are $R(p, z) = E \min(v \tilde{\gamma}^\theta p^{-\theta}, s^a) = \tilde{\gamma}^\theta p^{-\theta} \exp\left(\frac{\sigma^2}{2}\right) \Phi(\log(p^\theta \tilde{\gamma}^{-\theta} s^a) - \sigma^2) + s^a (1 - \Phi(\log(p^\theta \tilde{\gamma}^{-\theta} s^a)))$

We next write the firm's dynamic program. Let $V^{a,a}(p_{-1}, s_{-1}, a)$ be the firm's value of adjusting prices and inventories, given beginning-of-period stock of inventories s_{-1} , and inherited price p_{-1} , as well as productivity a . $V^{n,a}$ value of not adjusting prices and adjusting inventories, and $V^{a,n}, V^{n,n}$ similarly defined. The firm's value is the envelope of these four

options: $V = \max(V^{a,a}, V^{a,n}, V^{n,a}, V^{n,n})$. Then

$$\begin{aligned}
V^{a,a}(p_{-1}, s_{-1}, a) &= \max_{p, i \geq 0} \left[\frac{U_c}{P} [p\mathcal{R}(p, s_{-1} + i) - a\omega i - (\kappa_p + \kappa_s)W] + \beta EV(p'_{-1}, s'_{-1}, a') \right] \\
V^{n,a}(p_{-1}, s_{-1}, a) &= \max_{i \geq 0} \left[\frac{U_c}{P} [p_{-1}\mathcal{R}(p_{-1}, s_{-1} + i) - a\omega i - \kappa_s W] + \beta EV(p'_{-1}, s'_{-1}, a') \right] \\
V^{a,n}(p_{-1}, s_{-1}, a) &= \max_p \left[\frac{U_c}{P} [p\mathcal{R}(p, s_{-1}) - \kappa_p W] + \beta EV(p'_{-1}, s'_{-1}, a') \right] \\
V^{n,n}(p_{-1}, s_{-1}, a) &= \left[\frac{U_c}{P} p\mathcal{R}(p_{-1}, s_{-1}) + \beta EV(p'_{-1}, s'_{-1}, a') \right]
\end{aligned}$$

where i is the amount ordered. The law of motion for p_{-1} is $p'_{-1} = p$ if adjust p and p_{-1} if don't adjust p . The law of motion for $s'_{-1} = (s_{-1} + i - \min(\nu\gamma^\theta p^{-\theta}, s_{-1} + i))(1 - \delta)$ if order and $s'_{-1} = (s_{-1} - \min(\nu\gamma^\theta p^{-\theta}, s_{-1}))(1 - \delta)$ otherwise.

Intermediate good firms (wholesalers)

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

$$y = l$$

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

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

where W is the nominal wage rate. In equilibrium free entry drives intermediate's profits to zero and thus $\omega = W$.

Equilibrium

We impose a quantity-theory equation $M = \int_0^1 p(z)c(z)dz$ rather than derive the

demand for money explicitly in order to avoid introducing the labor-leisure etc. distortions associated with other specifications (CIA or MIU). Given M , we solve for $\tilde{\gamma}$ using

$$\int_0^1 p(z)c(z; \tilde{\gamma})dz = M$$

where $c(z, \tilde{\gamma})$ is the consumption of each variety given the multiplier $\tilde{\gamma}$. M is set equal to 1 in the steady-state.

A. Economy with no adjustment frictions

In the absence of frictions, and in the absence of $i \geq 0$ constraint, the firm's value is linear in s_{-1} (with slope ω), and the choice of p and $s^a (= s + i)$ reduces to:

$$\max_{p, s^a} \frac{U_c}{P} [p - \beta\omega(1 - \delta)] \mathcal{R}(p, s^a) - \frac{U_c}{P} \omega(1 - \beta(1 - \delta)) s^a$$

To understand this expression, recall that the only difference now between our problem and the original problem of a monopolistically competitive firm is that prices and inventories must be decided prior to the realization of demand. Thus, the choice of price is the same as in the original problem, but $\mathcal{R}(p, s^a)$ (as opposed to $p^{-\theta}$) is the demand curve and $\beta \frac{U_c}{P} \omega(1 - \delta)$ (as opposed to $\frac{U_c}{P} \omega$) is the real (expressed in term of marginal utility of consumption) marginal cost of supplying one extra good. The reason this last one differs from $\frac{U_c}{P} \omega$ is the fact that the saving from lowering your sales by 1 unit today only accrue next period when you can resell the additional unit. The firm thus takes the discounting and depreciation rate into account. The choice of s^a is also straightforward: on one hand higher s^a increases the probability of a sale, but the firm loses $\frac{U_c}{P} \omega(1 - \beta(1 - \delta))$ on every unspent additional unit of inventory to depreciation.

The first-order conditions in the frictionless case are

$$p = \frac{\varepsilon}{\varepsilon - 1} \omega \beta (1 - \delta) \text{ where } \varepsilon = -\frac{R_p p}{R} = \theta y_u \text{ and}$$

$$y_u = \frac{\int_0^{\frac{s^a}{\gamma^\theta p^{-\theta}}} v \gamma^\theta p^{-\theta} f(v) dv}{\int_0^{\frac{s^a}{\gamma^\theta p^{-\theta}}} v \gamma^\theta p^{-\theta} f(v) dv + s^a \left(1 - F\left(\frac{s^a}{\gamma^\theta p^{-\theta}}\right)\right)}$$

is share of output sold in states of v without stockouts.

The choice of s^a satisfies

$$1 - F\left(\frac{s^a}{\gamma^\theta p^{-\theta}}\right) = \frac{(1 - \beta(1 - \delta))}{\frac{p}{\omega} - (1 - \delta)\beta}$$

A higher markup $\frac{p}{\omega}$ thus makes the firm less likely to stockout as having an additional unit of inventories is more valuable.

B. Calibration

The period is a month, and thus the discount factor equal to $.96^{\frac{1}{12}}$. We assume preferences of the form $U(c_t, n) = \log c_t - \psi_t n_t$. We think of ψ_t as a preference shocks that in the various experiments will be allowed to vary with the monetary shock. We assume $\log(a)$ is

discretely distributed with $[-\bar{a}, 0, \bar{a}]$. The transition matrix is $\begin{bmatrix} 1 - \rho & \rho & 0 \\ \frac{\rho}{2} & 1 - \rho & \frac{\rho}{2} \\ 0 & \rho & 1 - \rho \end{bmatrix}$. Thus

the firm's log productivity is on average equal to 0 and deviations arrive with a Poisson-like proces with probability ρ . The discrete arrival of shocks allows the model to generate the large price changes observed in the data without a large selection effect¹². We choose ρ to match a frequency of price changes of 5 months given that our model has no motive for

¹²Midrigan (2006), Gertler-Leahy (2007).

temporary price changes. A recent paper by Midrigan and Kehoe (2008) shows that a model with a motive for temporary price changes (sales) in which prices change on average every 3 weeks, as in the data of prices in retail stores, but calibrated to match a rich set of features of the micro-price data, including the fraction of temporary versus permanent prices and the frequency with which sales revert to the pre-sale price, generates real effects of money similar to those of a model without a motive for sales in which 'regular' prices change every 5 months. The intuition why the real effects of money are larger is that temporary price changes most often return to their pre-change level and as a result do not respond to low frequency variations in the monetary policy. The upper bound on technology shocks, \bar{a} , is chosen to match a frequency of price changes of 10%, as documented by Bils-Klenow (2005) and Klenow-Kryvtsov (2008).

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

The fixed ordering cost, κ_s , together with the volatility of demand, σ , are chosen so that the model generates a monthly inventory-to-sales ratio of 1.4 months, as in the US retail trade sector and a correlation between changes in prices and quantities of -0.2 at the monthly frequency, as in the scanner price data for Dominicks. The last statistic was reported by Burstein and Hellwig (2007) for regular price changes at the monthly frequency using Dominick's data. The depreciation rate is set to 2.5%, a number in the range of estimates of inventory-carrying costs presented by Richardson (1995). Finally, the elasticity of demand, θ , is set equal to 3, in the mid-range of estimates in retail markets, and the discount factor to $\beta = 0.96^{\frac{1}{12}}$, given the monthly frequency of the model.

Table 5 reports the parameter values used. Table 6 reports the moments targeted. In addition to the 4 moments we use in calibration to pin down $\rho, \bar{a}, \sigma, \kappa_s$ (the rest of the parameters are assigned), we do well at matching additional features of the data. In particular, the model predicts that firms stockout 5% of the time and that 56% of firms order inventories

in every given week. These numbers are in the range of those reported in earlier work by Bils (2004), Aguirregabiria (2005) for the frequency of stockouts: 5%-8%, Aguirregabiria (1999): 79% frequency of orders in the data on chain-wide orders of a Spanish supermarket, and Alessandria et. al (2008) who use data from Hall-Rust (1999) of the orders of a large steel wholesaler and find a frequency of orders for domestic goods of 0.33 to 0.50 per month.

C. Decision rules

Before we proceed with our experiments, we briefly discuss optimal firm policy rules in our environment. We refer the reader to Aguirregabiria (1999) for a more detailed analysis of how fixed costs of ordering and price changes affect a retailer's optimal price and ordering decisions, and an existence/uniqueness proof.

Notice from the Bellman equations characterizing firm's dynamic program, that the optimality conditions governing the choice of inventories and prices are (to avoid cluttering the notation, we write these assuming away uncertainty about a which is set at its mean value as well as assuming the non-negativity constraint is not binding (it does not if a does not fluctuate):

$$p : \frac{U_c}{P} [R(p, s^a) + pR_p(p, s^a)] + \beta \int_0^\infty V_1(p, s'(v))f(v)dv = -\theta\beta(1 - \delta) \int_0^{\frac{s_{-1}+i}{\gamma^\theta p^{-\theta}}} V_2(p, s'(v))f(v)dv$$

$$i : \beta(1 - \delta) \int_0^{\frac{s_{-1}+i}{\gamma^\theta p^{-\theta}}} V_2((p, s'(v)) f(v)dv = \frac{U_c}{P} \left[\omega - p \left(1 - F \left(\frac{s^a}{\gamma^\theta p^{-\theta}} \right) \right) \right],$$

where recall $s^a = s_{-1} + i$ is the stock of inventories available after the ordering decision is made but before sales.

The firm thus orders inventories so as to equalize the marginal valuation of an additional unit of inventories next period to the marginal cost of acquiring it (which with stockouts is lower than the replacement cost, $\frac{U_c}{P}\omega$, as an additional unit of inventories allows the firm to avoid a stockout). Similarly, the firm chooses the price to equalize marginal revenue (which

here includes the savings of the price adjustment cost, as measured by $V_1(p)$ to the marginal valuation of inventories next period. Notice here again the sense in which inventory behavior can be informative about the wages (wholesale prices) the firm faces. The stickier wages are (the lower $\frac{U_c}{P}\omega$ in times of a monetary expansion), the more incentive the firm has to increase inventories s' , so as to lower their marginal valuation (the inventory frictions make V_2 is (K) -concave), and the more incentive to keep prices low.

Figure 6 plots the regions of the state (p_{-1}, s_{-1}) space in which the firm finds it optimal to adjust both inventory and price (heavy-shaded regions in the upper and lower left corners), adjust only inventory (intermediate-shaded region to the left), or adjust only price (light-shaded regions to the right). The region where inaction along both margins is optimal is the unshaded center region. We also plot the optimal price (solid line) and order decisions (we plot $s^a = s_{-1} + i$, the inventory available for sale, using dotted lines) conditional on the firm adjusting both prices and inventories. The three panels correspond to the low, medium, and high productivity of the retailer.

Notice that the adjustment regions are of the familiar (s, S) type: the firm adjusts its price/orders inventories, when its old price or current inventories are too far out of line. Conditional on ordering, the return point for inventories is independent of p_{-1}, s_{-1} as the two are reset (the order decision is drawn for a firm with a stock s_{-1} sufficiently low so that the irreversibility constraint does not bind). Finally, notice that the firm's optimal price, conditional on adjustment, is approximately equal to its frictionless optimum up to a point at which the irreversibility constraint on inventories binds: the firm's stock s_{-1} exceeds the return point). As s_{-1} increases above this point, the firm finds it optimal to run down its inventories by decreasing its price so as to sell more rapidly and avoid paying the depreciation cost. The model thus exhibits sluggish response of prices to negative productivity shocks as firms keep prices low exhaust the excess inventories.¹³ Finally, notice across the three panels that increasing the firm's cost of producing the final good, a , shifts the inaction region to the

¹³Alessandria, Kaboski, Midrigan (2008) study how this feature of the model can account for the slow pass-through of imported goods prices at the retail level following a large devaluation.

north-west and thus makes it less likely that the firm will order and more willing to tolerate higher prices. The optimal prices and orders also increase and fall respectively as the firm's cost increases.

Figure 7 plots the optimal decision rules conditional on adjustment/inaction in the s_{-1} and p_{-1} space, as well as their envelope that summarizes the choice that maximizes across the different options for $a = 0$. The left panel shows that if the firm does not adjust its price, the optimal inventory holdings fall with the past price as the firm expects to sell less in future periods. If the firm adjusts its price, its old price is irrelevant for its decision. When the firm's price is sufficiently far out of line, the firm adjusts it and this reflects on its optimal inventories. Similarly, the right panel shows that the firm's price, conditional on not ordering, decreases with higher inventory: if the stock available for sale is too low, the firm is likely to stock out and charges a high price as it perceives a lower elasticity of sales with respect to its price. Notice also the non-monotonicity in this price schedule: if the stock is very low, close to 0, the firm is better off keeping the price low today to avoid paying the menu costs of lowering the price next period: too few inventories make the gains from raising the price too small to warrant the adjustment costs on prices.

Figure 8 plots the (optimal, i.e., implied by the decision rules in Figure 7) expected inventory-to-sales ratio in the p_{-1} space, again for $a = 0$. In particular, we report $\frac{s^a - E_v \min(v\gamma^\theta p^{-\theta})}{E_v \min(v\gamma^\theta p^{-\theta})}$ conditional on ordering ($i > 0$, left panel) and not ordering ($i = 0$, right panel), as a function of the firm's past price and given the firm's optimal price adjustment decision conditional on the inventory decision in each panel. Notice in the left panel of Figure 8, that when the past price is sufficiently out of line, the firm adjusts it and orders an amount that is independent of the past price which in this case has no effect on the amount ordered. If the firm's past price is close to its optimum, the firm leaves it unchanged, but chooses an inventory level that, although decreases in the past price (as the solid line of Figure 7 shows), results in an expected inventory-to-sales ratio that increases in the firm's price. As in Bilal and Kahn (2000), a higher price makes an additional unit of sales more valuable as the gains from avoiding a stockout are larger. In the right panel of the Figure we plot the same inventory-sales ratio

conditional on the firm not ordering. In this case the firm's price increases the inventory-sales ratio by even more as it simply reduces sales without affecting inventories. This figure thus shows that in response to a decrease in markups, induced by a monetary expansion, firms would, holding all else constant, hold less inventories relative to sales, even conditional on ordering inventories.

4. Experiments

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

A. No real rigidities

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

$$\frac{W}{P} = -\frac{U_n}{U_c} = \psi C$$

and immediately rises in response to the increase in aggregate consumption with unitary elasticity. Moreover, the price of wholesale goods faced by the firms, expressed in units of marginal utility of consumption, $\frac{\omega U_c}{P} = \frac{W U_c}{P}$ (recall that this is the relevant price affecting ordering decisions) is irresponsive to the monetary shock: absent the countercyclical markups channel (and in the absence of fixed costs of ordering and irreversibility), the inventory-to-sales ratio would not respond to the monetary shock as the real cost of purchasing goods from the wholesale is constant during the transition.

Figure 9 plots the actual impulse responses in this benchmark economy. Here inventory-sales ratios drop, mainly because of the unexpected increase in consumption, but also because

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

B. Real rigidities

We next consider a preference shock ψ_t , that decreases the marginal disutility from work in the aftermath of the monetary expansion. In particular, we choose the path for ψ_t so that the nominal wage rate follows

$$\begin{aligned} W_t &= W_t^{*(1-\lambda)} W_{t-1}^\lambda, \text{ where} \\ W_t^* &= \psi P_t C_t \end{aligned}$$

is the wage rate in the benchmark economy. We think of the preference shock ψ_t as a shortcut and crude way of modeling features of the labor market that lead to a slower response of the real marginal cost of production.

Figure 10 computes impulse responses assuming $\lambda = \frac{1}{2}$. Notice that in this case real wages increase much less than C , and as a result the price of intermediate goods, $\frac{\omega U_c}{P}$, drops initially and then rises. Firms find it thus optimal to purchase inventories immediately, rather than wait. This is seen in the lower levels of the Figure which show that orders increase by 20%, whereas inventories increase on impact by 13%. Table 7 shows that the elasticity of real wages to consumption is 0.53, while that of inventory-to-sales to sales increases to 4.56,

much higher than in the data.

In Figure 11 we compute the impulse responses assuming $\lambda = 0.15$. This is the value that allows our model to most closely replicate the response of inventories observed in the data. In particular, in this case the elasticity of real wages to consumption is 0.92, while that of inventory-sales ratio to sales is -0.79.

Our simple experiments thus suggest that an elasticity of real wages to consumption of slightly below unity is what is required to account for the response of inventories to monetary shocks in the data.

5. Conclusions

As pointed out by Bils and Kahn, the behavior of inventories over the cycle may be informative about how markups and real marginal costs vary with the cycle. We apply their insight to quantify the role of sluggish adjustment of factor prices (real flexibilities) in accounting for the long delays with which prices respond to changes in the stance of monetary policy. We show that a model which combines the stockout-avoidance and (s,S) motives for holding inventories, calibrated to match salient features of the microeconomic data, requires an elasticity of real marginal cost with respect to output of slightly less than unity, in order to account for the behavior of inventories in the aftermath of monetary policy shocks.

References

- [1] Aguirregabiria, Victor, 1999. "The Dynamics of Markups and Inventories in Retailing Firms," *The Review of Economic Studies*, 66(2), 275-308.
- [2] Aguirregabiria, Victor, 2005. "Retail Stockouts and Manufacturer Brand Competition," mimeo.
- [3] Alessandria, George, Joseph Kaboski, and Virgiliu Midrigan, 2008. "Inventories, Lumpy Trade, and Large Devaluations," Federal Reserve Bank of Philadelphia Working Paper 08-3.

- [4] Ball, Laurence and David Romer, 1990. "Real Rigidities and the Non-Neutrality of Money," *The Review of Economic Studies*, 57(2), 183-203.
- [5] Bilts, Mark, 2004. "Studying Price Markups from Stockout Behavior," mimeo.
- [6] Bilts, Mark and Peter J. Klenow, 2004. "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947-985.
- [7] Bilts, Mark, and Klenow, Peter J. and Oleksiy Kryvtsov, 2003. "Sticky Prices and Monetary Policy Shocks," *Federal Reserve Bank of Minneapolis Quarterly Review*, Winter, 2-9.
- [8] Burstein, Ariel and Christian Hellwig, 2007. "Prices and Market Shares in a Menu Cost Model," mimeo.
- [9] Caplin, Andrew, 1985. "The Variability of Aggregate Demand with (S,s) Inventory Policies," *Econometrica*, 53(6), 1395-1409.
- [10] Christiano, Lawrence J. and Eichenbaum, Martin and Charles L. Evans, 1999. "Monetary policy shocks: What have we learned and to what end?," *Handbook of Macroeconomics*, in: J. B. Taylor & M. Woodford (ed.), *Handbook of Macroeconomics*, edition 1, volume 1, chapter 2, pages 65-148 Elsevier.
- [11] Christiano, Lawrence J. and Eichenbaum, Martin and Charles L. Evans, 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, University of Chicago Press, 113(1), 1-45.
- [12] Chari, V. V. and Patrick J. Kehoe and Ellen R. McGrattan, 2000. "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?," *Econometrica*, 68(5), 1151-1180.
- [13] Dotsey, Michael and Robert J. King, 2005. "Implications of state-dependent pricing for dynamic macroeconomic models," *Journal of Monetary Economics*, 52(1), 213-242.

- [14] Dotsey, Michael and Robert G. King, 2006. "Pricing, Production, and Persistence," *Journal of the European Economic Association*, 4(5), 893-928.
- [15] Fisher, Jonas and Andreas Hornstein, 2000. "(S, s) Inventory Policies in General Equilibrium," *The Review of Economic Studies*, 67(1), 117-145.
- [16] Friedman, Milton, 1968. "The Role of Monetary Policy," *American Economic Review*, 58(1), 1-17.
- [17] Gertler, Mark and Leahy, John, 2006. "A Phillips Curve with an Ss Foundation," NBER Working Paper 11971.
- [18] Hall, George and John Rust, 2000. "An Empirical Model of Inventory Investment by Durable Commodity Intermediaries," *Carnegie-Rochester Conference Series on Public Policy*, 52, 171-214.
- [19] Hall, George and John Rust, 2002. "Econometric Methods for Endogenously Sampled Time Series: The Case of Commodity Price Speculation in the Steel Market," mimeo.
- [20] Hornstein, Andreas and Harald Uhlig, 2000. "What is the Real Story for Interest Rate Volatility?," *German Economic Review*, 1(1), 43-67.
- [21] Khan, Aubhik and Julia Thomas, 2007a. "Inventories and the business cycle: An equilibrium analysis of (S,s) policies." *American Economic Review*, 97(4), 1165-88.
- [22] Kehoe, Patrick J. and Virgiliu Midrigan, 2007. "Sales and the real effects of monetary policy," Working Papers 652, Federal Reserve Bank of Minneapolis.
- [23] Khan, Aubhik and Julia Thomas, 2007b. "Explaining Inventories: A Business Cycle Assessment of the Stockout Avoidance and (S,s) Motives," *Macroeconomic Dynamics*, 11(5), 638-64.
- [24] Kimball, Miles, 1995. "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, 27(4), 1241-1277.

- [25] King, Robert G and Mark W. Watson, 1996. "Money, Prices, Interest Rates and the Business Cycle," *The Review of Economics and Statistics*, 78(1), 35-53.
- [26] Klenow, Peter J. and Oleksiy Kryvtsov, 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" forthcoming in *The Quarterly Journal of Economics*.
- [27] Klenow, Peter J. and Jon Willis, 2006. "Real Rigidities and Nominal Price Changes," mimeo..
- [28] Levy, Daniel, and Bergen, Mark, and Dutta, Shantanu, and Robert Venable, 1997. "The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains," *The Quarterly Journal of Economics*, 112(3), 791-825.
- [29] Nakamura, Emi and Jón, Steinsson, 2007. "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model," mimeo.
- [30] Nakamura, Emi and Jón, Steinsson, 2008. "Five Facts About Prices: A Reevaluation of Menu Cost Models," forthcoming in *The Quarterly Journal of Economics*.
- [31] Ramey, Valerie A and Daniel J. Vine, 2004. "Why Do Real and Nominal Inventory-Sales Ratios Have Different Trends?," *Journal of Money, Credit and Banking*, 36(5), 959-63.
- [32] Ramey, Valerie A. and Kenneth D. West, 1999. "Inventories," Handbook of Macroeconomics, in: J. B. Taylor & M. Woodford (ed.), Handbook of Macroeconomics, edition 1, volume 1, chapter 13, pages 863-923
- [33] Richardson, H. 1995. "Control Your Costs then Cut Them" *Transportation and Distribution*, 94-96.
- [34] Romer, Christina D. and David H. Romer, 2004. "A New Measure of Monetary Shocks: Derivation and Implications," *The American Economic Review*, 94(4), 1055-1084.

- [35] Solon, Gary and Barsky, Robert and Jonathan A. Parker, 1994. "Measuring the Cyclicity of Real Wages: How Important is Composition Bias," *The Quarterly Journal of Economics*, 109(1), 1-25.
- [36] Wen, Yi, 2008. "Input and Output Inventory Dynamics," Federal Reserve Bank of St. Louis Working Paper 2008-008A.

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

	Mean	Std	Serr corr	Correlation with output sales	
<i>NIPA monthly</i>					
Manufacturing and Trade	1.41	2.19	0.88	-0.83	-0.83
Retail	1.31	2.08	0.72	-0.49	-0.61
<i>NBER annual</i>					
Manufacturing	0.23*	3.90	0.31	-0.52	-0.66

Note: Data are taken from the BEA National Income and Product Accounts monthly data from January 1967 to December 1997 and the NBER Manufacturing Productivity Database from 1957 to 1996. Sales are defined as real final sales to domestic purchasers in NIPA data and real value of industry shipments for NBER data. Output is the sum of final sales and the change in the end-of-period real inventory stock. Inventory-to-sales ratio is defined as the ratio of the end-of-period inventory stock to final sales in that period. All data are HP filtered. Output, sales and inventory-to-sales ratio are defined in % deviations from respective HP trends. Inventory investment is defined in % points-of-output-fraction deviations from its HP trend. * weighted mean across industries in 1996. In 1957 it is 0.33.

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

	Level of Inventory-to-Sales Ratio			
	Low	Medium	High	Aggregate
	<i>Mean</i>			
Inventory-to-sales ratio	0.18	0.31	0.55	0.31
Sales share	0.47	0.27	0.26	1
	<i>Correlation with real GDP</i>			
Inventory-to-sales ratio	-0.41	-0.33	-0.47	-0.47
Sales share	0.12	0.09	0.03	0.13

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

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

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

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

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

	Across time	Cross-section
sales elasticity	-0.57	-0.12

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

Table 5: Parameter values

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


Table 6: Moments in data and Model

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

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


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

Figure 1: Countercyclicality of Inventory-to-Sales Ratio




Note: The NBER Manufacturing Productivity Database from 1957 to 1996.

Figure 2: Time variance of aggregate I/S ratio is not due to composition effects



The NBER Manufacturing Productivity Database from 1957 to 1996

**Figure 3: Responses of Federal Funds rate to its +1% innovation
January 1967- December 1996**



**Figure 4: Responses to +1% innovation to Federal Funds Rate
Manufacturing and Trade, January 1967 - December 1996**




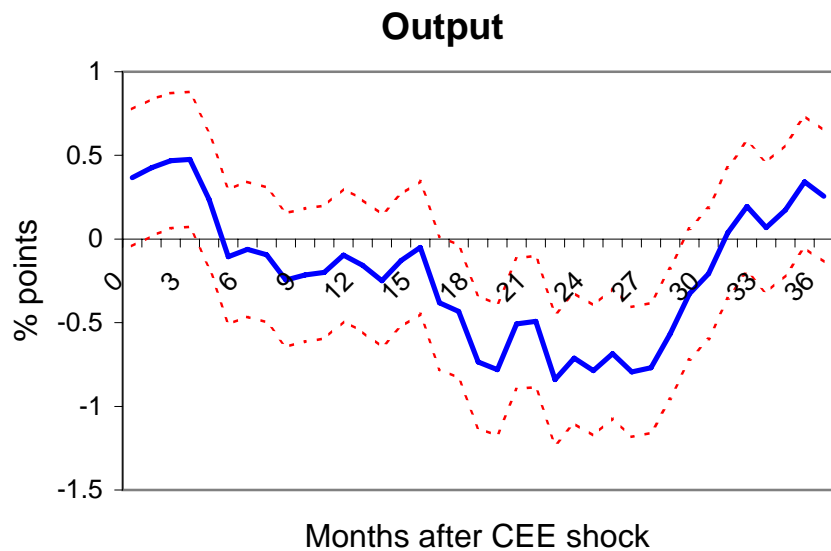


Figure 5: Responses to One S.D. Innovation to Federal Funds Rate, Manufacturing and Trade, January 1967 - December 1996

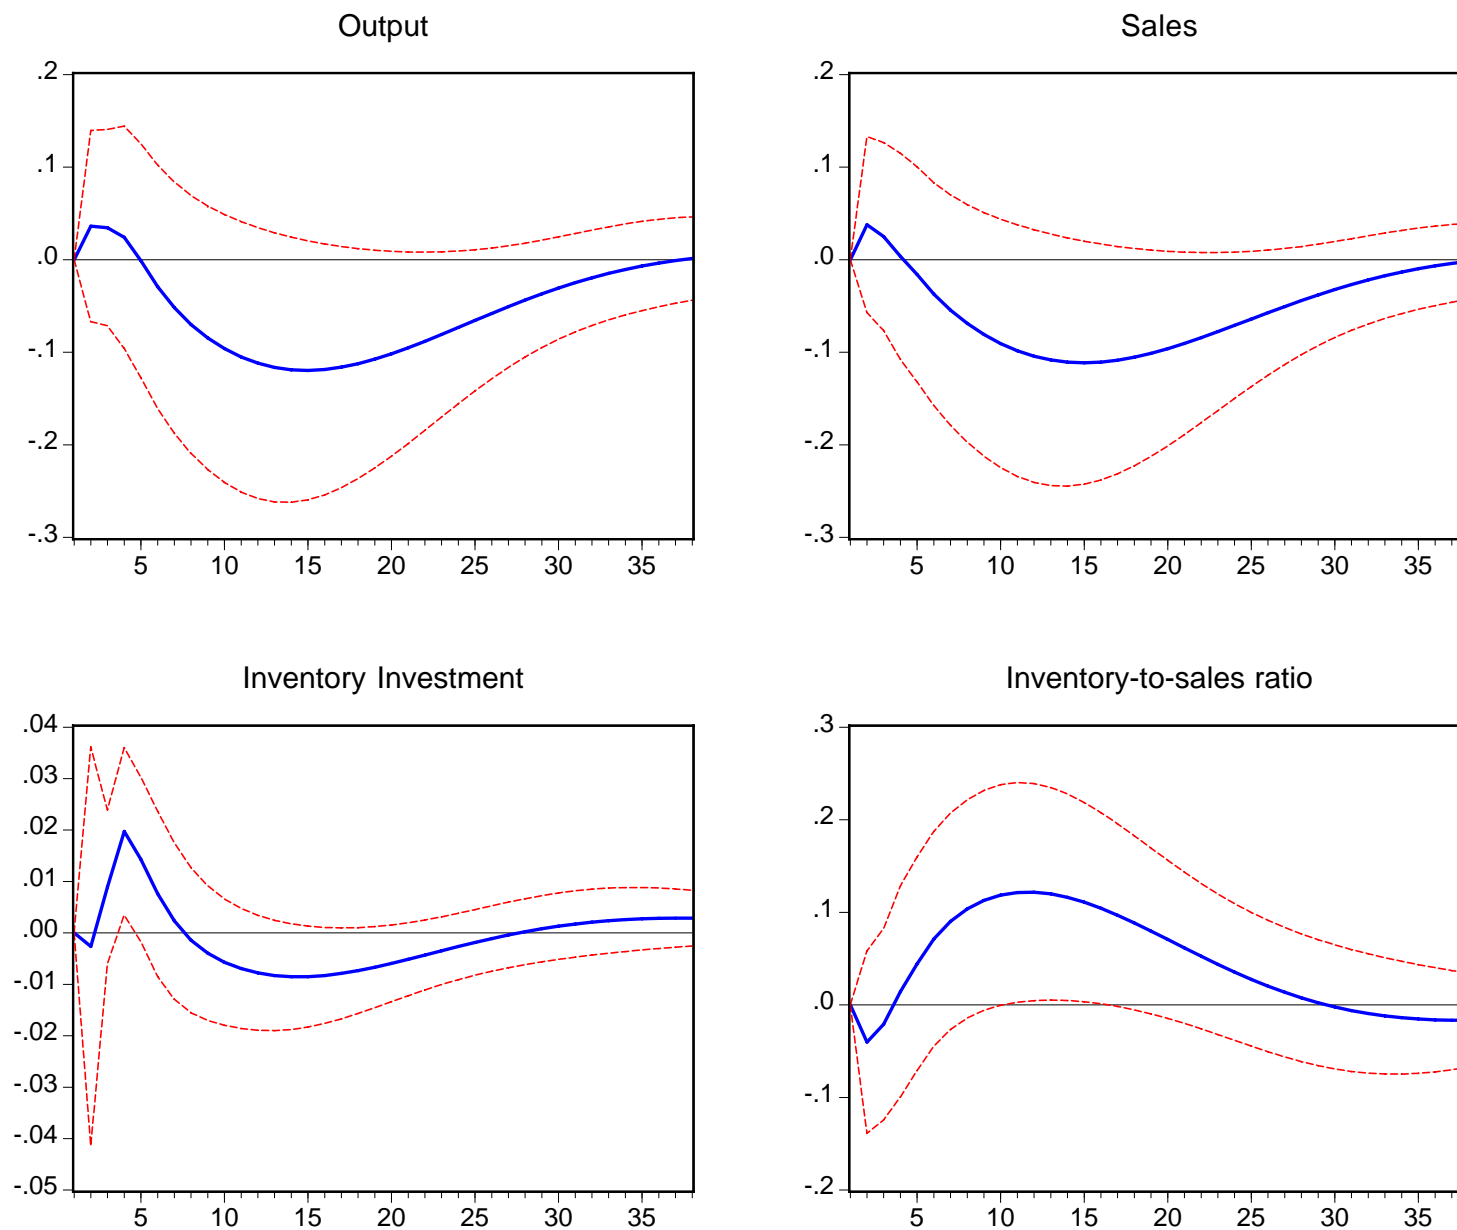


Figure 6: Inaction regions in model economy

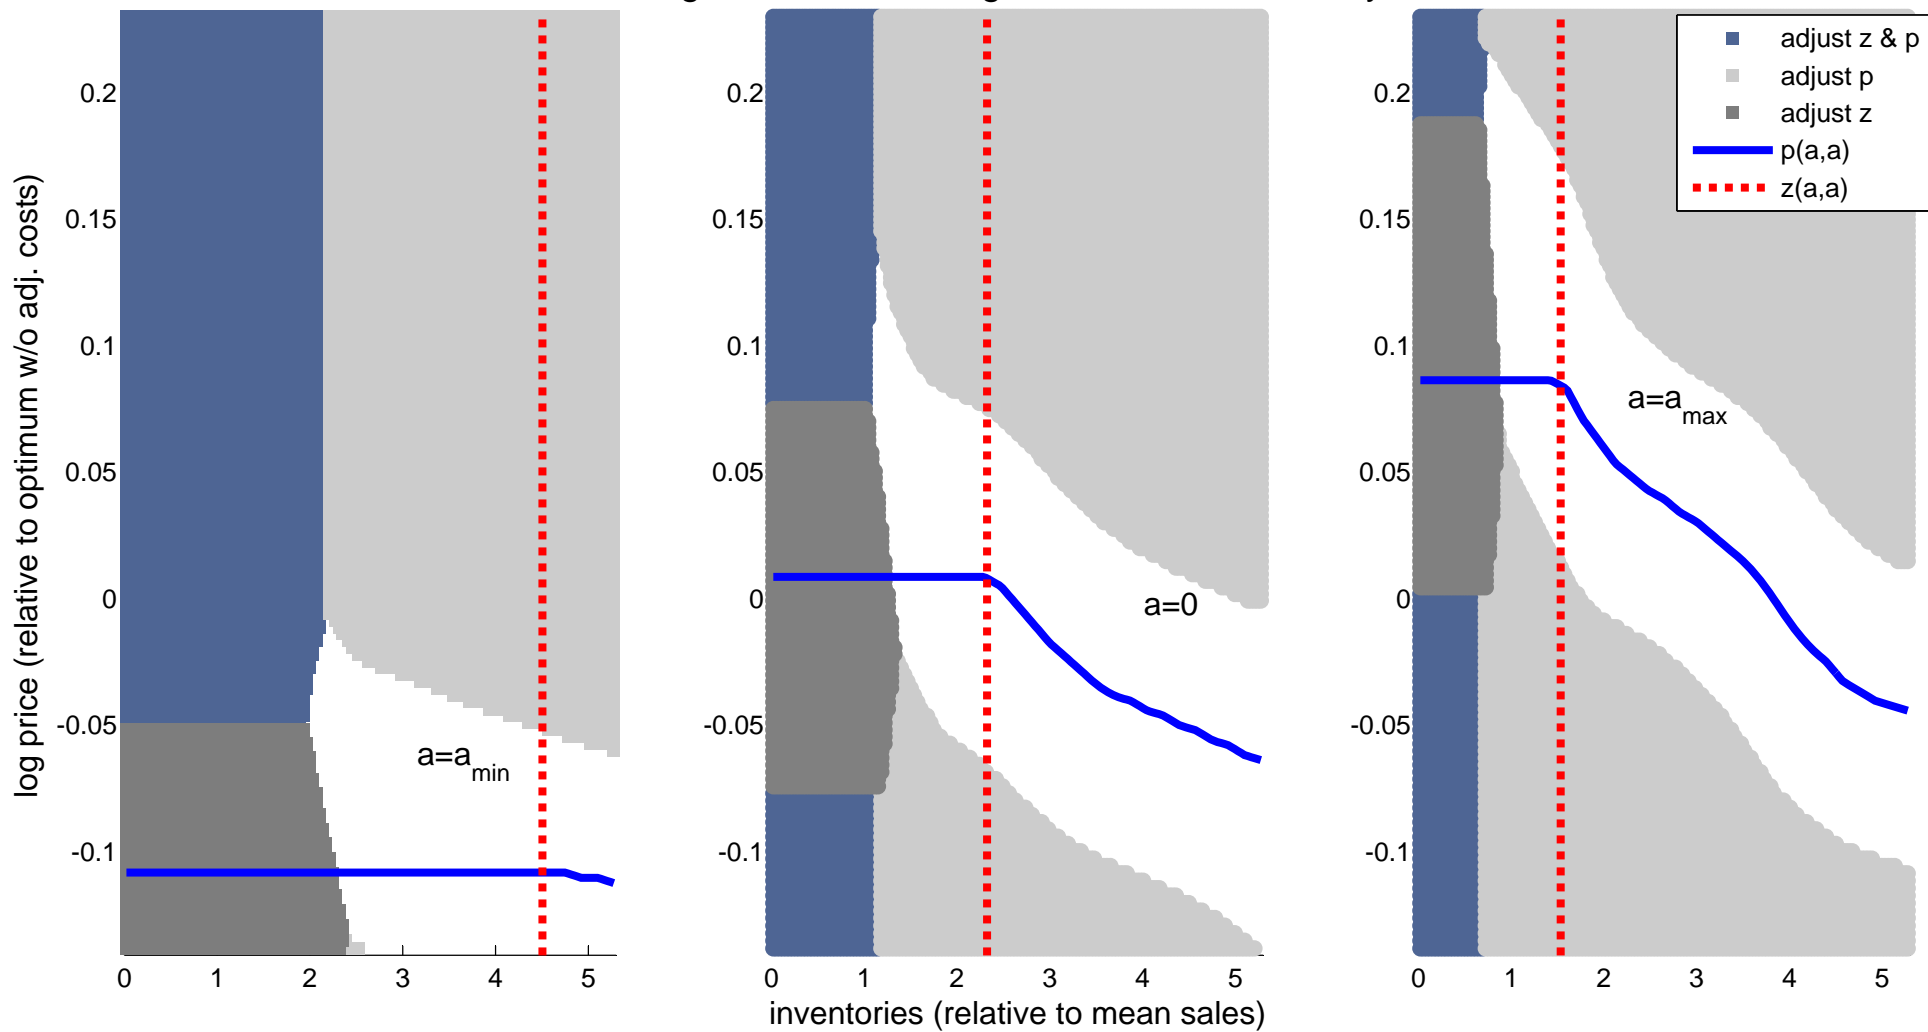


Figure 7a: Optimal inventory conditional on ordering




Figure 7b: Optimal price conditional on adjustment




Figure 8a: Expected I/S conditional on ordering and optimal price decision




Figure 8b: Exected I/S conditional on not ordering and optimal price adjustment decision




Figure 9: Benchmark



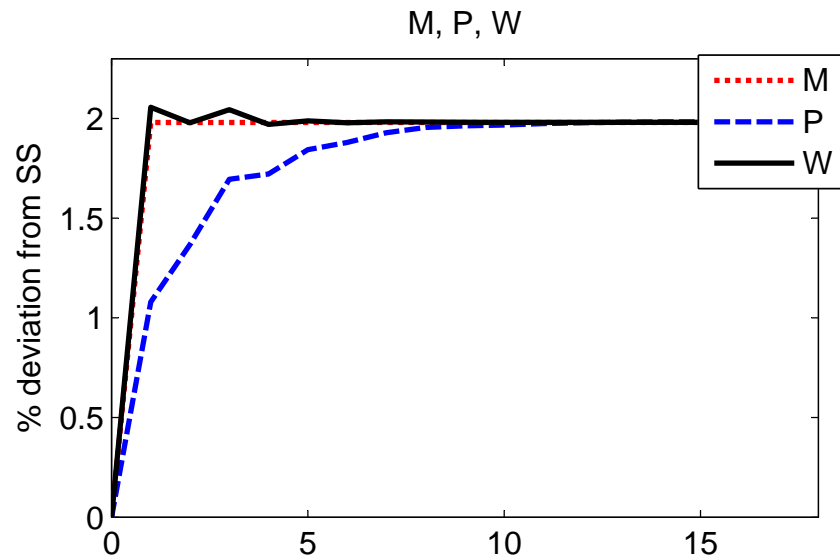


Figure 10: Large RR







Figure 11: Inventory-consistent RR

