Financing Constraints and Inventories

By

Ward Brown
and
Urs Haegler

DISCUSSION PAPER 367

December 2000

FINANCIAL MARKETS GROUP
AN ESRC RESEARCH CENTRE

LONDON SCHOOL OF ECONOMICS

ISSN 0956-8549-367

Financing Constraints and Inventories

Ward Brown[1]

and

Urs Haegler[2]

This Version: December 2000

Abstract

This paper puts forward the existence of of financing constraints as a possible explanation for two main empirical regularities about inventories; that (i) inventory investment is procyclical, and that (ii) the inventory-sales relationship displays highly positive serial correlation. There are no costs shocks, and in the numerical computations demand shocks are assumed to be serially uncorrelated. When financing constraints are not binding, the model predicts that the firm's optimal inventory investment is counter-cyclical. However, this prediction is reversed for a firm with binding financing constraints. Moreover, some persistence in the inventory-sales relationship is also generated by the model.

[1]International Monetary Fund, Washington, D.C.

[2]Department of Economics at Royal Holloway University of London, and Financial Markets Group, London School of Economics.

[3]The authors would like to thank John Moore, and Nobu Kiyotaki for their valuable comments.

Non-technical summary

There are two main empirical regularities about inventories. First, inventory investment is procyclical. Second, the inventory-sales relationship displays highly positive serial correlation. This paper puts forward a possible explanation for these observed empirical regularities: the existence of financing constraints. In our model, firms produce output with a convex production technology and pay convex storage costs on any unsold inventory. There are no costs shocks, and in the numerical computations demand shocks are assumed to be serially uncorrelated. These assumptions imply a procyclical target inventory level (as in the basic version of the linear-quadratic model). Therefore, when financing constraints are not binding, this model predicts that the firm's optimal inventory investment is counter-cyclical. However, this prediction is reversed for a firm with binding financing constraints. The intuition behind this result is that after a low demand realisation a constrained firm does not have sufficient cash (after paying storage costs) to enable it to produce as much as it had sold, and thereby return to the optimal target inventory level. Similarly, for a firm whose actual inventory level is below this optimum, a high demand realisation will cause it to produce more than it has sold in an effort to return to the optimum level as quickly as possible. Thus for constrained firms, inventory behaviour is consistent with both empirical regularities. It is procyclical, and the inventory-sales relationship is positively serially correlated.

1 Introduction

There are two main empirical regularities about inventories.[1] First, inventory investment is procyclical. Second, the inventory-sales relationship displays highly positive serial correlation. The standard workhorse model to study inventories - the linear-quadratic model (Holt, Modigliani, Muth, and Simon (1960)) - has difficulties to explain these empirical regularities. In that model, convex costs lead to the prediction that firms will use their inventories to smooth production, which implies counter-cyclical inventory investment. To explain the highly persistent relationship between inventories and sales, the benchmark model relies upon implausibly high parameters for either the convexity of the cost function or costs of inventory adjustment.

However, empirical work does generally support the linear-quadratic model's assumption of convex production costs.[2] Therefore, the literature has focused on two other possible explanations of these empirical facts, both of which can be incorporated into the linear-quadratic framework:[3] (i) serially correlated cost shocks,[4] or (ii) a flexible-accelerator mechanism which relies on serially correlated demand shocks.[5]

This paper puts forward a third possible explanation for the observed empirical regularities: the existence of financing constraints. It extends earlier work by the authors on this topic.[6] We model the problem in a dynamic-programming framework. We derive some general results analytically, but to demonstrate the results concerned with the above empirical regularities we use a numerical solution methods. In our model, firms produce output with a convex production technology and pay storage costs on any unsold inventory. There are no cost shocks, and in the numerical solution, demand shocks are assumed to be serially uncorrelated.

[1] See Ramey and West (1997) for a detailed discussion.

[2] Non-convex costs are of interest because they can explain procyclical inventory investment. But evidence for non-convex costs has been found only for a select group of industries (see, for example, Hall (1996), Bresnahan and Ramey (1994). Broader support for non-convex costs does not exist (see Ramey and West (1997) for a summary of the empirical literature).

[3] However, many of the papers which explore these alternative explanations do not use the linear-quadratic model.

[4] See, for example, Blinder (1986), Eichenbaum (1989).

[5] See, for example, Kahn (1987)) and Bils and Kahn (2000).

[6] See Brown and Haegler (1999), where we present a 'pen-and-paper' solution for a very simple numerical example, as well as an analytical proof of the 'excess-volatility-of-production' result.

These assumptions imply that when financing constraints are not binding, the firm's optimal inventory investment is counter-cyclical.

However, this prediction is reversed for a firm with binding financing constraints. The intuition behind this result is that after a low demand realisation a constrained firm does not have sufficient cash (after paying storage costs) to enable it to produce as much as it had sold, and thereby attain the optimal target inventory level. Similarly, for a firm whose actual inventory level is below this optimum, a high demand realisation will cause it to produce more than it has sold in an effort to reach the optimum level as quickly as possible. Thus for constrained firms, inventory behaviour is consistent with both empirical regularities. It is procyclical, and the ratio of inventories to sales is positively serially correlated.

An interesting feature of this model is that, unlike other models in the literature, its results do not depend upon an assumption of serial correlation in unobservable exogenous shocks to either cost or demand functions. Since a constrained firm's production affects the demand faced by producers of intermediate goods, our model suggests that inventory investment may have a more active impact on business cycles through its potential to amplify exogenous shocks. Further research is required to examine the impact of financing constraints in environments where exogenous shocks are serially correlated.

There is empirical support for the hypothesis that financing constraints affect the inventory investment of some firms. For example, Gertler and Gilchrist (1994) find evidence that small manufacturing firms draw down their inventory stocks heavily following a monetary contraction, whereas large firms appear to borrow in order to smooth the impact of a downturn on their inventory behaviour. Similarly, using U.S. panel data, Carpenter, Fazzari, and Petersen (1994) find that the inventory investment of small firms is more sensitive to cash flow than is the inventory investment of large firms.[7]

The rest of the paper is organized as follows. Section 2 describes our model. We then derive some general characteristics of the model analytically in Section 3, and in Section 4 we present two results concerning the dynamics of sales. This is followed by the description of the numerical solution and a discussion of the results in Section 5. Section 6 concludes.

[7] Other studies which report similar findings include Kashyap, Stein, and Wilcox (1993), Kashyap, Lamont, and Stein (1994), Guariglia (1996) and Guariglia and Schiantarelli (1995).

2 The Model

There are infinitely many time periods $t = 0, 1, 2, ...$ We consider a firm that produces a good for which, in each period, demand is random. The shock which governs the demand distribution in each period does not hinge on the price. In fact, we assume that demand is infinitely price-elastic, such that the goods sell at a given price p_t in period t. To facilitate the analysis further, we keep the sequence of prices constant, i.e. $p_t = p \; \forall \in \mathbb{N}$.

For notational reasons we denote the period-t demand realisation by z_{t+1}. In each period t the demand shocks are distributed on a subset $\mathcal{Z} \subseteq \mathbb{R}_+$, according to the probability density function $\phi : \mathcal{Z} \to \mathbb{R}_+$. The associated cumulative distribution function is denoted by $\Phi : \mathcal{Z} \to [0, 1]$. For convenience we define $\underline{z} = \inf \mathcal{Z} \geq 0$, and $\bar{z} = \sup \mathcal{Z}$ (which may be infinity). By parametrising ϕ with the demand shock from the previous period, we can allow for first-order serial correlation in the demand shocks. In this case the probability density functions are given by $\phi(z_{t+1}|z_t)$ and the cumulative distribution functions by $\Phi(z_{t+1}|z_t)$, for all t. We assume that if there exists first-order serial correlation in demand, it is positive.

At the beginning of period t, $t \geq 1$, the firm 'inherits' a non-negative stock of goods $g_t \in \mathcal{G} \subseteq \mathbb{R}_+$ and a non-negativeg9 that4t1sDhatefirsbutf189.99(th)980TD0Tc(z)01Tf134m0

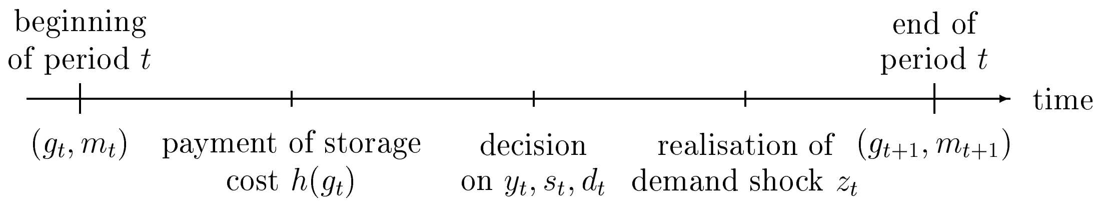

Figure 1: Timing of events

firm obtains only $p_s < p$ per unit of the good.

Next, the firm decides how to allocate the remaining funds (if any) between financing new production, building monetary reserves (precautionary savings), and paying dividends. Production y_t in period t is simply added to the stock of goods g_t carried over from the previous period. It leads to the firm providing a target inventory $n_t = g_t + y_t$ for sale in period t. The liquid funds retained are given by s_t, and we assume that these funds earn a safe gross rate of return of R. Finally, we denote by d_t the dividends paid out in period t.

Production costs are represented by a strictly convex and twice differentiable function $c : \mathbb{R}_+ \to \mathbb{R}_+$, which maps a non-negative production output y_t into the set of non-negative real numbers, with $c(0) = 0$, $c'(y_t) > 0$ and $c''(y_t) < 0$ for all $y \in \mathbb{R}_+$. Furthermore we assume that the average production costs exceed the scrap value at any level of production, that is $\frac{c(y_t)}{y_t} > p_s$ for all $y_t \in \mathbb{R}_+$ and all t.

This sequence of events in period t is summarised in Fig. 1.

When making its its reserving decisions, s_t, and its dividend decisions, d_t, the firm is subjected to the financing constraints,

$$s_t \geq 0 \tag{2.1}$$

$$d_t \geq 0 \tag{2.2}$$

for all $t \in \mathbb{N}$.

Moreover, each period it faces the budget constraint

$$h(g_t) + c(\max\{0, y_t\}) + p_s \min\{0, y_t\} + s_t + d_t \leq m_t. \tag{2.3}$$

Note that there is no non-negativity constraint on production y_t. As suggested by the non-negative domain of the cost function c, we do not allow the firm to reverse-engineer products into cash. However, it can sell goods at scrap value, which is interpreted as negative production in this paper.[8] Moreover, our definition of production is consistent with that used predominantly in the inventory literature (see, e.g., Blinder and Maccini (1991)). Since production is not directly observable in most firm data, it is defined empirically to be the change in inventories stocks plus sales.

The non-negativity constraint on dividends prevents the firm from raising equity capital. The non-negativity constraint on liquid reserves is a simple borrowing constraint.[9]

Once production has taken place the demand realisation, z_{t+1}, occurs. As this is a stockout-avoidance model we also impose a non-negativity constraint on the stock of goods. That is, the firm is not allowed to sell short output, or

$$x_{t+1} = \min\{z_{t+1}, n_t\}, \tag{2.4}$$

where $n_t = g_t + y_t$, the total amount of goods the firm makes available for sale.

The stock of cash (financial reserves) in period $t+1$ is the sum of period-t savings and returns on the latter, Rs_t, as well as the revenue generated from sales. Thus, the law of motion for cash holdings stock is

$$m_{t+1} = s_t + px_{t+1}, \tag{2.5}$$

The law of motion for the stock of goods is

$$g_{t+1} = g_t + y_t - x_{t+1}, \tag{2.6}$$

which can be rewritten as $g_{t+1} = n_t - x_{t+1}$.

Using (2.6), observable production is defined to be

$$y_t = g_{t+1} - g_t + x_{t+1}, \tag{2.7}$$

which again can be positive, zero or negative. Alternatively, can write

$$y_t = n_t - n_{t-1} + x_t.$$

[8]Alternatively, one could treat scrap sales simply as additional sales, the revenue of which would then be added to the right-hand side of the budget constraint. This has no consequences for the decisions made by the firm.

[9]Alternatively, it could be interpreted as restricting the firm's access to trade credit.

The analysis below is simplified considerably if we restrict the rate of return on cash to be smaller than the rate of time preference. The impact of the financing constraint would not be affected by setting $\beta R = 1$. However, our assumption that $\beta R < 1$ avoids technical problems related to indeterminacies of optimal solutions and to the unboundedness of return functions.

The objective of the firm is to choose production and savings to maximise the present discounted value of dividends, i.e.

$$\max E_0 \sum_{t=0}^{\infty} \beta^t d_t \tag{2.8}$$

subject to (2.1), (2.2), (2.3), (2.4), (2.5), (2.6), with the discount factor $\beta \in (0,1)$, where

$$d_t = m_t - h(g_t) - c(\max\{0, y_t\}) - p_s \min 0, y_t - s_t \geq 0,$$

for all t.

3 Characteristics of Value and Policy Functions

In this section we identify a number of properties that can be derived for the general model formulated in the previous section.

The choice of control variables y_t and s_t in each period t merely depends on the value of current state variables g_t and m_t, regardless of how the firm arrived there. Hence, the problem exhibits a recursive structure, which allows us to formulate the Bellman equation associated with the sequence problem outlined in Section 2. In each period t,

$$v(g_t, m_t, z_t) = \max_{y_t, s_t} \quad m_t - h(g_t) - c(\max\{0, y_t\}) - p_s \min\{0, y_t\} - s_t$$

$$+\beta \int_{\underline{z}}^{g_t + y_t} v(g_{t+1}, m_{t+1})\, d\Phi(dz_{t+1}|z_t) + \beta \int_{n_t}^{\bar{z}} v(0, m_{t+1})\, d\Phi(dz_{t+1}|z_t), \tag{3.9}$$

subject to (2.1), (2.2), (2.5), and (2.6).

Before examining the full problem we will consider the benchmark case that arises from relaxing constraint (2.2). We will refer to the resulting model as the unconstrained problem.

3.1 The Unconstrained Problem

In order to allow the firm to perfectly insure its desired production expenditure against negative demand shocks, it is sufficient to drop the non-negativity constraint on dividends d_t. If we did not retain the borrowing constraint (2.1), the firm's optimal policy would be to incur as much debt possible in the very first period. This is due to the return function being linear in m_t and the assumption that $\beta R < 1$.

With (2.1) still in place, the firm can simply raise equity finance if cash holdings are not sufficient to cover storage costs. In fact, it will always be necessary to do so whenever the latter exceed sales revenues. The reason is that optimal cash holding $s^\star(g, m, z) = 0$ in each period and for each state (g, m, z), because $\beta R < 1$.

To obtain the optimal production choice we formulate the modified Bellman equation

$$v(g, m, z) = \max_y m - h(g) - c(\max\{0, y\}) - p_s \min\{0, y\} \tag{3.10}$$

$$+\beta \int_{\underline{z}}^{g+y} v(g + y - z', pz') d\Phi(z'|z) + \beta v(0, p(g + y))[1 - \Phi(g + y|z)].$$

where we have already incorporated the laws of motion and the fact that, in equilibrium, $s^*(g, m, z) = 0$.

Differentiating with respect to y and using the envelope conditions, $v_g = -h'(g)$ and $v_m = 1$, we obtain

$$-c'(\max\{0, y\}) - p_s \min\left\{0, \frac{y}{|y|}\right\} - \beta h \Phi(g + y|z) + \beta p[1 - \Phi(g + y|z)] = 0 \tag{3.11}$$

as the Euler equation.

First, note that if g increases due to a drop in z by the same amount, $\Phi(g+y|z)$ increases. Optimality then requires y to drop.

Taking into account that $n = g + y$, the Euler equation can also be written in the form

$$-c'(\max\{0, n-g\}) - p_s \min\left\{0, \frac{n-g}{|n-g|}\right\} - \beta h \Phi(n|z) + \beta p[1 - \Phi(n|z)] = 0. \tag{3.12}$$

Viewed from this angle, we see that the optimal response to the same increase in g as before is an increase in n. Combining the two insights, it follows that a decrease in sales by one unit, which raises g by the same amount, leads to a decrease in y by *less* than one unit (as otherwise n could not have increased). Conversely, a increase in sales by one unit, which lowers g by the same amount, leads to a increase in y by *less* than one unit.

As a consequence, we obtain the classic production-smoothing result. In other words, due to the assumption of increasing marginal costs, production is less variable than sales in the absence of financing constraints that are relevant for a firm's production decisions.

3.2 The Constrained Problem

We now turn back to the original problem in which it is not possible to raise funds through negative dividends. This implies that there is an inequality constraint on the return function, which may be binding in some periods but not in others. In other words, optimal choices will not always be interior solutions, and we can therefore not proceed in the same way we did in the previous subsection.

We will, however, derive a few general properties of the solution. These properties ensure that the problem is well defined and provide insights as to the nature of the solution when financing constraints are imposed.

Property 1. The Bellman operator defined in (3.9) is a contraction mapping with a unique fixed point.

Due to the non-negativity constraints (2.1) and (2.2), the state subspace $\mathcal{G} \times \mathcal{M}$ for this problem is bounded below by zero. We will demonstrate in Proposition 3.1 below that this subspace is also bounded above by some finite values $\bar{g}$ and $\bar{m}$. Under this assumption it is straightforward to show that the conditions for Lemma 9.5 and Theorem 9.6 in Stokey, Lucas, and Prescott (1989) (pp.261-64) are satisfied. Thus, the Bellman operator defined on the right-hand side of (3.9) is a contraction mapping with a unique fixed point.

Property 2. The value function v is increasing in m and g.

Because of the law of motion for money (2.5), the one period return to the firm, and the feasibility constraints are all increasing in the state variable, m, v is also increasing in m by Theorem 9.7 in Stokey, Lucas, and Prescott (1989) (p.264). Similarly, since the law of motion for goods (2.6), and the firm's one-period return are also increasing in g, v is increasing in g as well.

Property 3. The value function v is concave and strictly concave in states where the financing constraints are binding.

Proof: See Appendix.

The following propositions and corollaries characterise optimal policy functions. They will prove useful in deriving a number of statements about the distribution of sales and net production later on. To derive we will simply assume that a stationary equilibrium exists.

PROPOSITION 3.1: *There exist upper bounds to the optimal target inventory, n, and the optimal cash retention, s, referred to as $\bar{n}$ and $\bar{s}$, respectively. If $m - \bar{s} > h(g) - c(\max\{0, y\}) - p_s \min\{0, y\}$ the remaining money is paid out as a dividend.*

Proof: Suppose the firm never pays dividends. Since $v(g, m, z)$ is then strictly concave in the first two arguments , the marginal returns to savings and to production will be decreasing in both state variables. The value of an infinitesimal unit of cash converges to βR as $m \to \infty$, since the probability of encountering a sequence of sales realisations in which the non-negativity constraint on dividends will be binding goes to zero. (Effectively, the firm becomes unconstrained.) Since $\beta R < 1$, there will be a level of cash retention, $\bar{s}$, beyond which the firm will prefer paying dividends over building further cash reserves. Similarly, as $g \to \infty$ the marginal value of inventories will eventually become smaller than the marginal cost, since the probability of selling the marginal unit goes to zero. Thus, there also exists a cap on the amount goods the firm puts up for sale, $\bar{n}$, beyond which the firm prefers paying dividends over further production. The existence of $\bar{s}$ and $\bar{n}$ implies the existence of endogenous upper bounds $\bar{g}$ and $\bar{m}$ to the state subspace $\mathcal{G} \times \mathcal{M}$ (see Property 1 above). □

COROLLARY 3.2: *There exist optimal policy functions, denoted by $s^\star(g, m, z)$ and $n^\star(g, m, z)$, which map each element of the state space into the space of feasible actions.*

Proof: The one-period return function is concave in the state variables, and the feasibility set for the choice variables is convex. In those states where no dividends are paid out, v is strictly concave in g and m. It follows that the maximum in (3.9) in this region is attained by unique choices of s_t and n_t. From Proposition 3.1 it follows that when dividends are positive, $s_t^* = \bar{s}$ and $n_t^* = \bar{n}$. □

COROLLARY 3.3: *The mean of production, $\hat{y}$, equals the mean of sales, $\hat{x}$.*

Proof: Assume $\hat{y} > \hat{x}$ is true. Then the firm would accumulate inventories indefinitely, contradicting the existence of $\bar{n}$.

Conversely, if $\hat{y} < \hat{x}$ were true, then $g_t \to 0$ as $t \to \infty$, which cannot be optimal since $v(g, m, z)$ is increasing in g. □

Intuitively speaking, this corollary holds because each unit of output that the firm produces is either sold on the 'regular' market, or it is given away at scrap value, which is accounted for as negative production.

The next proposition demonstrates that after an increase in either stock g or cash m held at the beginning of a period, the firm will never reduce its choice of target inventory n or reserving s. Thus, in a sense, n and s weakly increase with the operating capital of the firm.

PROPOSITION 3.4: *The policy functions $n^*(g, m, z)$ and $s^*(g, m, z)$ are non-decreasing in financial slack, for all $z \in \mathcal{Z}$.*

Proof: See Appendix.

We end this section with an additional observation to provide some intuition for the main results below. In the range of sales that do not lead the firm to stock out, an increase in previous-period sales x_- never causes a reduction in the firm's 'overproduction', defined as $y - x_-$. Put differently, in response to the increase in x_-, production y is augmented by at least as much as the increase in x_-. This is a consequence of Proposition 3.4. If one increases sales by one unit, the firm will produce at least one unit more to replace it, and it has additional funds to retain or expand production further.

4 Some Further Analytical Results

Due to the presence of financing constraints it is not straightforward to find analytical answers to some of the questions that are of interest to us here. This is true in particular for the two main stylised facts that have emerged from the empirical literature, as outlined in the introduction. Therefore, we analyse the model numerically in the next section.[10] Before doing so, however, we present two results regarding the dynamics of sales, which are less challenging to prove. This is followed by some reflections on the nature of the financing constraints.

4.1 The Dynamic Behaviour of Sales

Let us first assume that in equilibrium the covariance of production and sales exceeds the variance of sales. The next proposition demonstrates that under these circumstances, there is positive correlation between sales and changes in inventories.

PROPOSITION 4.1: *If the covariance between production and sales is larger than the variance of sales, changes in the stock of inventory put up for sale, Δn_t, covary positively with sales, x_t.*

Proof: Writing out the expression for $cov(\Delta n, x)$ we obtain

$$\begin{aligned} cov(\Delta n, x) &= cov(n, x) - cov(n_{-1}, x) \\ &= cov(g + y, x) - cov(n_{-1}, x) \\ &= cov(y, x) + cov(n_{-1} - x, x) - cov(n_{-1}, x) \\ &= cov(y, x) - var(x)\,, \end{aligned}$$

which is positive by assumption. □

Another interesting result is that the model generates positive serial correlation in sales despite the fact that the underlying demand process i.i.d.

PROPOSITION 4.2: *Sales exhibit positive first-order autocorrelation.*

Proof: Defining X as the space of sales realisations and $\psi(x)$ as the probability density on X, it needs to be shown that

$$\int_X (x - \hat{x}) \left[\int_X (x' - \hat{x}) \psi(x'|x) dx' \right] \psi(x) dx > 0, \tag{4.13}$$

where x and x' refer to the sale realisation in the current and the next peridod, respectively, and $\psi(x'|x)$ to the probability density of x', conditional on the current-period sales level being x. The left-hand side of inequality (4.13) can be simplified

[10] The reader may want to consult Brown and Haegler (1999) for an analytical proof of the 'excess-volatility-of-production' result.

to

$$\int_X \left[\int_X (x' - \hat{x}) \psi(x'|x) dx' \right] \psi(x) dx, \tag{4.14}$$

It follows from Proposition 3.4 that an increase in the sales realisation leads to a target inventory which is at least as high as that without the increase. Therefore, the conditional mean, $\xi(x) \equiv \int_X (x' - \hat{x}) \psi(x'|x) dx'$, is nondecreasing in x. Moreover, since the firm does not always put up $\bar{n}$ for sale, it must be true that $\xi(\bar{n}) > 0$. This implies that there is a sale, $\omega \in [\underline{z}, \bar{n}]$ such that for all $x > \omega$ the conditional expectation of next period's differential between sales and the unconditional mean of sales is positive, i.e. $\xi(\bar{n}) > 0$. Note that positive weight is attached to higher values of x and the same negative weight (in absolute terms) to lower values of x. Hence, (4.14) must be positive. □

4.2 The Nature of Financing Constraints

To what extent do the results in this paper depend on the type of financing constraints used? Here we offer only informal arguments that our results will not alter under different assumptions about the financing constraints. The basis for our argument is that regardless how the financing constraint is modeled, it will be binding primarily after low sales realisations. Thus 'underproduction' will still be associated only with low sales realisations, which is the central feature of the results in this paper.

Specifically, consider two different ways to model the financing constraints: (i) a supply of external finance which is perfectly elastic, but at a higher cost than internal finance; (ii) the firm is able to enter information-constrained insurance contracts.

The first case, where the firm faces a hierarchy of finance, is closest to our model. Indeed, the financing constraints in our model could be interpreted as representing a finance hierarchy in which the premium on external funds is so high that it is never optimal for the firm to use external finance. Suppose instead that the premium were not so high. In this case there will be three different regions. In the first region, internal funds will be so low that the benefit of the marginal good put up for sale is high enough to warrant the use of external finance. In the second region the firm is still constrained, but the value of the marginal good put up for sale is not high enough to justify the use of external funds. In this region the firm behaves exactly as it would in our model. In the third region the firm is unconstrained. Note that, because external finance is more costly, the optimal amount of goods put up for sale will be lower when the marginal unit is financed externally than when the marginal unit is financed with internal funds. Thus, if the firm begins with the unconstrained amount of goods for sale and has a low sales realisation it will underproduce. In other words, underproduction will still be associated with low sales realisations.

The second case is further from our model. Thus, our argument is only suggestive. Suppose that lenders can observe inventories at only two points in time: after sales *and* payment of storage costs, and just before sales. In other words, lenders can observe the firm's production decision, y_t. However, lenders cannot observe sales and storage costs. In this setup there is scope for the firm to report a bad sales realisation when in fact there has been a good sales realisation. Typically, the optimal contract in such a setup places restrictions on the observable decision variable in order to ensure that the borrower truthfully reports good outcomes. This usually implies that in bad outcomes the level of the decision variable is set lower than it would be if all variables were observable in order to introduce a cost to misreporting a good outcome as a bad outcome. In the information constrained framework, the optimal constract would likely force the firm to set production lower for low sales realisations. Thus, this informal argument suggests that the association between low sales and underproduction would remain, thereby preserving the excess variance of production result.

5 Numerical Analysis

First we describe how the general model of the previous sections is adapted for the purpose of numerical computations. This encompasses the specification of the production-cost and storage-cost functions, as well as of the distribution of the demand shocks. Moreover, a number of variable transformations are required to ensure that state and choice spaces are compact, and we need to formulate penalty functions to represent the financing constraints.

The production-cost function c is assumed to be of the quadratic type, that is

$$c(\max\{0, y\}) = c_1 \max\{0, y\} + \frac{c_2}{2} \max\{0, y\}^2,$$

with parameters $c_1 > 0$ and $c_2 > 0$. For the sake of simplicity we impose linearity on the storage-cost function h, such that, with a slight abuse of notation

$$h(g) = hg.$$

We assume the demand shocks to be log-normally distributed. Thus, the realisation of demand in period t, z_{t+1}, is determined by the autoregressive process

$$\ln z_{t+1} = \rho \ln z_t + \epsilon_t,$$

where ϵ_t is a serially uncorrelated, normally distributed random variable with zero mean and variance σ^2, and $\rho \in [0, 1)$ is the autocorrelation coefficient.

As the support of z_{t+1} is not compact, we transform the demand shock into a new variable

$$Z_{t+1} \equiv \frac{1}{2}\left[\tanh(\ln z_{t+1}) + 1\right].$$

This implies that Z_{t+1} is distributed on the unit interval.

Note that for a target inventory $n \equiv g + y$, the level of ϵ at which the firm just stocks out, denoted by $\bar{\epsilon}$, is given by the equation

$$\rho \ln z + \bar{\epsilon} = \tanh^{-1}(2n - 1),$$

or

$$\bar{\epsilon} = \ln \sqrt{\frac{n}{1-n}} - \rho \ln z.$$

Furthermore, we transform ϵ and $\bar{\epsilon}$ themselves into $\nu = \frac{\epsilon}{\sigma\sqrt{2}}$ and $\bar{\nu} = \frac{\bar{\epsilon}}{\sigma\sqrt{2}}$, respectively. As ν has a probability density function equal to $e^{-\nu^2}$, this allows us to apply directly the Gauss-Hermite quadrature rule in the numerical integration below. It follows that the law of motion for the transformed demand-shock variable is

$$Z'(Z, \nu) = \frac{1}{2}\left[\tanh\left(\rho \ln \frac{Z}{1-Z} + \sigma\sqrt{2}\nu\right) + 1\right].$$

The problem of a non-compact state space arises with the cash state variable, too. Therefore, we also transform cash holdings m to $M \equiv \tanh(\ln m)$ to ensure that (transformed positive) cash holdings are confined to the interval $[-1, 1]$. For a given M, the actual cash holding is

$$m = \exp(\tanh^{-1}(M)) = \sqrt{\frac{1+M}{1-M}}.$$

Given state (g, m, Z) and choices (y, s), we can now write the return function as

$$\Pi(g, m, Z|y, s) = m - hg - \left[c_1 \max\{0, y\} + \frac{c_2}{2} \max\{0, y\}^2\right] - p_s \min\{0, y\} - s.$$

As for the financing constraints we follow McGrattan (1993) in using the penalty-function approach to represent them numerically. That is, in order to capture the non-negativity constraint on dividends we will modify the Bellman by adding the penalty function

$$\frac{\gamma}{3} \min\{0, \Pi(g, m, Z|y, s)\}^3 \tag{5.15}$$

to the return function, for some 'large' parameter $\gamma > 0$.[11] For the non-borrowing constraint $s \geq 0$ we do the same by adding

$$\frac{\xi}{3} \min\{0, s\}^3,$$

for some 'large' parameter ξ.

This procedure results in the modified Bellman equation

$$v(g, m, Z) = \max_{y,s} \Pi(g, m, Z|y, s) + \frac{\gamma}{3} \min\{0, \Pi(g, m, Z|y, s)\}^3 + \frac{\xi}{3} \min\{0, s\}^3$$

$$+\frac{\beta}{\sqrt{\pi}} \int_{-\infty}^{\bar{\nu}} v(n - Z'(Z, \nu), Rs + pZ'(Z, \nu))e^{-\nu^2} d\nu + \frac{\beta}{\sqrt{\pi}} \int_{\bar{\nu}}^{\infty} v(0, Rs + p\nu)e^{-\nu^2} d\nu.$$

Differentiating with respect to y and s, and making use of the envelope conditions yields Euler equations

$$-\left[c_1 \max\left(0, \frac{y}{|y|}\right) + c_2 \max(0, y) + p_s \min\left(0, \frac{y}{|y|}\right)\right] \left[1 + \gamma \min\{0, \Pi(g, m, Z|y, s)\}^2\right]$$

$$-\frac{\beta}{\sqrt{\pi}} \int_{-\infty}^{\bar{\nu}} h \left[1 + \gamma \min\{0, \Pi(n - Z'(Z, \nu), Rs + pZ'(Z, \nu)|y', s')\}^2\right] e^{-\nu^2} d\nu$$

$$+\frac{\beta}{\sqrt{\pi}} \int_{\bar{\nu}}^{\infty} p \left[1 + \gamma \min\{0, \Pi(0, Rs + p\nu|y', s')\}^2\right] e^{-\nu^2} d\nu = 0,$$

and

[11] Typically, values of 10^3 or 10^4 are chosen.

$$
-\left[1+\gamma \min\{0, \Pi(g, m, Z|y, s)\}^2 - \xi \min\{0, s\}^2\right]
$$

$$
+\frac{\beta}{\sqrt{\pi}}\int_{-\infty}^{\bar{\nu}} R\left[1+\gamma \min\{0, \Pi(n-Z'(Z,\nu), Rs+pZ'(Z,\nu), Z'|y', s')\}^2\right] e^{-\nu^2} d\nu
$$

$$
+\frac{\beta}{\sqrt{\pi}}\int_{\bar{\nu}}^{\infty} R\left[1+\gamma \min\{0, \Pi(0, Rs+pn, Z'|y', s')\}^2\right] e^{-\nu^2} d\nu = 0,
$$

respectively.

To find the policy functions $y^*(g, m, Z)$ and $s^*(g, m, Z)$ that solve this system of equations, we apply the finite-elements method (see, e.g. Reddy (1992),McGrattan (1993) for conceptual details). MATLAB is the software package used to implement the computations.[12]

Note that we can replace next period's choices y' and s' above by

$$
y' \equiv y^*[g'(g, m, Z), m'(g, m, Z), Z'(Z, \nu)]
$$

and

$$
s' \equiv s^*[g'(g, m, Z), m'(g, m, Z), Z'(Z, \nu)],
$$

respectively.

The main objective of the paper is to show that our model with financing constraints can generate an outcome consistent with the two stylised facts mentioned in the introduction. In order to bias our numerical analysis against this hypothesis, we rule out the possibility of positive serial correlation in demand, and compute the model with i.i.d. shocks only, i.e. for $\rho = 0$. This has the added benefit that there are only two state variables, g and m. We do not claim that serial correlation in demand shock does not play a role in generating the empirical observations, but we deem it preferable to isolate the effect of financing constraints from other influences.

The first set of parameters used is given in Table 1. The resulting policy functions for y, n and s are depicted in Figures 2, 3 and 4, respectively. The graph for the value function can be viewed in Figure 5. Inspection of the value function appears to confirm our result from Section 3 that it is strictly concave at lower levels of g and m, where the firm is likely to be financially constrained, and to become linear at higher levels (unconstrained region).

[12] The programming code is available upon request from u.haegler@rhul.ac.uk.

Parameter	Value
β	0.95
p	2
c_1	1
c_2	0.7
p_s	0.8
h	0.8
R	1
σ	1.5

Table 1: Parameters for Constrained Model 1)

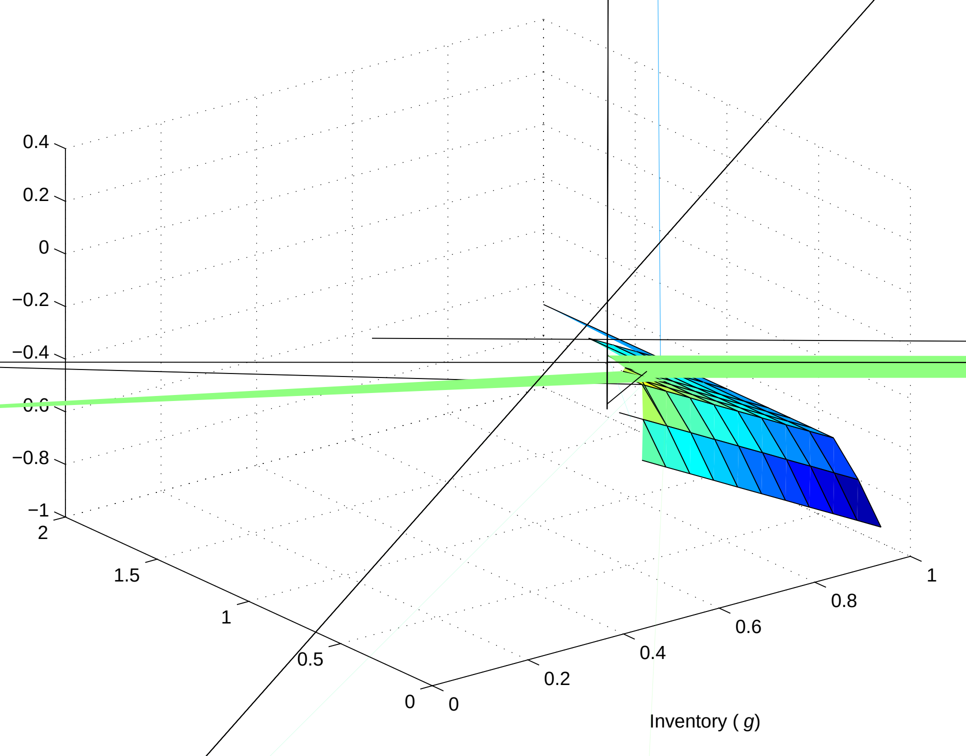

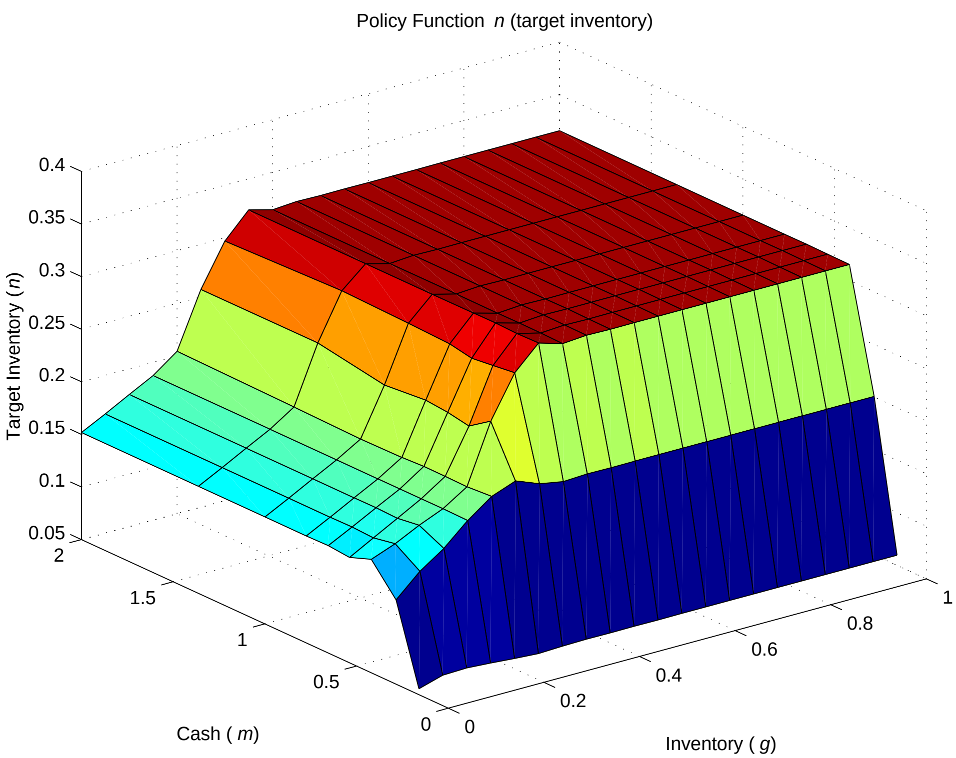

Figure 3: Policy Function for Target Inventory n

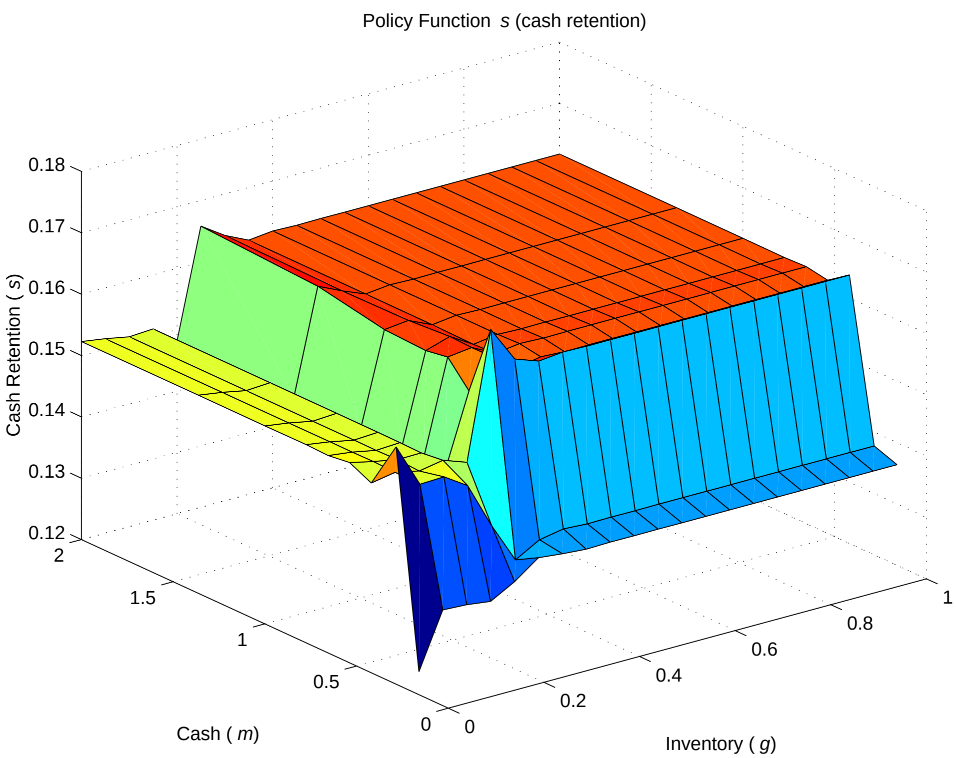

Figure 4: Policy Function for Cash Holdings s

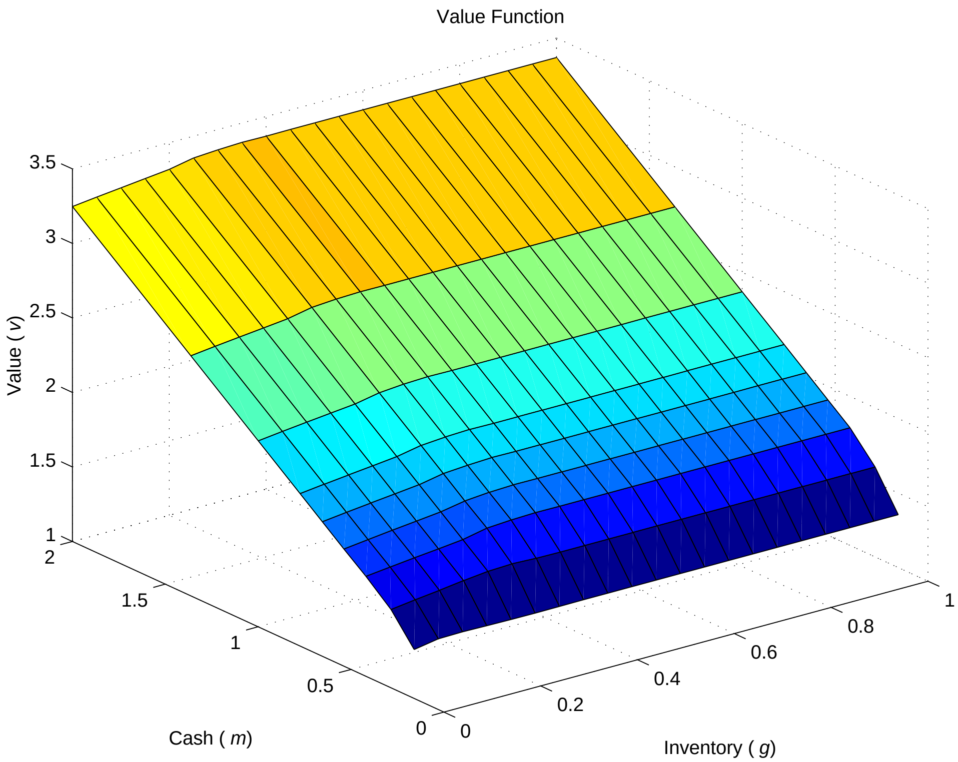

Figure 5: Value Function

Measure	Value
Mean of production y	0.1240
Mean of sales x	0.1251
Var(y)/Var(x)	1.8379
1st order autocorrelation coefficent of x	0.3648
1st order autocorrelation coefficent of $n-\theta x$	0.2507
2nd order autocorrelation coefficent of $n-\theta x$	0.1603
correlation coefficent between n and x	0.3810
correlation coefficent between s and x	0.3465

Table 2: Simulation Results for Constrained Model 1

The next stage of the numerical analysis is to use the policy functions and a random sequence of demand-shock realisations to generate time series for output, target inventories and sales. The results of these simulations are reported in Table 2.

A number of insights can be gained from these results. Firstly, we note that the mean of production is virtually the same as that for the mean of sales, which confirms Corollary 3.3. Secondly, we find that the model is consistent with the stylised fact that production is more volatile than sales, as expressed by the ratio of the two variances. Thirdly, there is positive serial correlation in sales. This is not surprising, as in the previous section we were able to prove this result analytically.

The next observation is in line with the second significant stylised fact we are interested here, namely the persistence of the inventory-sales relationship. This is highlighted by the positive values of the first-order and second-order autocorrelation coefficients for the series $n_t - \theta x_t$, where $\theta \equiv \frac{\text{mean}(n)}{\text{mean}(x)}$.[13] It is true that the individual values, or even their sum, are not as high as those that are empirically, which tend to be well above 0.7. However, it has to be kept in mind that the simulations are based on i.i.d. demand shocks. It is to be expected that by allowing for sufficiently large positive serial correlation in demand shocks one would obtain values in that range as well.[14] A typical pattern for the inventory-sales relationship

[13] Ramey and West (1997) demonstrate that $n_t - \theta x_t$ is a stationary relationship if the sales process is non-integrated, which is the case here.

[14] We remind the reader that in order to generate the persistence result Ramey and West (1997) assume the demand shocks to be a unit-root process.

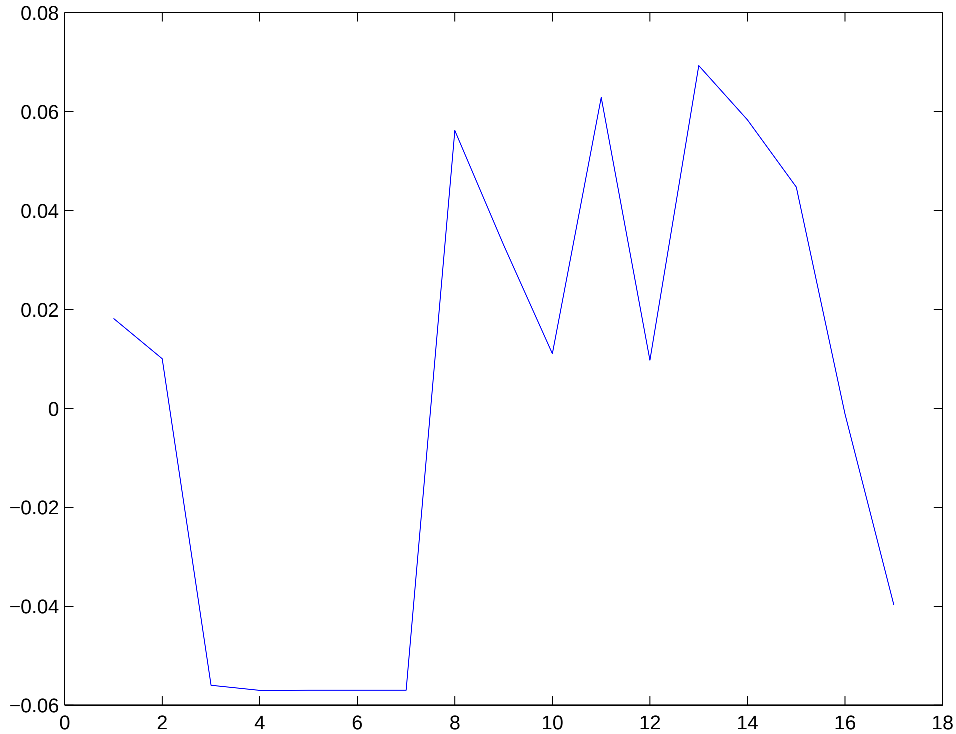

Figure 6: An Inventory-Sales Relationship Process

over time is depicted in Figure 6.

Finally, we point out the positive correlations between sales and (target) inventories (0.3810) as well as between sales and cash reserves (0.3465).

To verify that the financing constraint are indeed crucial in obtaining these results, we compare them with those the emerge from the simulation of a model where there they are absent. More precisely, we simulate a version of the unconstrained model outlined in Subsection 3.1, using the same parameters as before. The results of this exercise are given in Table 3.

First note that the last figure is zero because in the unconstrained model cash holdings are held constant at zero. Holding monetary reserves is unnecessary as funds can be raised more cheaply through equity if necessary.

Mean production and mean sales are again identical as expected, but they are higher than in the constrained model. The reasons are that the firm is never forced to produce less than desired, and that it is more confident to provide more

Measure	Value
Mean of production y	0.1759
Mean of sales x	0.1759
Var(y)/Var(x)	0.7943
1st order autocorrelation coefficent of x	-0.0793
1st order autocorrelation coefficent of $n-\theta x$	-0.0698
2nd order autocorrelation coefficent of $n-\theta x$	0.0079
correlation coefficent between n and x	-0.7546
correlation coefficent between s and x	0

Table 3: Simulation Results for Unconstrained Model

goods for sale on average. It is less worried about not stocking out and incurring high storage costs, so that no resources have to be allocated to precautionary cash holdings.

The variance of production is now smaller than the variance of sales, which means that there is production smoothing rather than production bunching.

We also lose the second stylised fact (persistence in the inventory-sales relationship). The first-order autocorrelation coefficient is negative and the second-order coefficient very close to zero.

Finally, in contrast to the constrained model there is negative serial correlation in sales here, and target inventories correlate negatively with sales.

Thus, as nothing else has changed in the model, it is clear that the presence of financing constraints is responsible for the results which to a large extent are completely the opposite of those obtained in the unconstrained model.

As a final experiment, we return to the constrained model and change one of the parameters. We examine the impact of a decrease in the (regular) price of output p from 2 to 1.8, leaving all other values unaltered. Table 4 displays the results of this new simulation.

It should come as no surprise that average output and sales are lower than under the higher price. However, qualitatively speaking, there are no changes to the picture that emerged from the first simulation. The ratio of variances still clearly exceeds 1, and all the correlation measures are still positive. The only

Measure	Value
Mean of production y	0.0833
Mean of sales x	0.0833
$\text{Var}(y)/\text{Var}(x)$	1.7004
1st order autocorrelation coefficent of x	0.3083
1st order autocorrelation coefficent of Δn	0.2257
2nd order autocorrelation coefficent of Δn	0.1314
correlation coefficent between n and x	0.5517
correlation coefficent between s and x	0.3081

Table 4: Simulation Results for Constrained Model 2

difference is that, with the exception of the correlation between n and x, they tend to be slightly lower than before.

6 Conclusions

This paper has presented a model of inventory investment which uses financing constraints to explain two empirical regularities of inventories: the procyclicality of inventory investment, and the positive serial correlation of the inventory-sales ratio. The model shows that then a firm is financially constrained, it may unable to produce as much as it has sold following particularly poor demand realisations. Under these circumstances, production is lower than sales. Conversely, when inventories are below their optimum level, a firm produces more than it sells after a relatively good demand realization in order to return to the optimum inventory level as quickly as possible. This behaviour implies that inventory investment is procyclical, and that the inventory-sales ratio exhibits positive serial correlation. These results do not stem from assumptions regarding the cost function. When the firm is unconstrained, neither empirical regularity is predicted by the model.

Equally interesting, these results are also not driven by assumptions of serially correlated exogenous shocks. In our model there is no cost shock, and the demand shock is assumed to be i.i.d. This is unlike other models in the literature, and suggests that inventory investment may have a more active role business cycles through its potential to transform or amplify exogenous shocks. A future avenue of research is to examine how financing constraints affect inventory investment in environments where exogenous shocks are serially correlated.

A Appendix

Proof of Property 3

That v is concave is standard and stems from the fact that the Bellman functional equation $v = Tv$ is a contraction mapping, with an operator T that preserves concavity. This is because integration in (3.9) preserves concavity and the return function is concave as well (see Stokey, Lucas, and Prescott (1989) [p.265] for details).

To show strict concavity in the constrained region we adopt the method of a concavity proof in Abel (1985).[15]

Consider two firms, labelled A and B, respectively. Suppose the non-negativity constraints on dividends are binding in their respective current inventory-cash states (g_A, m_A) and (g_B, m_B). Firm A's optimal choice is assumed to be (y_A, s_A) and that of firm B is (y_B, s_B). Define a function $\tilde{c}$ that combines the cost function with the revenue function from scrap sales, i.e. $\tilde{c}(y) \equiv c(\max\{0, y\}) - p_s \min\{0, y\}$. Note that $\tilde{c}$ is convex. Since financing constraints are binding we have

$$m_i = h(g_i) + \tilde{c}(y_i) + s_i, \tag{A.16}$$

$i = A, B$. Assume that $g_A \leq g_B$ and $m_A \leq m_B$, with at least one inequality holding strictly. Under binding financing constraints it then follows that $n_A < n_B$, that is the 'wealthier' firm is able to attain a higher target inventory.

Next, form a convex combination of the two firms is, that construct a new firm which is in state $(\theta g_A + (1-\theta)g_B, \theta m_A + (1-\theta)m_B)$. A policy choice of $(\theta y_A + (1-\theta)y_B, \theta s_A + (1-\theta)s_B)$ may not be optimal for this firm, but it is feasible. To show this one can apply the following argument. Feasibility requires

$$\theta m_A + (1-\theta)m_B \geq c(\theta y_A + (1-\theta)y_B) + h(\theta g_A + (1-\theta)g_B) + \theta s_A + (1-\theta)s_B.$$

Due to equations (2.3) and the fact that $y_i = n_i - g_i$, $i = A, B$, this can be rewritten as

[15] In that paper production occurs with a lag. Consequently, only the stock of inventories which the firm has at the beginning of the period are available for sale. Moreover, the stock of inventories is the only state variable, as firms do not hold cash variable and there is no serial correlation in demand shocks.

$$\theta[c(n_A - g_A) + h(g_A)] + (1-\theta)[c(n_B - g_B) + h(g_B)] \geq$$
$$c(\theta(n_A - g_A) + (1-\theta)(n_B - g_B)) + h(\theta g_A + (1-\theta) g_B).$$

Due to the convexity of $c(\cdot) + h(\cdot)$ in n_i and g_i, this condition is met (Jensen's inequality).

If the demand shock in the current period, z', is greater than both n_A and n_B, both firms stock out which leads to next period's states $(g'_A, m'_A) = (0, s_A + pn_A)$ and $(g'_B, m'_B) = (0, s_B + pn_B)$, respectively. The same is true, however, for the convex combination, whose new state would then be $(g'_\theta, m'_\theta) = (0, \theta(s_A + pn_A) + (1-\theta)(s_B + pn_A))$. As v is concave the value of the convex-combination firm next (and hence in the current) period, denoted by v_θ, is therefore not smaller than the convex combination of the two firms, given by $\theta v_A + (1-\theta) v_B$.

Similarly, if the demand shock z' lies below both n_A and n_B, none of the firms, including the convex-combination firm, will stock out. Next period's states are then $(g'_A, m'_A) = (n_A - z', s_A + pz')$ and $(g'_B, m'_B) = (n_B - z', s_B + pz')$ for firms A and B, respectively, and $(g'_\theta, m'_\theta) = (\theta n_A + (1-\theta) n_B - z', \theta s_A + (1-\theta) s_B + pz')$ for the convex-combination firm. Its next-period state is just the convex combination of (g'_A, m'_A) and (g'_A, m'_A). Because of the concavity of v, $v_\theta \geq \theta v_A + (1-\theta) v_B$ in this case as well.

If the demand shock z' lies between n_A and n_B, firm A stocks out but firm B does not. Their states next period are then $(g'_A, m'_A) = (0, s_A + pn_A)$ and $(g'_B, m'_B) = (n_B - z', s_B + pz')$, respectively. If $z' < \theta n_A + (1-\theta) n_B$, $(g'_\theta, m'_\theta) = (\theta n_A + (1-\theta) n_B - z', \theta s_A + (1-\theta) s_B + pz')$, and if $z' \geq \theta n_A + (1-\theta) n_B$, $(g'_\theta, g'_\theta) = (0, \theta(s_A + pn_A) + (1-\theta)(s_B + pn_B))$. In either case we can imagine that, in a first stage, shares θ and $1-\theta$ of demand z' are covered by divisions a and b, respectively, of the convex-combination firm, where divivsion a corresponds a proportion θ and division b a proportion $1-\theta$ of the entire firm. Its value would then just be the same as the convex combination of the values of firm A and B. However, in a second stage, excess demand from division a can be shifted to division b and (partially) satisfied there. This leads not only to higher overall revenues but also to a reduction in goods depreciation. Hence, with this demand constellation, $v_\theta > \theta v_A + (1-\theta) v_B$. In expected terms, the valuation of the convex-combination firm is therefore strictly higher than the convex combination of the values of A and B. □

Proof of Proposition 3.4.

We will first present the proof for $n^*(g,m,z)$. The statement clearly holds for states in which the non-negativity constraint on dividends is not binding, as in this region additional cash is simply used to pay dividends, without changing the amount $\bar{n}$ made available for sale. For states in which the non-negativity constraint on consumption is binding we prove the statement by contradiction. Note that in those states the return function is equal to zero, and we can therefore write

$$\begin{aligned} v(g,m,z) &= \beta \int_{\underline{z}}^{n^*(g,m,z)} v(n^*(g,m,z) - z', s^*(g,m,z) + pz', z') d\Phi(z'|z) \\ &\quad + \beta[1 - \Phi(n^*(g,m,z))] v(0, s^*(g,m,z) + pn^*(g,m,z), z'). \end{aligned} \quad \text{(A.17)}$$

Assume that in some (g,m,z) the optimal policy is to choose values (n,s). Also suppose that, for some small $\epsilon > 0$, the optimal policy in $(g, m+\epsilon, z)$ is $(N, S_\epsilon) \equiv (n-\mu, s+\epsilon+\tilde{\mu})$, for some small decrease in output $\mu > 0$, which reduces production costs by some $\tilde{\mu}$. This implies

$$\begin{aligned} &\int_{\underline{z}}^{N} v(N - z', S_\epsilon + pz', z') d\Phi(z'|z) + \int_{N}^{\bar{z}} v(0, S_\epsilon + pN, z') d\Phi(z'|z) > \\ &\int_{\underline{z}}^{n} v(n - z', s_\epsilon + pz', z') d\Phi(z'|z) + \int_{n}^{\bar{z}} v(0, s_\epsilon + pn, z') d\Phi(z'|z), \end{aligned} \quad \text{(A.18)}$$

where $s_\epsilon \equiv s + \epsilon$. But then the firm could do better by shifting funds $\tilde{\mu}$ from production to savings in state (g,m,z) as well, i.e. for $(N,S) \equiv (n-\mu, s+\tilde{\mu})$

$$\begin{aligned} &\int_{\underline{z}}^{N} v(N - z', S + pz', z') d\Phi(z'|z) + \int_{N}^{\bar{z}} v(0, S + pN, z') d\Phi(z'|z) > \\ &\int_{\underline{z}}^{n} v(n - z', s + pz', z') d\Phi(z'|z) + \int_{n}^{\bar{z}} v(0, s + pn, z') d\Phi(z'|z). \end{aligned} \quad \text{(A.19)}$$

Hence, (n,s) cannot be an optimal choice in state (g,m,z), which contradicts the initial assumption.

To show that (A.18) does imply (A.19), note first that for very small μ, the decrease of the first term of the left-hand side due to the change μ in the (upper) integral boundary is offset by the increase of the second term of the left-hand side due to the same change μ in the (lower) integral boundary.

Since v is strictly concave when the non-negativity constraint on consumption is binding, the stockout terms (without probabilities) of the two inequalities compare as follows:

$$v(0, S_\epsilon + pN, z') - v(0, s_\epsilon + pn, z') \leq v(0, S + pN, z') - v(0, s + pn, z'). \quad \text{(A.20)}$$

Moreover, conditional on not stocking out, reducing n in favour of s is at least as valuable at lower savings as it is at higher savings, which is due to storage costs. Hence, $\forall\ z' \leq N$,

$$v(N-z', S_\epsilon+pz', z')-v(n-z', s_\epsilon+pz', z') \leq v(N-z', S+pz', z')-v(n-z', s+pz', z'). \tag{A.21}$$

This is true since the change of stockout probability is the same on both sides of the inequality. This results in identical reductions in expected storage costs. This 'relief' on cash holdings has a weakly higher when the latter are lower (m instead of $m+\epsilon$, which is again due to the concavity of v.

The gross unit return on additional savings μ is 1 $\forall\ z$, whereas the gross unit return on an additional goods provision of μ is strictly smaller than 1 for at least some z that do not lead to a stockout.

An analogous argument can be made when there is a jump from state (g, m) to a state $(g+\epsilon, m)$, as this would be equivalent to the previous situation, in the sense that there is again an increase in the firm's operating capital.

As for the function $s^*(g, m)$ assume that there is a range of states over which the function is decreasing. The only way s^* could be decreasing in the constrained region is if there was an overproportionate increase of n in response to a small increase of funds available (be it in the form of more cash or of higher goods inventories). Due to strict concavity of the value function, however, the negative effect on the expected value of v next period due to a decrease in s would more than offset the positive effect of an increase in n. We conclude that such a policy cannot be optimal. □

References

ABEL, A. B. (1985): "Inventories, Stock-Outs and Production Smoothing," *Review of Economic Studies*, LII, 283–293.

BILS, M., and J. KAHN (2000): "What Inventory Behavior Tells Us About Business Cycles," *American Economic Review*, 90(3), 458–481.

BLINDER, A. S. (1986): "Can the Production Smoothing Model of Inventory Behaviour be Saved?," *Quarterly Journal of Economics*, CI, 431–453.

BLINDER, A. S., and L. J. MACCINI (1991): "Taking Stock: A Critical Assessment of Recent Research on Inventories," *Journal of Economic Perspectives*, 5, 73–96.

BRESNAHAN, T., and V. RAMEY (1994): "Output Fluctuations at the Plant Level," *Quarterly Journal of Economics*, 109, 593–624.

BROWN, W., and U. HAEGLER (1999): "Financing Constraints and Inventories: An Explanation of Why Production is More Volatile Than Sales," mimeo.

CARPENTER, R. E., S. M. FAZZARI, and B. C. PETERSEN (1994): "Inventory Investment, Internal-Finance Fluctuations, and the Business Cycle," *Brookings Papers on Economic Activity*, pp. 75–138.

EICHENBAUM, M. S. (1989): "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment," *American Economic Review*, 79, 853–864.

GERTLER, M., and S. GILCHRIST (1994): "Monetary Policy, Business Cycles, and the Behaviour of Small Manufacturing Firms," *Quarterly Journal of Economics*, 95, 309–340.

GUARIGLIA, A. (1996): "The Effects of Finacial Constraints on Inventory Investment: Evidence form a Panel of UK Firms," Discussion Paper, University of Essex.

GUARIGLIA, A., and F. SCHIANTARELLI (1995): "Production Smoothing, Firm Heterogeneity and Financial Constraints: Evidence from a Panel of UK Firms," mimeo, University of Essex.

Hall, G. J. (1996): "Non-convex Costs of Capital Utilization: A Study of Production and Inventories at Automobile Plants," mimeo.

Holt, C. C., F. Modigliani, J. F. Muth, and H. Simon (1960): *Planning Production, Inventories and Work Force*. Prentice-Hall, Englewood-Cliffs.

Kahn, J. A. (1987): "Inventories and the Volatility of Production," *American Economic Review*, 77, 667–679.

Kashyap, A. K., O. A. Lamont, and J. C. Stein (1994): "Credit Conditions and the Cyclical Behaviour of Inventories," *Quarterly Journal of Economics*, 109, 565–591.

Kashyap, A. K., J. Stein, and D. W. Wilcox (1993): "Monetary Policy and Credit Conditions: Evidence from the Composition of External Finance," *American Economic Review*, 83, 78–98.

McGrattan, E. (1993): "Solving the Stochastic Growth Model With a Finite Element Method," *Federal Reserve Bank of Minneapolis Staff Report*, no. 164.

Ramey, V., and K. West (1997): "Inventories," mimeo.

Reddy, J. (1992): *An Introduction to the Finite Element Method*. McGraw-Hill, New York.

Stokey, N. L., R. E. Lucas, and E. Prescott (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge.