Inventories in Motion: A New Approach To Inventories Over The Business Cycle*

Michael McMahon†

Date: February 15, 2014

Abstract

I propose an inventories-in-motion concept which represents a new approach to inventories over the business cycle. This channel has previously been ignored by macroeconomists. I build a general equilibrium business cycle model in which inventories arise naturally as a result of gaps between production of goods and their consumption as goods are distributed. These inventories are actively managed and adjusted to meet consumption and investment needs in the economy. Although conceptually very simple, I show that such inventory behaviour matches a number of stylised facts of aggregate inventories. Nonetheless, my model does not admit an important role for inventory management improvements in declining macroeconomic volatility in the last 30 years.

Keywords: Inventories; business cycle; Great Moderation.

JEL Codes:

*I acknowledge, without implicating, the advice and guidance of Francesco Caselli, Silvana Tenreyro and Francesco Giavazzi. I have also benefited from comments from Bernardo Guimaraes, Stephen Hansen, Rachel Ngai, Chris Pissarides, Morten Ravn, Kevin Sheedy and Alwyn Young, as well as many seminar participants. I am indebted to Wouter den Haan for his course which introduced me to numerical methods. Any remaining errors are, of course, my own.

†University of Warwick, CEPR, CAGE (Warwick), CEP (LSE), CfM (LSE) and CAMA (ANU). Email m.mcmahon@warwick.ac.uk

1 Introduction

A bottle of beer purchased by a consumer in a shop, or at a bar, was produced some time before this consumption. In fact, the traditional beer distribution chain involves four distinct steps. Beer is produced in the brewery and becomes a finished good (although it often gets transported in large trucks to a separate plant for bottling and labeling). The bottles or cans of beer get shipped to a beer distributor who allocates them to regional wholesalers. These wholesalers break up the shipments into smaller units and pass the beer onto the local retail units (bars and shops). Finally, the retailer provides the goods in smaller quantities for the consumer to enjoy. This process is not instantaneous - there are large lags at each stage of the process.[1]

Beer distribution is a key example used in the supply chain management literature to describe the importance of inventories in the distribution chain and has given rise to a supply-chain game used in the teaching of inventory management (see Forrester (1961) and Sterman (1989)). The reason is that from the moment it is produced and as it moves along the distribution chain, the beer is an inventory.[2] It is the management of such distribution inventories, or 'inventories-in-motion', that has received considerable attention by the supply chain field for the last 30 years (Gilmore 2008).[3]

In this paper, I develop a DSGE model of inventories based on such distribution lags; I show that such a model gives rise to macroeconomic inventory behaviour that matches the aggregate data. This is important because inventory adjustment has long been recognised as a major source of macroeconomic volatility (accounting for almost half of the volatility of GDP growth), and yet macroeconomic analysis of inventories have typically not focused on inventories arising from actions to optimize the flow of goods through the distribution chain. My first contribution, therefore, is to solve a general equilibrium model of inventories in which inventories arise naturally in the distribution chain between production and consumption. I assume that goods ordinarily take one period between production and consumption, but I allow the representative agent access to a technology that facilitates (costly) early delivery which overcomes

[1] In fact, Budweiser, as an example, having reduced the lags involved in distribution of its beer, now labels bottles according to when the beer was produced - this is their "Born On" date. The lag between the "Born On" date and the date it is available for purchase is usually between four and ten weeks, though it can be up to a year.

[2] In fact, even before it is produced, beer gives rise to inventory holdings. The hops, malt barley, yeast, rice and water that make up the ingredients of the beer are likely to be held as input inventories by the brewery. I do not examine input inventories in this paper, though the analysis in this paper could be extended to consider such inventories.

[3] As Copacino (1988) says, and the inspiration for the title of this paper: 'Over the past decade, we have witnessed profound changes in many aspects of logistics. None of the developments, however, has been as striking as the recent trend toward managing 'inventory in motion' - that is, managing inventory while it is still in transit instead of waiting until it arrives at the warehouse'.

these distribution lags; such early delivery is used to smooth consumption.

I begin in Section 2 with a brief review of the importance of inventories in macroeconomic analysis, as well as the existing literature. I show that despite the importance of inventories to the macroeconomy, there is no canonical model of their behaviour but rather a number of different approaches to model them. Each of these models is judged against a number of key stylised facts of inventory behaviour. I do not propose that my inventory-in-motion model is a canonical model. Rather, my aim is to develop a reasonably simple model which matches the aggregate data, and particularly one that captures the consequential nature of some of the distribution chain inventories associated with the inventory-in-motion concept. This is important because such inventories form a large part of inventory analysis in other fields such as Management Science and Operations Research and yet have received much less attention from macroeconomists.

To illustrate the simple mechanism of inventory-in-motion, in Section 3 I explain the basic modelling concept within the context of a 3-period economy without, for simplicity, capital and labour supply decisions. I then embed the inventories-in-motion concept in an infinite-horizon DSGE model with elastic labour supply and costly adjustment of capital (see Section 4). Although conceptually simple, the main challenge of the DSGE model is that it involves occasionally-binding non-negativity constraints which make the model an unsuitable candidate for solution via log-linearisation. I, instead, use an extended version of the Parameterised Expectations Algorithm (PEA) pioneered by Wright and Williams (1982, 1984), and which is ideal for dealing with potentially binding constraints (Christiano and Fisher 2000).

An objection to the distribution cycle I propose is that while the typical unit of time in business cycle analyses is one quarter (driven mainly by data availability), firms rarely have a distribution cycle that is three months long. I follow the literature that examines the effects of time aggregation on macroeconomic modelling, such as Aadland (2001) and Heaton (1993) and calibrate a higher frequency (monthly) model and then aggregate the data to explore the consequences at quarterly frequency.

I compare the quarterly behaviour generated by my model with the behaviour from an equivalently calibrated standard Real Business Cycle (RBC) model as well as a model with the distribution delays, but no early delivery technology. I find that the introduction of delays between production and consumption naturally generate pro-cyclical inventory behaviour and so, despite being simple and highly stylised, such inventories are able to match a number of the key facts about inventory behaviour at the macro level. However, if we just introduce the time delay in distribution, we also substantially reduce the volatility of output and especially hours. The control of inventories-in-motion increases the volatility of labour hours and also amplifies the procyclicality of inventory investment. These comparisons are discussed in Section 6.

More recent interest in inventories has also been influenced by the potential role played by inventory improvement techniques in the decline of the macroeconomic volatility that has occurred since at least the 1980s (this is the so-called 'Great Moderation', or 'Great Stability' in the UK).[4] Improved inventory management, facilitated by technological developments such as barcodes, scanners, Radio Frequency Identification Tags (RFID), and Electronic Data Interchange (EDI), has resulted in quicker and easier management of inventory both within and, importantly, between warehouses, and is one of many competing explanations for this increased stability (McCarthy and Zakrajsek 2007).[5] My second contribution, in Section 7, is to map the most salient features of these inventory management improvements into the baseline inventory model to assess the role played by improved inventory management in explaining the Great Moderation. Others, such as Khan and Thomas (2007b), McCarthy and Zakrajsek (2007) and Iacoviello, Schiantarelli, and Schuh (2009), have pursued this question but have used difference approaches and different motives for inventories. My model is built upon precisely the form of inventories affected by many of the technological improvements in supply chain management. I find, consistent with the earlier papers, that there is little evidence to support the idea that inventory management techniques were a driving force in reducing aggregate volatility. Nonetheless, these management improvements appear to have played an important role in matching other changes in the aggregate data such as a lower level of the inventory-sales ratio. These other changes are not generated by the 'good luck' explanation which emphasises reduced volatility of shocks hitting the economy. I also stress the important role played in the Great Moderation by the decline in the covariance between components of demand, sectors of the economy, or types of product in the economy.

2 Inventories and the Macroeconomy

The importance of private inventories in the behaviour of business cycles is well known (the survey article by Blinder and Maccini (1991b) contains the main the references in this regard). Defined for the sake of national accounting as 'materials and supplies, work in process, finished goods, and goods held for resale' (Bureau of Economic Analysis 2008), Figure 1 that despite making up, on average, less than 1% of nomi-

[4]See, for example, McConnell and Perez-Quiros (2000), Stock and Watson (2002), and Blanchard and Simon (2001) for US evidence, and Benati (2004) for UK evidence

[5]Two well-known inventory management approaches associated with such inventory-in-motion are the 'Just-in-Time production' approach and 'The Walmart Approach'.

nal GDP[6] and contributing only about 2% of average GDP growth (0.1pp), inventory investment has accounted for 43% of the volatility of real GDP growth.

Figure 1: Role of Inventories in the Macroeconomy: U.S. 1960-2007

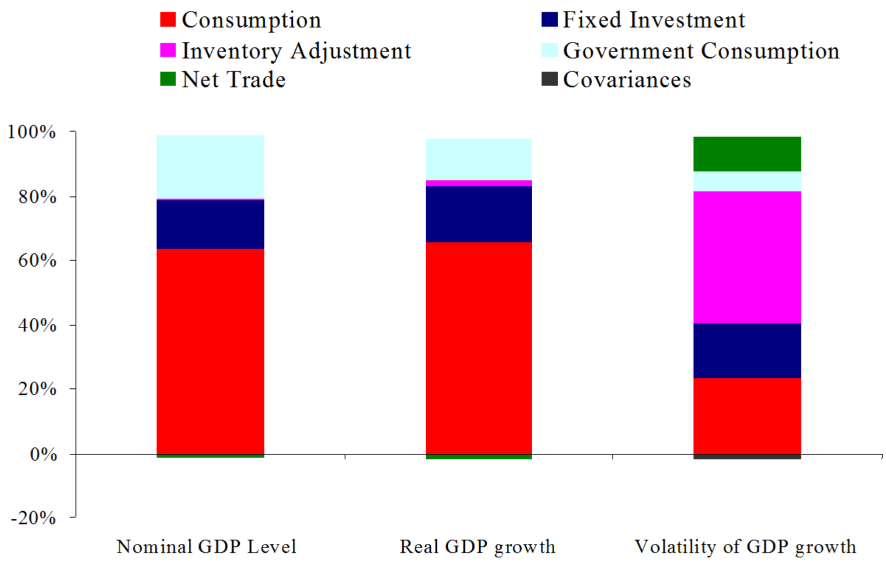

Given this importance, macroeconomists have for a long time been interested, both empirically and theoretically, in the behaviour of inventories at the aggregate level; Metzler (1941) and Abramavitz (1950) are key early references in this regard.[7] Numerous empirical studies (see, for example, (Blinder 1986) have documented a number of stylised facts about inventory behaviour in addition to their key role in the fluctuations in GDP growth. Two early established key facts are that production is more volatile than sales and production and inventory investment are positively correlated. Khan and Thomas (2007b) and (Bils and Kahn 2000) also emphasise that the inventory-sales ratio is counter-cyclical. Wen (2005) shows that inventory investment is only procyclical at business cycle frequency while it is counter-cyclical at higher frequencies. There is little support for a negative relationship between interest rates and inventory investment (Maccini, Moore, and Schaller 2004). These facts poses a first challenge for any model of inventory behaviour over the business cycle.

In fact, it was the fact that production is more volatile than sales that undermined the production-smoothing model which was the main focus of macroeconomic researchers following Holt, Modigliani, Muth, and Simon (1960). Blanchard (1983), looking within the automobile sector, concluded that inventories are, in fact, a desta-

[6]The range for the share of inventory investment in GDP over the sample is [-2%,+4%].

[7]Although Christiano and Fitzgerald (1989) argue that successfully modelling business cycle movements could (and should) proceed without incorporating any speculative inventory holding.

bilising force on output. Miron and Zeldes (1988) and Ramey (1991) are attempts to refine this framework to fit the facts but nonetheless this approach has received less interest in recent years.

There remains little consensus on a canonical model for aggregate inventory behaviour. Many inventory models were restricted to partial equilibrium analyses because of the difficulties in solving the general equilibrium dimension.[8] Early general equilibrium models necessitated modelling short-cuts to admit an internal solution of the model such as including inventories as a factor of production (Kydland and Prescott 1982, Ramey 1989), as part of a sales function ((Bils and Kahn 2000)), or as part of household utility (Kahn, McConnell, and Perez-Quiros 2002).[9]

The two main approaches used in recent analyses are those which assume a stock-out avoidance motive and the (s, S) model of fixed costs of ordering goods. The former, pioneered by (Kahn 1987), assumes that firms hold inventories to avoid a (costly) stockout in which demand exceeds products available for sale and showed in a partial equilibrium framework that serially correlated demand shocks could explain the fact that production was more volatile than sales. The latter, is used in many applications including inventory analysis (for example, Blinder (1981) uses the approach to model retail inventories) and is based on the idea that with a fixed cost of ordering goods, the optimal behaviour for firms is to bunch orders and follow a rule whereby they only reorder once inventories fall to an optimally-determined lower bound (s) and then replenish the inventories to an upper bound level (S).

Embedding these motives in general equilibrium frameworks require complex numerical methods although both have proven reasonably successful approaches. Shibayama (2008) embeds this motive in a general equilibrium business cycle model and finds that his model can help to match the behaviour of inventories at the aggregate level. Khan and Thomas (2007a) compare general equilibrium approaches to the stockout avoidance motive and their approach to the (s, S) model.[10] They find that, under reasonable assumptions about shocks hitting the economy, the (s, S) model performs better.

My emphasis in this paper is on inventories-in-motion and their active management within a general equilibrium business cycle framework. This concept, whilst of great

[8]After the second World War, the study of optimal inventory policy at the level of the firm became an active area of research in management science; early papers include Arrow, Harris, and Marschak (1951), Bellman (1956), and Mills (1957).

[9]Iacoviello, Schiantarelli, and Schuh (2009) use both of these simplifications to facilitate the Bayesian estimation of their model.

[10]Khan and Thomas (2007b) get around the difficulties posed by the fact that aggregation (across time and/or goods) of a large number of firms behaving according to an (s, S) inventory policy would lead to aggregate (macro) behaviour that mimics the (s, S) and is inconsistent with observed data by assuming firm-level idiosyncratic shocks each period to the cost of reordering which help to generate more reasonable aggregate effects.

interest in fields such as Management Science and Operations Research has received little attention in macroeconomics. This idea is also closely linked to the types of improvements that we have seen in the distribution and logistics industry and are related to the technologies that it is suggested have improved inventory management to such a degree that it has affected macroeconomic volatility. It is therefore surprising that macroeconomists have largely ignored the idea.

The key idea is that inventories arise naturally as a result of gaps between production of goods and their consumption as goods are distributed.[11] But such inventories are not passively ignored but rather are actively managed and adjusted to meet consumption and investment needs in the economy. This is very intuitive and differs from earlier approaches which assume inventories would not arise unless there is a decision for them to arise; I assume that inventories will arise and the action is how actively to manage them.

Their are a few papers in international trade that also focus on distribution and delivery lags such as Ravn and Mazzenga (2004) and Burstein, Neves, and Rebelo (2003). The closet paper to mine is Alessandria, Kaboski, and Midrigan (2008) which emphasises the use of inventories as a result of delivery lags and economies of scale. They focus on international trade, rather than the business cycle, and argue that time lags between the order and delivery of goods mean that firms engaged in international trade have a greater inventory management problem which means that looking only at transport costs understates international trade costs substatntially. While this paper also employs a one period delivery lag, the focus (on international trade) differs from my business cycle focus and the analysis is partial equilibrium.

The use of early delivery offers a margin similar to the use of labour in Burnside, Eichenbaum, and Rebelo's (1993) labour hoarding model. In their model, firms make employment decisions in advance of realisations of shocks, but once the shock is realised, they adjust through costly variations in the labour effort workers are asked to supply. In my model, labour is not predetermined but output takes a period to be delivered unless the margin offered by early delivery is employed.

Let me preemptively defend a number of the modelling choices that I make. Firstly, I focus on inventories-in-motion and in the model these arise as finished good inventories but, although the existing literature has argued over which form of inventory holding is the most important (and thus the correct one to model), this is not as restrictive as it may seem at first. For simplicity, I model a single good and a single sector

[11] There is no reason to think that this simply refers to transportation delays but also items awaiting processing, delays in order completion, and also finished goods on shelves (or pallets) awaiting consumers.

and do not explicitly model each level in the distribution chain. Therefore, inventories that arise as part of the distribution chain, consist of all retail and wholesale inventories, finished goods in the manufacturing sector; it is not economically important whether the final goods are stored in a room at the manufacturing plant, on pallets in a wholesalers, or on a retailer's shelves (Summers 1981).

Khan and Thomas (2007b), however, argue that the focus should be on manufacturing inventories rather than retail or wholesale inventories. My model can, at least, capture finished good manufacturing inventories. Once again, the distinction between different types of manufacturing inventories is not always economically relevant. For example, if a steel manufacturer sells steel bars to another firm who holds them as inventory to use later in production, the bars are counted as 'Materials & Supplies'; on the other hand, if the steel bar producer had held the metal bars on their premises, they would be listed as 'Finished Goods'. Nonetheless, the analysis of inventories-in-motion management of input inventories and intermediate goods sectors is an obvious future extension along the lines of Maccini and Pagan (2006).

One may be tempted to draw comparisons eith the seminal 'time-to-build' approach of Kydland and Prescott (1982). In their work inventories enter as a factor of production. My model shares some similarities in that the amount of goods that are available for consumption today, depends on the inventories brought into the period. Though similar to putting inventories in the production function, there are some important differences.For example, in my model, it is the marginal product of capital (and its expected future path) that will determine the behaviour of inventories, while the causality runs the other way with the inventories in the production function models; the positive correlation between inventories and output is driven by inventories in their model, while output will cause inventories directly in my model.

I shall proceed by assuming that productivity shock are the only exogenous source of business cycle variation, and prices are flexible. Given Gali's (1999) arguments that demand shocks are a more likely source of business cycle variation, I have explored, and it is the subject of a future extention, a model along these lines. One difficulty is numerical in that an additional shock adds another state variable which increases the computational burden of the model considerably. Moreover, without sticky prices, demand shocks cannot generate volatility that is of similar orders of magnitude to my supply models presented here. An increasing importance of demand shocks may be key to match the observed falling covariance during the Great Moderation as suggested by Barnichon (2007).

Overall, I believe that inventories appear to be playing a major role in the movement of macroeconomic measures such as GDP and therefore warrant attention from macroeconomists in order to more fully understand their behaviour. I also believe that the main motive in this paper captures an important channel that gives rise to inventories and that I will show matches a number of the key facts about inventory behaviour:

Fact 1 *Inventory adjustment is a small component of GDP growth on average, but it contributes a great deal to its volatility;*

Fact 2 *Sales are less volatile than production;*

Fact 3 *Production and inventory investment are procyclical;*

Fact 4 *The inventory-sales ratio is counter-cyclical.*

3 A Simple 3-period Economy With Inventories As Freight

To illustrate my inventory-in-motion, or freight inventories, mechanism I first explore the idea within a 3-period, stochastic economy ignoring, for the sake of simplicity, the role of capital and labour supply decisions. The simple economy is characterised by a single consumption good from which utility is derived; the consumption choice is given by c_t and consumer preferences are given by:

$$\begin{aligned} \mathbb{U} &= \mathbb{E}\left[U(c_1) + \beta U(c_2) + \beta^2 U(c_3) \right] \\ \text{where } U(c_\tau) &= \frac{c_\tau^{1-\gamma}}{1-\gamma} \end{aligned}$$

In each period, output of consumption goods is given by:

$$y_t = a_t \qquad t = 1, 2, 3$$

where a_t is productivity and labour is supplied inelastically and normalised to 1.

Productivity can take three values - high (a_H), medium (a_M) and low (a_L) with $a_H \geq a_M \geq a_L$. In period 1 and 3, I assume that productivity takes the medium value a_M with certainty. The only uncertainty is regarding the realisation of the endowment in the second period which is given by by the following probability distribution function:

$$a_2 = \begin{cases} a_H & \text{w.p. } p_H \\ a_M & \text{w.p. } 1 - p_L - p_H \\ a_L & \text{w.p. } p_L \end{cases} \tag{1}$$

I consider two types of shocks:

1. *Productivity shocks* - shocks in which actual productivity differs from expected productivity;

2. *Expectations/news shocks* - shocks which change the agent's view about the likely realisation (p_L or p_H or both). Of course, the full effects of these shocks depends on whether the realisation is (close to) what was expected.

The concept of freight inventories is as follows: goods, once produced, take one period to be distributed free of charge and, for simplicity, there is no cost of holding inventories in this form. The amount of goods in the pipeline between period t and $t+1$, or carried into period $t+1$ is given by f_{t+1}.[12] The agent has the option to distribute goods more quickly at a cost; the agent can choose a fraction ι_t ($0 \leq \iota_t \leq 1$) of the output in period t to deliver immediately at a cost which is paid in period t.[13] Agents enter period 1 with a given freight inventory (f_1) and then inventories in freight is given by:

$$f_{t+1} = (1 - \iota_t).y_t$$

Consumption in each period is given by those freight carried from the last period, the immediately delivery output of the current period but less the cost of early delivery:

$$c_t = \iota_t.y_t + f_t - J(\iota_t)\, y_t$$

where $J(\iota_t)$ is the per unit cost of immediate delivery. I assume these costs are convex and given by $J(\iota_t) = w.(\iota_t)^2$, where w is the cost function parameter such that higher w will increase the cost of immediate delivery. I calibrate the w parameter to rule out $\iota_t > 1$.

[12] Stocks are measured as beginning of period.

[13] An alternative approach which leads to an equivalent set up assumes that firms choose an amount of the good to deliver early ($\zeta_t = \iota_t y_t$).

3.1 Model Solution

The optimisation is:

$$\begin{aligned}
\max_{\{c_1,c_2,c_3;\iota_1,\iota_2,\iota_3\}} \mathbb{U} &= \mathbb{E}\big[\, U(c_1) + \beta U(c_2) + \beta^2 U(c_3) \,\big] \\
\text{s.t. } c_1 &= [\iota_1 - J(\iota_1)] \,.a_M + f_1 \\
c_2 &= [\iota_2 - J(\iota_2)] \,.a_2 + (1-\iota_1).a_M \\
c_3 &= [\iota_3 - J(\iota_3)] \,.a_M + (1-\iota_2).a_2 \\
\iota_1 \geq 0 \quad \iota_2 \geq 0 \quad \iota_3 \geq 0
\end{aligned}$$

As $\iota_t \geq 0$, I solve the model using Kuhn-Tucker optimisation given this potentially binding constraint. The full Kuhn-Tucker optimisation is presented in Appendix A while here I outline the recursive approach to the solution.

In the final period, there is no uncertainty and, as there are no future periods, the agent wishes to bring forward as much of y_{t+3} as possible (otherwise it is wasted). To this end, the agent uses early delivery technology until the marginal cost of delivering a unit early is 1 (the benefit of delivering early). Given $J(\iota_t) = w.(\iota_t)^2$, $\iota_3 = \frac{1}{2w}$ and they receive $[\iota_3 - J(\iota_3)] = \frac{1}{4w}$ of their output. I constrain $w \geq \frac{1}{2}$ in order to ensure that $\iota_3 \leq 1$.

This means that once uncertainty is resolved at the start of period 2, the agent can make full plans for period 2 and 3. The model, therefore, is solved in following steps:

1. The agent enters with f_1 and a set of beliefs about the possible outcomes for productivity in period 2:

$$a_2 = \begin{cases} a_H & \text{w.p. } p_H \\ a_M & \text{w.p. } 1 - p_L - p_H \\ a_L & \text{w.p. } p_L \end{cases} \tag{2}$$

2. Based on these, the agent selects c_1 and ι_1 taking into account how the ι_1 decision (see below and the Appendix A for more detail on this), which determines the amount of freight carried into period 2, affects the marginal utility of consumption in period 2.

3. Knowing inventories brought to period 2, and the realisation of the productivity shock in period 2, the agent can then optimay choose c_2, ι_2, c_3 and ι_3.

The only difficulty that arises is that if $\iota_1 \geq 0$ ($\kappa_1 = 0$), the consumers behaviour is given by a two-equation system that includes an intertemporal Euler equation:

$$\begin{aligned} U_c'(c_1)(1 - J'(\iota_1)).a_M &= \mathbb{E}[\beta U_c'(c_2)].a_M \\ c_1 &= [\iota_1 - J(\iota_1)].a_M + f_1 \end{aligned}$$

The difficulty here is that $\Omega \equiv \mathbb{E}[\beta U_c'(c_2)]$ depends on the possible values of the productivity shocks, their associated probabilities and the chosen values for ι_1 (which affects the amount of resources brought into period 2). This does not have an easy closed form solution making the system in period 1 difficult to solve. However, given I can solve explicitly for the optimal outcome in periods 2 and 3 conditional on a given ι_1, I actually calculate explicitly the relationship between Ω (expected period 2 marginal utility) and ι_1.[14] I then approximate this relationship using a 5th-order polynomial in ι_1 and use this polynomial in place of Ω in the system of equations.

3.2 Analysis of Model Results

I now use the model to examine how the economy, and particularly inventories, respond to different productivity and expectations shocks. I also highlight the freight inventories channel by examining behaviour under different parameterisations of the environment. The baseline model parameterisation, along with the alternative parameter values I explore, are outlined in Table 1.

Table 1: Parameterization of 3-Period Model

Description	Parameter	Baseline Value	Alternative Value
Coefficient of risk aversion	γ	2	1
Early Delivery Costs	w	0.5	1
Initial Inventories	f_1	0.5	1
Time Discount Factor	β	0.9	0.5
Low Productivity	a_L	0.5	
Medium Productivity	a_M	1	
High Productivity	a_H	1.5	
Prob(Low Productivity)	p_L	0.1	0.3, 0.8
Prob(Medium Productivity)	$1 - p_{-}p_H$	0.8	0.4, 0.1
Prob(High Productivity)	p_H	0.1	0.3, 0.8

[14]This involves, for each possible realisation of productivity in period 2, calculating the choice of c_2 for a each possible value of ι_1 on a discretized grid between 0 and 1. Using the probabilities attached to each possible outcome, and the marginal utility of c_2, I generate the $\mathbb{E}[U_c'(c_2)]$.

The baseline model assumes that the agent strongly expects period 2 output to be at the Medium level; the probability of a_H and a_L is 0.1 in each case meaning that a_M is expected with probability 0.8. In the baseline case, realised productivity is set to the Medium level and so this 'expect Medium, realise Medium' scenario is listed dubbed 'MM' where the first capital letter indicates the expected level (with probability 0.8) and the second capital letter indicates which realisation occured.[15]

As shown in Figure 2, this is a constant output scenario. The growth in consumption in the final period is to be expected as it is the terminal period and so everything possible will be consumed; this means that inventories are run down and the cost of early delivery increases. There is some growth in consumption between period 1 and 2 and this is driven by risk aversion ($\gamma = 2$) and despite the fact that a_L is unlikely ($p_L = 0.1$), and there is a symmetric small chance of a_H. Nonetheless, in the face of potentially lower future consumption, the risk averse agent builds up inventories as insurance; when the downside does not materialise, the agent consumes more in future periods by running down these inventories.

Figure 2: Baseline Model: Expect a_M ($p_L = p_H = 0.1$), Realise a_M

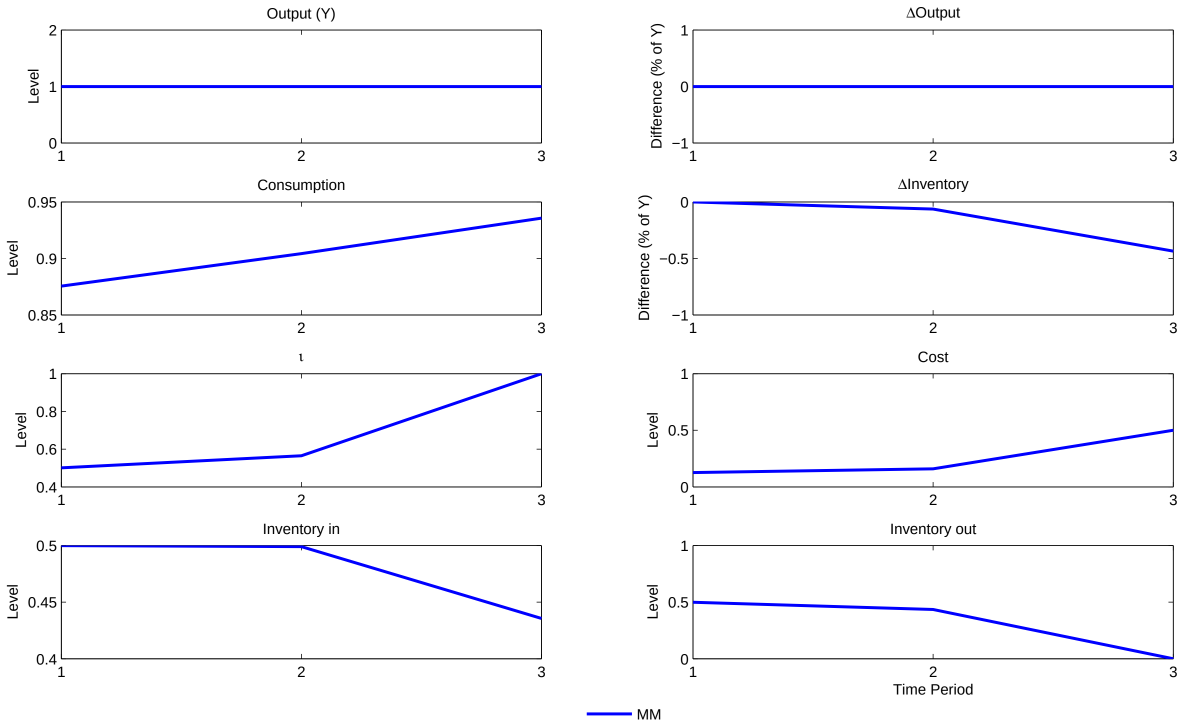

To illustrate this last point, I explore two alternative scenarios called MM_g, in which I reduce the risk aversion parameter to $\gamma = 1$, and MM_h in which I leave the probability of a_M at 0.8 but remove the downside uncertainty ($p_L = 0$ and $p_H = 0.2$).

[15]MM means 'expect Medium, realise Medium', while HL means 'expect High, realise Low'.

I report both these scenarios in Figure 3 and as I report the results as % deviation from the baseline scenario (Δ variables reported as differences), I suffix the model name with an 'r' ($MM_h r$). Reducing γ means that consumption is almost completely flat across period 1 and 2; in fact, the agent is more willing to transfer consumption across periods meaning c_3 is slightly higher (this agent's willingness to smooth across time allows them to reduce the total amount spent on delivery costs). When there is no downside but a 20% chance that a_M is realised (model MM_h), the agent actually consumes more in the first period facilitated through early delivery meaning there are less freight inventories. Once a_M is realised, the agent then consumes less (and rebuilds inventory) in period 2.

Figure 3: Lower Risk Aversion ($MM_g r$) and No Downside Risk ($MM_h r$): Behaviour of Variables Relative to Baseline Model

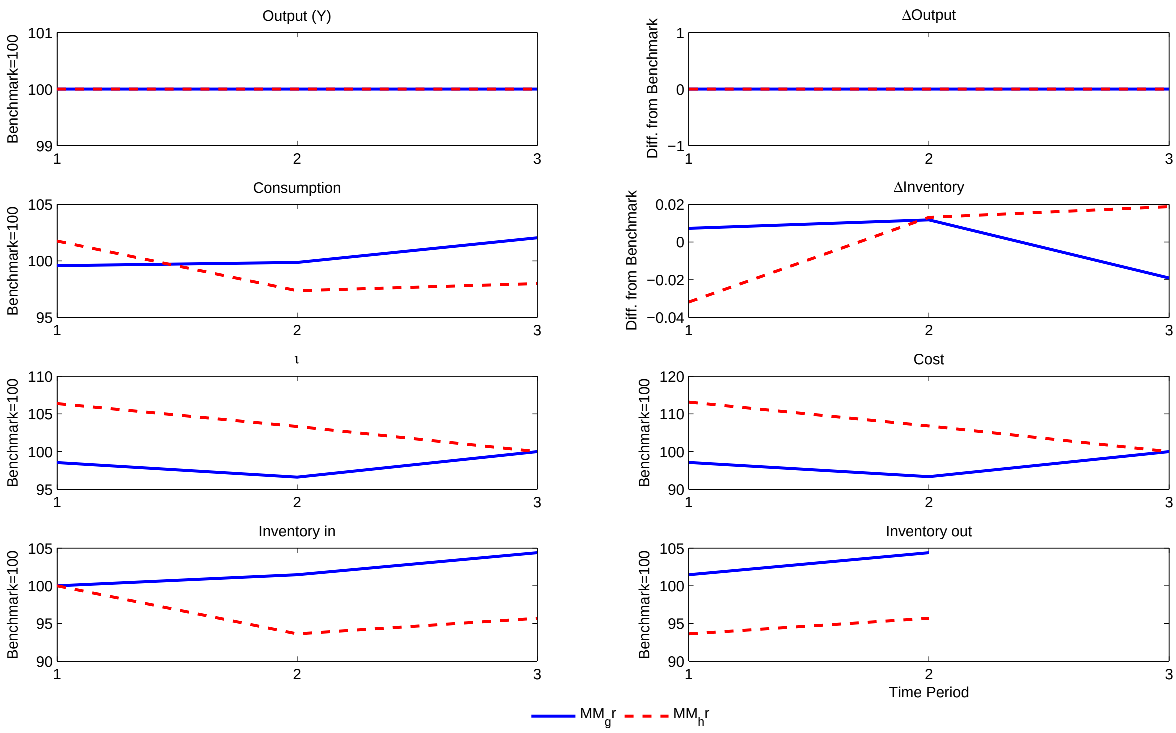

Figure 4 shows the relative behaviour of the model under 3 different parameter specifications; I examine the effect of higher delivery costs ($MM_w r$), higher initial inventories ($MM_f r$) and greater impatience ($MM_b r$). Doubling the delivery costs parameter ($w = 1$ compared with $w = 0.5$) makes early delivery less desirable and so more inventories are carried over (less delivered early). The greater costs also exhaust resources leading to lower consumption in each period. Making the agent more impatient ($\beta = 0.5$ compared with $\beta = 0.9$) has the obvious effect of tilting the consumption profile so that it falls over time.[16] This is facilitated by much greater use of early delivery and the associated de-stocking (despite the associated costs of early delivery). Finally, increasing the inventories brought into period 1 ($f_1 = 1$ compared with $f_1 = 0.5$) acts like a wealth shock in the model; consumption and inventories are both higher across all the periods as the agent uses less early delivery to smooth the higher starting wealth across all periods. (If f_1 were instead low, the agent would choose to bring forward delivery of some goods in order to smooth consumption.)

Figure 4: Higher Inventory Costs ($MM_w r$), Higher Initial Inventories ($MM_f r$) and lower β ($MM_b r$): Behaviour of Variables Relative to Baseline Model

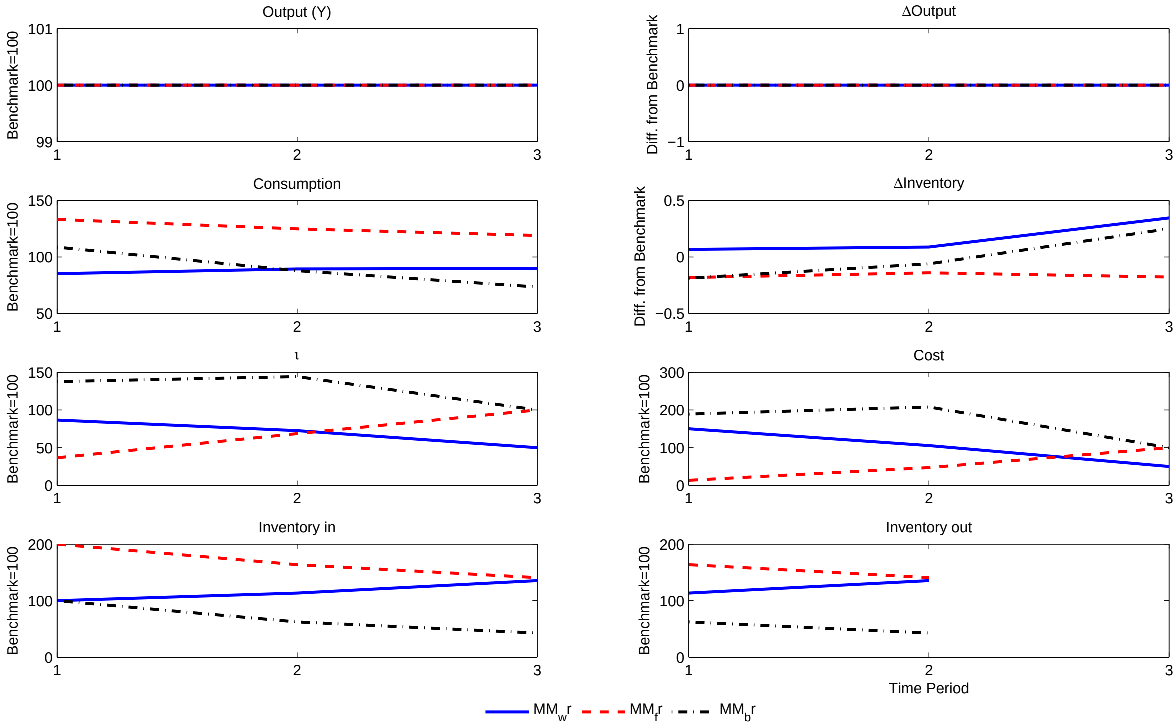

I now consider the effect of different realisations of productivity in the period 2 when everything else is as in the benchmark model (MM) and a_M was the most like outcome. Figure 5 shows that a positive productivity shock boosts output and, *ceteris paribus*, increases the freight inventory generating a positive correlation between output and inventory adjustments. In order to spread the windfall over the last two periods, the agent will bring forward delivery (ι increases). This mechanism directly leads to a positive correlation between output growth and the change in inventories. A negative productivity shock has the mirror image effects.

To illustrate the effects of an expectations shock in the 3-period model, I consider

[16] The less the agent values consumption in the next period, the more likely she is to delivery goods early (and so she is less likely to hold inventories).

Figure 5: Productivity Shock - High (MHr) and Low (MHr) Realisations of a_2: Behaviour of Variables Relative to Baseline Model

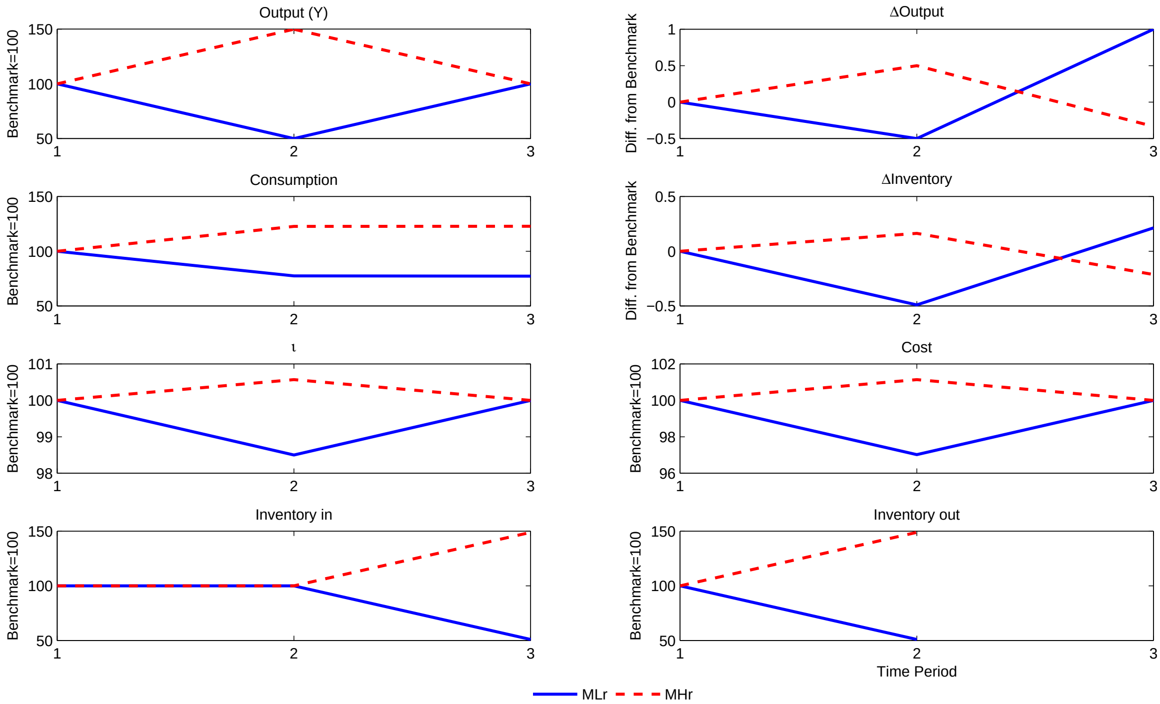

the effects of the agent getting either a positive expectations shock, which entails the agent believing in period 1 that a_H is most likely in period 2 ($p_H = 0.8$, $p_L = 0.1$), or a negative expectations shock (a_L is most likely with mirrored probabilities). While I can use these cases to examine effect of expectations on behaviour in period 1, the dynamic paths of the variables will depend on the actual realisation of productivity. The results presented in Figure 6 assume that, in both cases, the realisation is a_M; this means that the only difference between these and the baseline model is the expectation in period 1.

Expecting a high (low) outturn in period 2, the agent will increase (reduce) consumption in period 1 by increasing (reducing) ι_1. Early delivery is used to smooth consumption even where current production is low relative to expected future productivity. Without it, expecting higher income in the future would not affect consumption today. Of course once productivity is realised as a_M, the high (low) expectation economy behaves as if there were a negative (positive) productivity shock.

3.3 Storage of Goods

All the scenarios so far are such that $\iota > 0$. This outcome is the most likely given that current consumption will always be preferred to future consumption because of

Figure 6: Expectations Shock I - Expect a_H (HMr) or a_L (LMr), Realise a_M: Behaviour of Variables Relative to Baseline Model

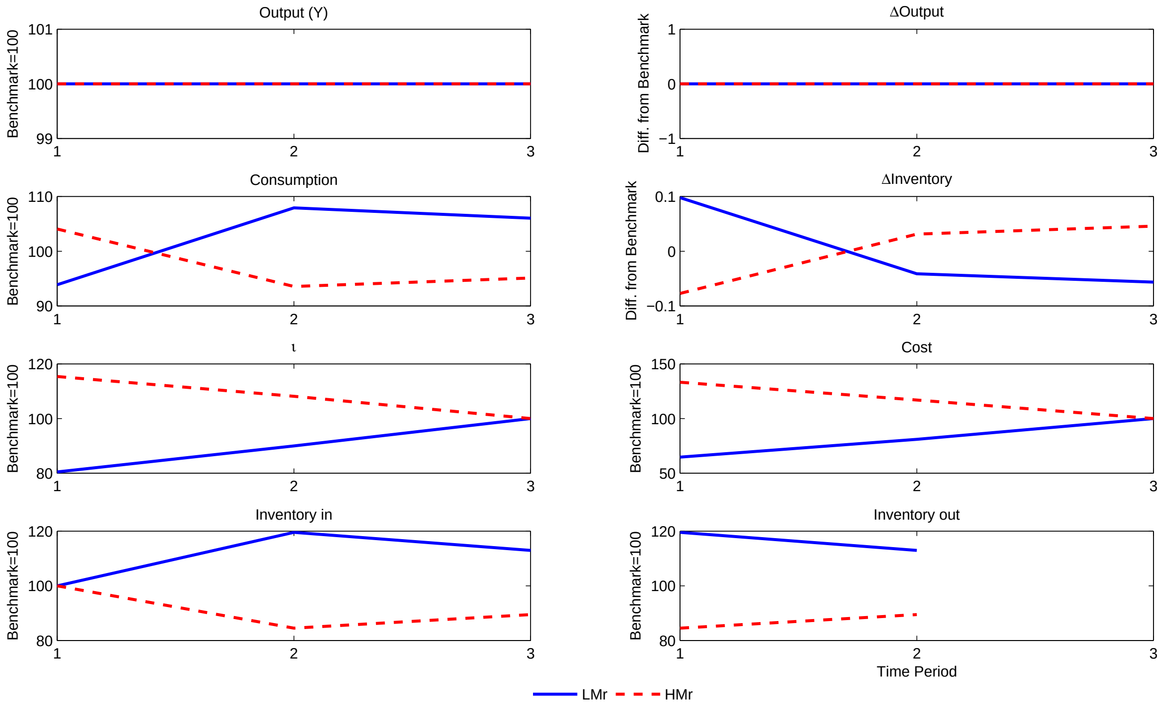

discounting, and the final period acts as a windfall which the agent wishes to bring forward (y_3 is not subject to uncertainty and will be brought forward as much as possible to prevent it being wasted). However, in other circumstances, for example if the agent enters with a very large amount of inventories or the final period has low, or even no, output, the agent may prefer to defer consumption; the non-negativity constraint may bind in early periods ($\iota = 0$). In such cases, the agent might wish to use inventories as a storage device; in times of high productivity, people may choose to store some of the good for consumption at a later time when productivity is expected to be lower.[17]

While I have not considered such an inventory motive in the 3-period model, I do introduce it in the DSGE model when I cannot parameterize the model to ensure $\iota_t > 0$. The amount of goods stored in period t and carried into period $t+1$ will be given by s_{t+1} (for storage). The cost of s_{t+1}, paid in period $t+1$, is given by a simple iceberg cost (υ) which can be thought of as loss, theft or damage resulting from the

[17] A similar story is discussed at the firm level by Blinder and Maccini (1991a). The effects of this type of inventory behaviour may be purely deterministic. For example, consider a manufacturer who knows that demand for their goods will be especially high in January, but also that productivity may be extremely low over Christmas (as many workers are on holidays at times over Christmas and New Years). The producer uses the relatively high productivity of November and the first weeks of December to build up extra inventories in order to meet demand when productivity is low.

stock storage.[18] Crucially for the solution of the model, storage is also required to be non-negative - $s_{t+1} \geq 0$.

4 Dynamic Stochastic General Equilibrium Model

Although the simple model above is useful to introduce the inventory-in-motion concept, it ignores general equilibrium effects and so I now embed freight inventories into an infinite-horizon DSGE model. While I keep the notation equivalent to the 3-period model, there are still many variables, parameters, and functions to keep track of; Table 2 provides a description of the main model variables and parameters.

4.1 Model Setup

The economy is characterised by a single sector which produces storable consumption goods from which utility is derived.[19] This good is produced using capital and labour:

$$y_\tau = a_\tau \left(n_\tau\right)^\alpha k_\tau^{1-\alpha}$$

where a_τ is Total Factor Productivity (TFP), the driving exogenous process in the model, which is given by an AR(1) in logs:

$$\ln a_\tau = \rho \ln a_{\tau-1} + \varepsilon_t$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon)$.

The consumption choice is given by c_τ and consumer preferences are given by:

$$\begin{aligned} \mathbb{U} &= \mathbb{E}_t \left[\sum_{\tau=t}^{\infty} \beta^\tau \left(U\left(c_\tau\right) + bN\left(1-n_\tau\right)\right) \right] \\ \text{where } U\left(c_\tau\right) + bN\left(1-n_\tau\right) &= \frac{c_\tau^{1-\gamma}}{1-\gamma} + b\frac{\left(1-n_\tau\right)^{1-\eta}}{1-\eta} \end{aligned} \tag{3}$$

where b is a parameter which determines the relative weight attached to disutility of labour in the consumer's decision.

[18]The basic analysis is unchanged by the inclusion of a convex cost term $\left(q.(s_{t+1})^2\right)$ to capture warehouse expenses and output losses through the use of labour in the physical handling of the storage.

[19]The distinction between durable and storable is an important one, but not one I address in the model. The goods in my model do not provide a service flow, although I assume that all goods can be stored. In reality, all goods can be stored, although the horizon over which they may be stored differs greatly.

Table 2: Parameters and Variables Used in the Model

Variable	Description
Exogenous Variables And The Underlying Processes	
a_t	TFP variable in period t
ε_t	TFP shock $\varepsilon_t \sim N(0, \sigma_\varepsilon)$
Endogenous Variables and Functions	
y_t	Output in period t
k_t	Capital stock for use in period t
n_t	Labour input in period t
c_t	Consumption in period t
s_{t+1}	Storage carried into period $t+1$
ι_t	Early delivery variable (%) in period t
$J(\iota_t)$	Cost function to bring forward delivery to period t, $= w.\iota^2$
d_{t+1}	Inventories carried into period $t+1$
s_t	Stored inventories from period $t-1$
$f_{\tau+1}$	Freight carried into period $t+1$
Ξ_{t+1}	Adjustment costs of investment - paid in period t
$U(c)$	Utility function in consumption
$N(1-n)$	Utility function in leisure
$Sales_t$	Consumption + investment in period t
$\frac{D_{t+1}}{Sales_t}$	Inventory-sales ratio in period t
Parameters	
α	Labour share of output
$\varkappa$	Parameter of adjustment costs
β	Time discount rate
γ	Parameter on utility from consumption (c)
η	Parameter on utility from leisure $(1-n)$
b	Relative weight of disutility of labour in utility
w	Cost of early delivery parameter
υ	Iceberg cost of storing goods
δ	Depreciation rate
ρ	TFP shock persistence
Numerical Solution	
¥	Update parameter used in numerical solution
$\Omega_n(a_t, k_t, D_t; \omega)$	Expectation approximation function
ω	Expectation approximation coefficients
$\Theta_n(a_t, k_t, D_t; \theta)$	Expectation approximation function
θ	Expectation approximation coefficients
$\Psi_n(a_t, k_t, D_t; \psi)$	Expectation approximation function
ψ	Expectation approximation coefficients

The agent has the option to determine how quickly goods are distributed, as well as an option to store consumption goods. The model is such that each only one of these motives can operate within a period; the agent will never wish to consume more, and thus pay for immediate delivery, and at the same time, wish to defer consumption from today and so pay to store the goods. As before, the amount of goods in the pipeline in period t and carried into period $t+1$ are given by f_{t+1}, the goods stored, at a cost υ, in period t and carried into period $t+1$ are given by s_{t+1}. $J(\iota_t)$ is the per unit cost of immediate delivery.

Capital can, at a cost, be made from the consumption good. The cost of converting the consumption good into a capital good (capital adjustment costs). This is assumption is vital for the inventory motives to operate as well as to capture the positive correlation between consumption and investment in reaction to an expectations shock.[20] I assume that capital adjustment costs are given by:

$$\frac{\varkappa}{2}\left(\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)^2$$

where $k_{t+1}-(1-\delta)k_t$ is net investment and I label the percentage change in the capital stock as $\Xi_{t+1}\left(=\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)$. The assumed form of adjustment costs mean that it is costly to actively adjust (up or down) the level of capital stock. It is not costly, however, beyond the loss of production capability, to allow depreciation to passively reduce the level of capital stock. This means that in steady-state there is a cost to pay and adjustment costs affect the level of consumption in steady state.[21]

Within each period, the goods available are used for consumption, storage or investment and are given by: (i) output produced today but delivered immediately (less costs); (ii) stored goods brought forward from the last period; and (iii) normally delivered goods that have been in freight from the last period but, from which, we subtract the lost output due to adjustment costs. Therefore, the goods flow constraint in each

[20]Without such adjustment costs, all intertemporal consumption smoothing, and particularly storage, can be achieved via capital adjustment - storage of goods as goods (rather than machines) would never be desirable as the return on storage is negative (cost). Similarly, following a positive expectations shock, the agent would normally increase consumption and then, only when the shock is realised, increase investment.

[21]The alternative, such that there is no adjustment costs in steady-state, means that there is an adjustment cost to allowing depreciation to take place. If we consider that the adjustment costs derive from the costs of rearranging or delaying production during new investment installation, then allowing depreciation to take place passively does not seem like it should create costs (beyond loss of production capability).

period is:

$$\begin{aligned} c_\tau + s_{\tau+1} + k_{\tau+1} - (1-\delta)k_\tau &= (\iota_\tau - J(\iota_\tau))y_\tau + (1-\upsilon)s_\tau + f_\tau - \frac{\varkappa}{2}\left(\frac{k_{\tau+1} - (1-\delta)k_\tau}{k_\tau}\right)^2 \\ where\ f_{\tau+1} &= (1-\iota_\tau)y_\tau \\ y_\tau &= a_\tau (n_\tau)^\alpha k_\tau^{1-\alpha} \end{aligned}$$

In addition, the model is subject to the following two occasionally/frequently non-negativity constraints:

$$\begin{aligned} s_{\tau+1} &\geq 0 \\ \iota_\tau &\geq 0 \end{aligned}$$

In order to keep track of total inventories, be they inventories-in-motion or storage inventories, I define the state variable $D_t = (1-\upsilon).s_t + f_t$, 'total beginning of period inventories in period t'. The full optimisation of the model is laid out in Appendix B; the necessary first order conditions for an equilibrium in period t are:

$$U_c'(c_t)\left(1 + \frac{\varkappa}{k_t}\Xi_{t+1}\right) = \beta^2 \mathbb{E}_t\left[U_c'(c_{t+2})(1-\iota_{t+1})MPK_{t+1}\right]$$

$$\ldots + \beta\mathbb{E}_t\left[U_c'(c_{t+1})\left((\iota_{t+1} - J(\iota_{t+1}))MPK_{t+1} + (1-\delta) + \frac{\varkappa}{k_{t+1}}\left(\frac{k_{t+2}}{k_{t+1}}\right)\Xi_{t+1}\right)\right] \quad (4)$$

$$N_n'(1-n_t) = U_c'(c_t)(\iota_t - J(\iota_t))MPL_t + \mathbb{E}_t\left[\beta U_c'(c_{t+1})(1-\iota_t)MPL_t\right] \quad (5)$$

$$\kappa_t = \mathbb{E}_t\left[\beta U_c'(c_{t+1})y_t\right] - U_c'(c_t).(1-J'(\iota_t))y_t \quad (6)$$

$$\mu_t = U_c'(c_t) - \beta(1-\upsilon)\mathbb{E}_t\left[U_c'(c_{t+1})\right] \quad (7)$$

$$c_t + s_{t+1} + k_{t+1} - (1-\delta)k_t = (\iota_t - J(\iota_t))y_t + (1-\upsilon)s_t + f_t - \frac{\varkappa}{2}(\Xi_{t+1})^2 \quad (8)$$

$$\ln a_t = \rho \ln a_{t-1} + \varepsilon_t \quad (9)$$

$$MPL_t = \alpha a_t (n_t)^{\alpha-1} k_t^{1-\alpha} \quad (10)$$

$$MPK_{t+1} = (1-\alpha)a_{t+1}(n_{t+1})^\alpha k_{t+1}^{-\alpha} \quad (11)$$

$$\Xi_{t+1} = \left(\frac{k_{t+1} - (1-\delta)k_t}{k_t}\right) \quad (12)$$

$$\kappa_t \geq 0, \quad \iota_t \geq 0 \quad (13)$$

$$\mu_t \geq 0, \quad s_{t+1} \geq 0 \quad (14)$$

Though they appear more complicated, these are standard first order conditions. Equation (4) is the intertemporal optimality condition for investment in which the cost of investment today (resulting from lower consumption as a result of the investment and the adjustment cost) is set equal to the benefit of that investment (which is the MPK_t and the impact on adjustment costs in the future) which will be spread between period $t+1$ and $t+2$ depending on the choice of how much to deliver next period. Equation (5) is the intratemporal Euler equation which determines the optimal amount of labour taking account of the fact that the benefit to working (the MPL_t) is also split between this period and the next period depending on how much early delivery is used. Equation (8) is the goods constraint, (9) is a standard AR(1) productivity process, while (10), (11) and (12) are definitions of variables to make the conditions more readable.

The key equations in this model are the Kuhn-Tucker multiplier equations associated with the early delivery and storage decisions (equations (6) and (7) respectively), and the relevant non-negativity constraints ((13) and (14)). When it is desirable to deliver goods early ($\iota_t > 0$), $\kappa_t = 0$ and (6) becomes $U_c'(c_t)(1 - J'(\iota_t)) = \mathbb{E}_t[\beta U_c'(c_{t+1})]$ and is the relevant Euler equation for inventories. This equation implies that the expected benefit from bringing forward one more unit of consumption at a cost of $(1 - J'(\iota_t))$, should equal the marginal cost in terms of lower expected utility of consumption in the next period. In this case, $\mu_t > 0$ with no storage being used. On the other hand, where storage is desirable and no goods are delivered early, equation (7) becomes the relevant Euler equation balancing the cost of giving up one unit of consumption today with the benefit of higher consumption in the next period (after taking into account the costs of storage).

GDP Measurement

In this model economy, output can be decomposed into its expenditure components using the accounting identity which is derived from the budget constraint:

$$\begin{aligned} y_t &= c_t + k_{t+1} - (1-\delta)k_t + (D_{t+1} - D_t) + J(\iota_t)\, y_t ... \\ &\quad ... + \frac{\varkappa}{2}(\Xi_{t+1})^2 + \upsilon s_\tau \end{aligned} \tag{15}$$

Output growth is given by the percentage change $(\frac{\Delta y_\tau}{y_{\tau-1}})$.

However, this measurement does not correspond to the measurement of GDP in the economy; business costs (early delivery, storage, and capital adjustment) are intermediate consumption by firms and so need to be subtracted from output. This is true if the costs have to be paid formally (such as business consultancy costs for the installation of new investment) and are hence recorded, but it is also true if the costs were lost output (the use of worker time but without producing the usual physical output).[22] Therefore:

$$\underbrace{y_t - J(\iota_t).y_t - \frac{\varkappa}{2}(\Xi_{t+1})^2 - \upsilon.s_\tau}_{GDP_\tau} = \underbrace{c_t + k_{t+1} - (1-\delta)k_t}_{Sales_\tau} + \underbrace{(D_{t+1} - D_t)}_{\Delta Inventory_\tau} \tag{16}$$

In the analysis presented below, I use this definition of GDP. GDP growth contributions, which correspond to the Bureau of Economic Analysis (BEA) data and allow the adding up of variance of GDP growth, are given by:

$$\frac{\Delta GDP_t}{GDP_{t-1}} = \frac{\Delta c_t}{GDP_{t-1}} + \frac{\Delta(k_{t+1} - (1-\delta)k_t)}{GDP_{t-1}} + \frac{\Delta(D_{t+1} - D_t)}{GDP_{t-1}} \tag{17}$$

Steady-State

The deterministic steady-state of this model, in which productivity and output are constant and there is no uncertainty, is characterised by:[23]

$$\begin{aligned} c_t &= c^* \ \ \forall t \\ a_t &= 1 \ \ \forall t \\ n_t &= n^* \ \ \forall t \\ k_{t+1} &= k^* \ \ \forall t \end{aligned}$$

This steady-state is one in which it is optimal to take early delivery of some goods ($\iota^* > 0$ and $s^* = 0$). The reason for this is that leaving goods in the distribution chain has an implicit cost given by time discounting; if we wait until the next period to receive the goods, the utility that we derive from consuming the goods is lower than if we consume the same goods today. Optimal immediate delivery of goods balances the costs of early delivery with the loss of utility through discounting. Optimality, derived

[22]In the latter case, firm output would actually be measured as GDP and no intermediate consumption would be recorded. Another case concerns where the firm output is measured perfectly but intermediate consumption is mismeasured because the in-house provision of the services is not properly accounted for. In this case, the statistical authorities might attribute all output (y_t) to final value-added but instead mismeasure the expenditure side of the economy. In this case they will need to add a statistical discrepancy equal to the unmeasured parts of spending $\left(J(\iota_t)\,y_t + \frac{\varkappa}{2}(\Xi_{t+1})^2 + \upsilon s_{\tau+1}\right)$.

[23]This deterministic steady state solution is not the same as the zero shock outcome to the model; the difference results from the fact that the expectation function will take account of the probability of different shocks and a sustained period without shocks will not cause them to update their views. In other words, there is no learning in the model.

using $\kappa_t = 0$, and (6), requires is:

$$1 - J'(\iota^*) = \beta \tag{18}$$

The higher the discount rate, the more goods we choose to deliver immediately. Further, the intratemporal Euler equation for labour allocation (equation (5)), the intertemporal Euler equation for investment (equation (4)), and the budget constraint (equation (8)) provide three equations in the three remaining unknown steady-state choice variables (c^*, n^*, k^*). Here, I write these three equations in terms of steady-state variables using the reporting variables of y^*, MPL^*, MPK^*, and D^*:

$$\begin{aligned}
N'_n(1-n^*) &= U'_c(c^*)\,MPL^*(\iota^*(1-\beta) - J(\iota^*) + \beta) && (19)\\
1 + \frac{\varkappa\delta}{k^*}(1-\beta) &= \beta^2(1-\iota^*).MPK^* + \beta(\iota^* - J(\iota^*))MPK^* + \beta(1-\delta) && (20)\\
c^* + \delta k^* + \frac{\varkappa}{2}\delta^2 &= (1 - J(\iota^*))y^* && (21)\\
y^* &= (n^*)^{\alpha}(k^*)^{1-\alpha} && (22)\\
MPL^* &= \alpha(n^*)^{\alpha-1}.(k^*)^{1-\alpha} && (23)\\
MPK^* &= (1-\alpha)(n^*)^{\alpha}(k^*)^{-\alpha} && (24)\\
D^* &= (1-\iota^*)(n^*)^{\alpha}(k^*)^{1-\alpha} && (25)
\end{aligned}$$

Numerical Solution

The difficulty in solving this system is that the presence of occasionally binding constraints means that the method of log-linearisation around the steady-state is not possible. That solution method obtains a (local) approximate solution but it will not capture the effect of kinks in the policy function where constraints occasionally bind. I therefore apply the approach of the parameterized expectations algorithm (PEA).

PEA, which was first used by Wright and Williams (1982, 1984) and popularised by Den Haan and Marcet (1990, 1994), replaces the conditional expectation in the Euler equation with an approximation. The approximation is fine-tuned iteratively until the approximation leads to behaviour that is consistent with the expectations. It is particularly well-suited for the solution of models with occasionally binding constraints, as the parameterized expectation means we do not have to solve separately for the policy and multiplier functions (Christiano and Fisher 2000). I use the non-stochastic variant of PEA which involves explicitly calculating, at pre-selected grid points for possible values of the state variables, the conditional expectation using Gaussian quadrature methods.

I need to approximate the three expectations that appear in the necessary first order conditions of the model using Ω, Θ, and Ψ as the approximation functions as follows:

$$\mathbb{E}_t\left[U_c'\left(c_{t+1}\right)\left(\left(\iota_{t+1}-J\left(\iota_{t+1}\right)\right)MPK_{t+1}+(1-\delta)+\frac{\varkappa}{k_{t+1}}\left(\frac{k_{t+2}}{k_{t+1}}\right)\Xi_{t+1}\right)\right]\approx \Omega_n\left(a_t,k_t,D_t;\omega\right) \qquad (26)$$

$$\mathbb{E}_t\left[U_c'\left(c_{t+2}\right)\left(1-\iota_{t+1}\right)MPK_{t+1}\right]\approx \Theta_n\left(a_t,k_t,D_t;\theta\right) \qquad (27)$$

$$\mathbb{E}_t\left[U_c'\left(c_{t+1}\right)\right]\approx \Psi_n\left(a_t,k_t,D_t;\psi\right) \qquad (28)$$

I use an exponentiated polynomial in the three state variables - TFP (a_t), capital (k_t), and total inventories (D_t) - to approximate the conditional expectation. Once I have an estimate of these expectations, the model can easily be solved for any set of state variables. My expectation functions add an extra level of difficulty in that the $\Theta\left(a_t,k_t,D_t;\theta\right)$ approximation contains the expectation of a variable from two periods ahead (c_{t+2}). To avoid the extra computational burden that fully solving for this expectation would bring, I use the approximation of marginal utility given by (28). Full details of the numerical solution method is provided in Appendix C.

5 Time Aggregation and Calibration

Time Aggregation

Goods in the model take a period to be produced and, in the absence of early delivery, a further period to be distributed. Of course, anecdotal evidence suggests that while this process takes months or even years in some industries, there are other industries in which the distribution cycle should be measured in hours or days. Moreover, not all the delays from order to goods being on shelves in industries such as apparel can be attributed to my inventory-in-motion concept. Unfortunately I am not aware of macroeconomic data on the extent of inventory-in-motion, but it seems unlikely that the average distribution cycle is one quarter.

I, therefore, calibrate a monthly model and then aggregate the data according to BEA standards to match the standard quarterly frequency of business cycle analysis. Of course, without hard data, the choice of one month, rather than one week, or even one day, may seem arbitrary. One month is chosen such that the inventory-sales ratio will be approximately be 1 at a monthly horizon which is consistent with Maccini, Moore, and Schaller's (2004) estimate that the average inventory level corresponds to four weeks of sales.[24] One month also, usefully, makes the aggregation to quarterly frequency easier and reduces the computational time required to analyse my model.

In fact, there is a literature which endorses modelling decisions at a finer frequency given that temporal aggregation of data generated at higher frequencies can have important effects on the behavior of economic time series and evaluation of economic models. Marcellino's (1999) key message is that greater care should be taken with specifying a temporal frequency for theoretical models in order to then compare the predictions of their model to statistical properties of the data. Christiano, Eichenbaum, and Marshall (1991) emphasise temporal aggregation as a cause of Type I errors in empirical tests of models of permanent income hypothesis.[25] I follow Heaton (1993), who focuses on time non-separabilities, and Aadland (2001), looking a labour market behaviour, in calibrating a higher frequency business cycle model and then comparing the time aggregated model series with the data.

One important implication of modelling the business cycle at a monthly, rather quarterly, horizon is that in the baseline calibration (see below), I make use of microeconomic evidence that labour supply is more elastic at higher frequencies than at lower frequencies. A greater willingness to work longer means that the amplification of TFP shocks with inventory control is greater. It also makes it easier to match the annual flow costs arising from inventories without making the use of inventories prohibitively expensive. This is because stock measures are invariant to the horizon of the period under consideration; inventories at the end of 2010 (annual frequency) is the same as inventories at the end of December 2010 (monthly frequency), and the same as inventories at the end of December 31st 2010 (daily frequency). So using a monthly frequency to match an annual flow cost of iventories which equals 12% of the annual value of inventories, a monthly charge of only 1% of the inventory holdings is required (as the total costs for the year will be the sum over the 12 months). By lowering the associated cost at a monthly frequency, it encourages the holding of inventories.

Taking care of this distinction between flow and stock variables is important. In the paper I follow, as closely as possible, the BEA procedures for temporally aggregating data; flow variables are measured as the within quarter sum while stock variables are systematically sampled.[26] In the analysis below, GDP, Consumption, hours, costs and investment are measured as the flow variables while fixed capital and inventories are

[24] Without a role for materials and supplies, and work in progress inventories, both of which are included in the NIPA estimates of inventory holdings, it is not possible for me to match the exact ratio of inventories to sales in the NIPA data.

[25] The interested reader should consult Sims (1971), Geweke (1978), Christiano (1985), Marcet (1991), Lippi and Reichlin (1991), Rossana and Seater (1992) and Granger and Siklos (1995) as other papers in this temporal aggregation bias literature.

[26] In reality, the flow variables are the quarterly sum of systematically sampled flow data.

systematically sampled at the beginning of period. Some flows, such as TFP (z), ι and s are expressed as quarterly averages.

Calibration

I follow Aadland and Huang's (2004) method of consistent higher frequency calibration which is designed to ensure that steady-state values of temporally aggregated flows are consistent across high and low frequency calibrations. The parameter values presented in Table (3) are used to solve my Baseline TFP model of inventories.[27]

Table 3: Parameters in Baseline Model

Description	Parameter	Monthly Value
Utility and Production functions		
Relative weight of disutility of labour in utility	b	3
Parameter on utility from consumption (c)	γ	1
Parameter on utility from leisure ($1-n$)	η	0.66
Time Discount Factor	β	0.998
Labour share of output	α	0.68
Cost Parameters		
Iceberg cost of storing goods	υ	0.001875
Depreciation rate	δ	0.008
Early Delivery Costs	w	0.1
Parameter of adjustment costs	$\varkappa$	181
Productivity Process		
TFP shock persistence	ρ_z	0.9
TFP shock standard deviation	σ_z	0.01

The additively separable isoelastic utility function given by equation (3) requires three calibrated parameters: γ defines the CRRA parameter ($\frac{1}{\gamma}$ = intertemporal elasticity of substitution), b determines the relative weight on the marginal utility of leisure and η is related to the Frisch elasticity of labour supply. The first two of these are straightforward to calibrate. The weight parameter b is 3 in order to ensure that steady-state labour supply is 25% of total monthly time. Standard choices for the CRRA parameter are between 1 and 8; following King and Rebelo (1999), I begin by choosing $\gamma = 1$ as this makes the model solution slightly easier and is a typical value used in the business cycle literature.

Typical quarterly values of η range between 1 and 4 (a smaller η means there is a bigger the impact of wage shocks on labour supply). However, following Aadland and

[27] I have also examined the robustness of the results to alternative calibrations.

Huang (2004), I choose the monthly value for labour supply elasticity ($\frac{1}{\eta}$) such that the agent is more willing to substitute labour from month to month than from quarter to quarter consistent with micro-evidence from studies such as Browning, Hansen, and Heckman (1999) and MaCurdy (1983). I set $\eta = 0.66$ which is the midpoint of Macurdy's estimated range of labour supply elasticity at a monthly frequency.

I calibrate the time discount factor using the long-term real interest rate:[28]

$$\beta^{annual} = \frac{1}{1+\bar{r}} \qquad \beta^{quarterly} = \frac{1}{(1+\bar{r})^{\frac{1}{4}}} \qquad \beta^{monthly} = \frac{1}{(1+\bar{r})^{\frac{1}{12}}}$$

where $\bar{r}$ is the average long-term interest rate. Setting $\bar{r} \approx 3\%$, $\beta^{annual} = 0.970$, $\beta^{quarterly} = 0.993$ and $\beta^{monthly} = 0.998$.

Using data on annual nominal capital stock from the Bureau of Economic Analysis (BEA), the average capital to GDP ratio in the US was about 3 between 1960 and 2007. As the numerator in this ratio is a stock variable, it is not affected by the monthly frequency. The denominator, however, is affected by this choice; the ratio of capital to monthly GDP that I try to match in my model is around 33. Together with the choice of β, the depreciation rate is set to match this capital-output ratio. I use a depreciation rate of 10% per annum, as in King and Rebelo (1999), the monthly depreciation rate δ is chosen to be 0.008. Following Chari, Kehoe, and McGrattan (2000), I chose the investment adjustment cost parameter (χ) to match the relative volatility of investment (to output volatility) found in the US data (3.2 between 1960 and 2007); the chosen monthly value of χ is 181. I set the coefficient on labour in the Cobb-Douglas production technology α to be 0.68 in order to match the average labour share in the US between 1960 and 2007; this is unaffected by the frequency choice.

The typical approach in DSGE models to calibrate the parameters of the productivity process (ρ, σ) is to estimate an AR(1) using a quarterly series for TFP in the economy:

$$\ln a^{Q}_{T+1} = \rho^{Q} . \ln a^{Q}_{T} + \varepsilon_{T+1}$$

where ρ^{Q} and σ^{Q}_{ε}.

As Lippi and Reichlin (1991) show, the estimate of persistence of shocks to GDP is changed by temporal aggregation and, importantly, this change is not necessarily systematic - it may go either way depending on the underlying data. Therefore, in order to calibrate the monthly AR(1) process for TFP in my model, I use a quarterly estimate of US TFP and fit an ARMA(1,1) which is the approximate time-series process for

[28] Other authors choose β to match a steady-state capital-output ratio.

quarterly TFP which is measured as the average of the monthly levels. I then use a Monte-Carlo exercise to find values for the monthly AR(1) that yield quarterly estimates consistent with the data when we aggregate to the quarterly frequency.[29] The monthly parameters I use are $\rho^M = 0.9$ and $\sigma_\varepsilon^M = 0.01$. which correspond to quarterly parameters of $\rho^Q = 0.83$ and $\sigma_\varepsilon^Q = 0.012$.

To calibrate the value of υ, I follow Khan and Thomas (2007b) who use the estimates for the carrying costs of inventories provided by Richardson (1995). In order to include only those costs for storage of goods that my model explicitly covers, I concentrate on the costs of deterioration and pilferage; these are estimated to be in the range of $3\% - 6\%$ of inventory value for the year. However, since my model only covers about 50% of total inventories in the economy, I rescale these values and use the range $1.5\% - 3\%$. Using the middle of this range, this corresponds to a monthly cost of 0.1875 as a percentage of inventory value ($\upsilon = 0.001875$). I assume that the cost function for early delivery is quadratic, $J(\iota) = w.\iota^2$ and I choose w such that the average inventory-sales ratio corresponds to approximately 4 weeks as already discussed. This entails using $w = 0.1$.

6 Macroeconomic Behaviour of Inventories-in-Motion

Policy Functions

I first examine the estimated policy functions for the inventory decision variables in the baseline model. Figure 7 plots, in the top and bottom panels respectively, the optimal choice of ι_t and s_{t+1} (the control variables affecting active inventory management in period t). The plots are drawn for combinations of the TFP variable over the interval $[0.975, 1.025]$ (corresponding to 2.5 standard deviations from steady-state) and start of period inventories are allowed to vary on the interval $[0.91, 2.12]$ which represents steady-state stocks $\pm 40\%$; capital is held constant at its steady-state level.

The policy functions display the kink which arises as a result of hitting the non-negativity constraint and justifies the non-linear solution method I use. Whenever TFP is at, or below, its steady-state level of 1, and even when it is slightly above, the agent will wish to bring forward consumption using early delivery as described in section 4.1. It is only when TFP is particularly high that the agent might consider storing the goods. A log-linear approximation of this model would not capture this

[29] The quarterly values are derived from an ARMA(1,1) model estimated using a logged US TFP series for the period 1961 Q1 to 2006 Q4; the series is detrended using an HP-filter with the smoothness parameter set to 1600 as suggested by Ravn and Uhlig (2001)) for quarterly data.

behaviour.

Figure 7: Optimal Choice of Inventory Control Variables For Different Values of Initial State Variables: Initial capital stock is assumed to be at their steady-state level

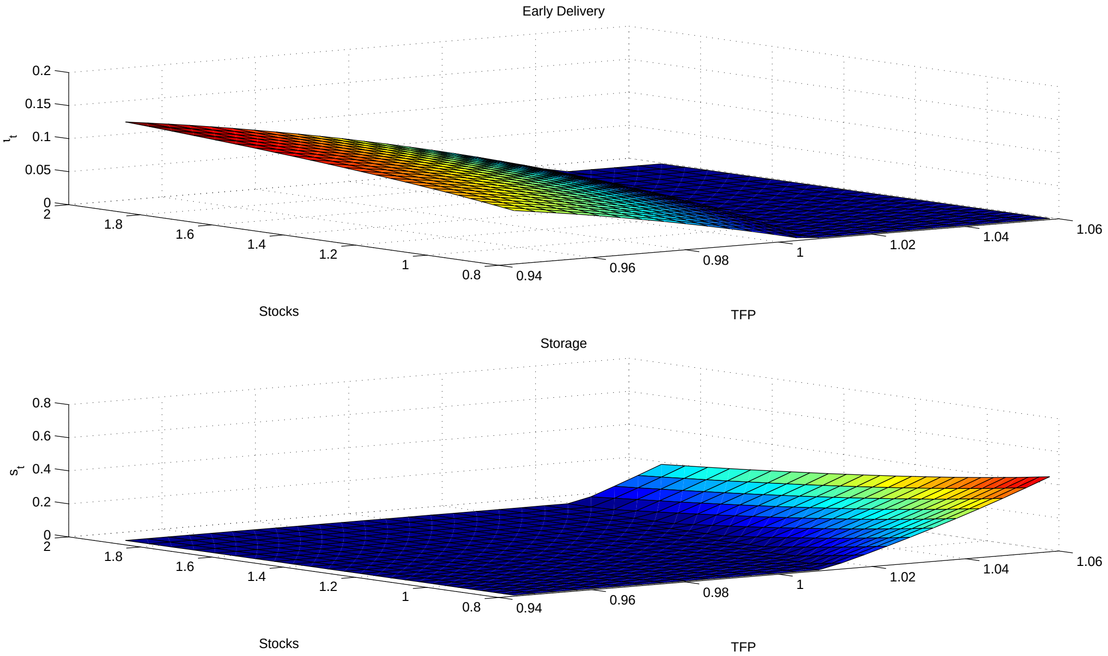

Figure 8 repeats the analysis but in this case capital stock varies over the interval $[41.45, 62.176]$ (which corresponds to $\pm 20\%$ of the steady-state level of capital) and the beginning-of-period stock of inventories is assumed to be at its steady-state level; TFP continues to vary on the interval $[0.975, 1.025]$. In this case, there are two important non-linearities. The first is the kink that is to be expected (and was discussed above) where the capital stock is low, and, at the same time, so is TFP. In such a case, the agent does not wish to bring forward any of the freight and hits the non-negativity constraint. The second is that, for any given capital, the optimal ι_t is an inverted-U shape function of capital stock (though higher capital, *ceteris paribus*, leads to more consumption being brought forward today).

Comparison to two alternative models

In this section I compare the baseline inventory model to two alternative models. The first is a standard RBC model in which there is no lag in availability of goods for consumption, but there are capital adjustment costs and the model is otherwise identical to the baseline inventories model (including in terms of calibration). This model, called 'RBC Model', is solved using PEA that follows closely the solution of the baseline model; it cannot follow it exactly as there are only 2 state variables and

Figure 8: Optimal Choice of Inventory Control Variables For Different Values of Initial State Variables: Initial holdings of inventory are assumed to be at their steady-state level

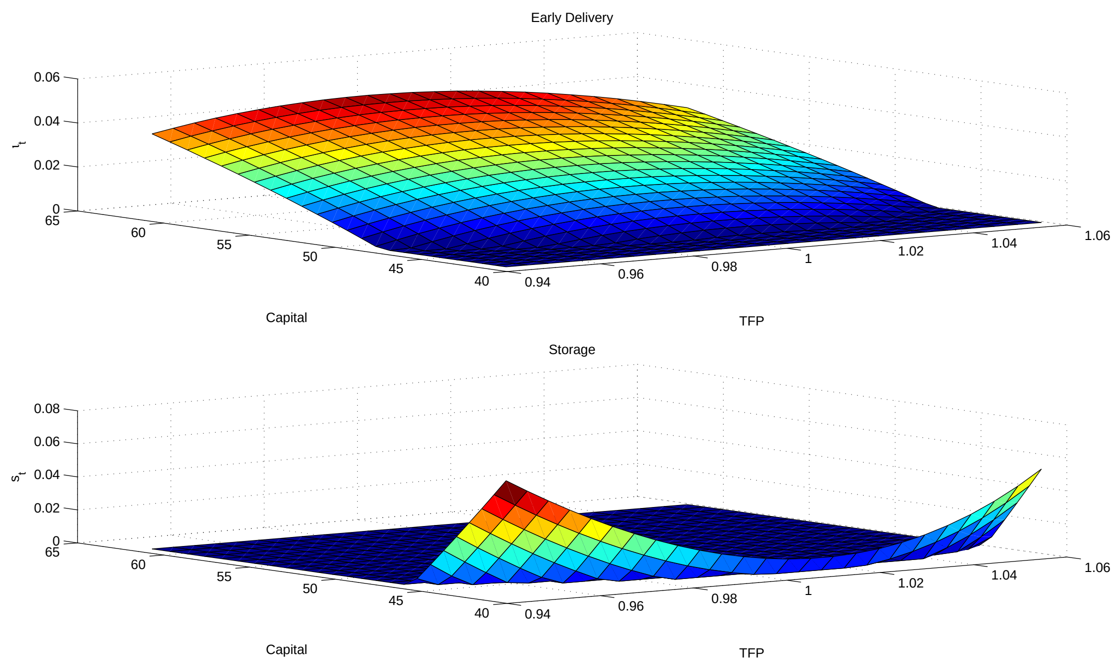

there is only a single expectation to approximate. There are no inventory holdings in this model. Appendix D outlines the necessary conditions for an equilibrium in this model.

Secondly, in order to examine the role played by active control of inventory-in-motion, as distinct from the effect of the lag structure in distribution, I also compare the baseline inventories model to a model in which all goods are subject to the one month delay in delivery but inventories cannot be actively adjusted. In this model, which I will refer to as 'Pipeline Model', inventories are a natural consequence of the pipeline delay and, therefore, this model differs from the baseline model only in that the agent does not exercise any control over inventories.

For the first comparison, I simulate each model using the *same* exogenous shock process for 3000 months; I then aggregate each series to yield a quarterly frequency series of 1000 quarters (all reported results use the quarterly data). Table 4 confirms the achieved calibration in terms of capital and inventory to sales ratios. It also shows that the first characteristic of inventories is matched in my model; inventory adjustment is, on average, a negligible component of GDP (rounding to zero).

Table 5 shows that despite being such a small share on average, inventory adjustment accounts for 25% of the variance of GDP growth in my Baseline Inventories Model. The Pipeline Model actually generates an even larger share of the variance

Table 4: Result of Simulations I: Key ratios and Shares of GDP

	$\frac{K}{GDP}$	$\frac{D}{Sales}$	$\frac{\Delta D}{GDP}$	$\frac{C}{GDP}$	$\frac{I}{GDP}$
Baseline Model	3.0	1.0	-0.0	0.8	0.2
Pipeline Model	3.0	1.0	-0.0	0.8	0.2
RBC model	3.2	0.0	0.0	0.8	0.2

Table 5: Result of Simulations II: Contributions to the Variance of GDP Growth

	GDP		C		I		ΔStocks		Covariance	
	Var	%	Contrib	%	Contrib	%	Contrib	%	Contrib	%
Baseline Model	6.14	100	0.07	1	3.94	64	1.55	25	0.58	9
Pipeline Model	1.36	100	0.01	1	1.22	90	0.52	39	-0.39	-29
RBC model	6.30	1.00	0.07	1	5.12	81	0.00	0	1.11	18

of GDP growth which suggests that some of the large contribution of inventory adjustment to the variance of GDP growth does not, necessarily, represent an active inventory motive and may, in fact, simply be generated by natural delays between production and consumption. However, the Pipeline model generates much smaller variance of GDP growth. All models attribute too much variation in GDP growth to investment and not enough to consumption - this is a typical issue with this type of DSGE model.

Next, I examine the relative volatility of the HP-filtered (log) series from the simulated model; I do this within each model and also across the models using the Baseline Model as the benchmark. The within model analysis is presented in Table 6 and it shows that all models do reasonably well. As mentioned above, I cannot fully match the relative volatility of consumption and all models predict consumption that is too smooth relative to GDP. The main advantage in terms of fit that comes with the active control of inventories in the baseline model is that it increases the volatility of labour hours closer to that in the data. As all models have the same assumed labour supply elasticity ($\eta = 0.66$), active inventory management increases the volatility of hours relative to output by inducing greater labour effort in times of high TFP than is the case where the agent cannot gain from this effort by bringing consumption forward.

Table 6: Volatility of HP-filtered GDP and Components Relative to GDP, by Model

	GDP	C	I	Labour	TFP	$\frac{D}{Sales}$
Baseline Model	1.00	0.18	3.74	0.72	0.54	0.23
Pipeline Model	1.00	0.14	4.17	0.17	1.13	0.31
RBC Model	1.00	0.17	3.73	0.72	0.52	N/A

Table 7: Volatility of HP-filtered GDP and Components Relative to GDP, across Model

	GDP	C	I	Labour	TFP	$\frac{D}{Sales}$
Baseline Model	1.00	1.00	1.00	1.00	1.00	1.00
Pipeline Model	0.47	0.35	0.53	0.11	1.00	0.64
RBC Model	1.02	0.96	1.02	1.03	1.00	N/A

In Table 7 I compare the level of volatility of the HP-filtered time-series across the models using the Baseline Model as the basis for comparison. As expected, the model with inventory management, when subject to precisely the same shocks, has more internal amplification of the shocks relative to the Pipeline Model; the standard deviation of HP-filtered GDP is more than twice as high as the equivalent model in which the agent cannot adjust their inventory behaviour. The main difference comes through the induced labour effort. The consequence of these differences is a more volatile inventory-sales ratio. Table 7 also shows that the RBC model actually generates slightly more volatilty in absolute terms relative to the Baseline Model but it cannot, obviously, match the behaviour of inventories.

In order to better understand the dynamic response of the economy to TFP shocks, in Figure 9 I plot the dynamic response of the main variables in the three models. I assume that, having been at its zero-shock, stochastic steady state, there is a one standard deviation positive technology shock that takes place in period $t = 1$.

The earlier results are clear from these impulse responses. The increase in TFP induces a larger increase in hours in the RBC and Baseline models, but almost nothing in the Pipeline model. The increase in both TFP and hours push the MPK up which amplifies the increase in investment. In both the Baseline and Pipeline models, there is, as a direct result of the increase in output an increase in inventories (and positive inventories investment). However, by being able to actively manage inventories means the agent can reduce the amount brought forward and so smooth some of the higher productivity over future periods; less goods are brought forward for immediate consumption which amplifies the increase in inventories. This generates a more marked increase in the inventory-sales ratio.

The results of a negative shock, not reported here, are similar. What is noticeable is that storage is not used in response to either shock in the baseline model. Nor does it appear to be regularly used in the simulations. There are two reasons for this: (1) storage is less desirable than capital investment for carrying excess goods from one period to the next period (because capital investment provides a positive return in terms of more output the next period), and (2) in steady-state, some goods are

Figure 9: Impulse Response Functions: Effect of a 1 Standard Deviation Positive TFP Shock, Relative to own Steady-State (SS)

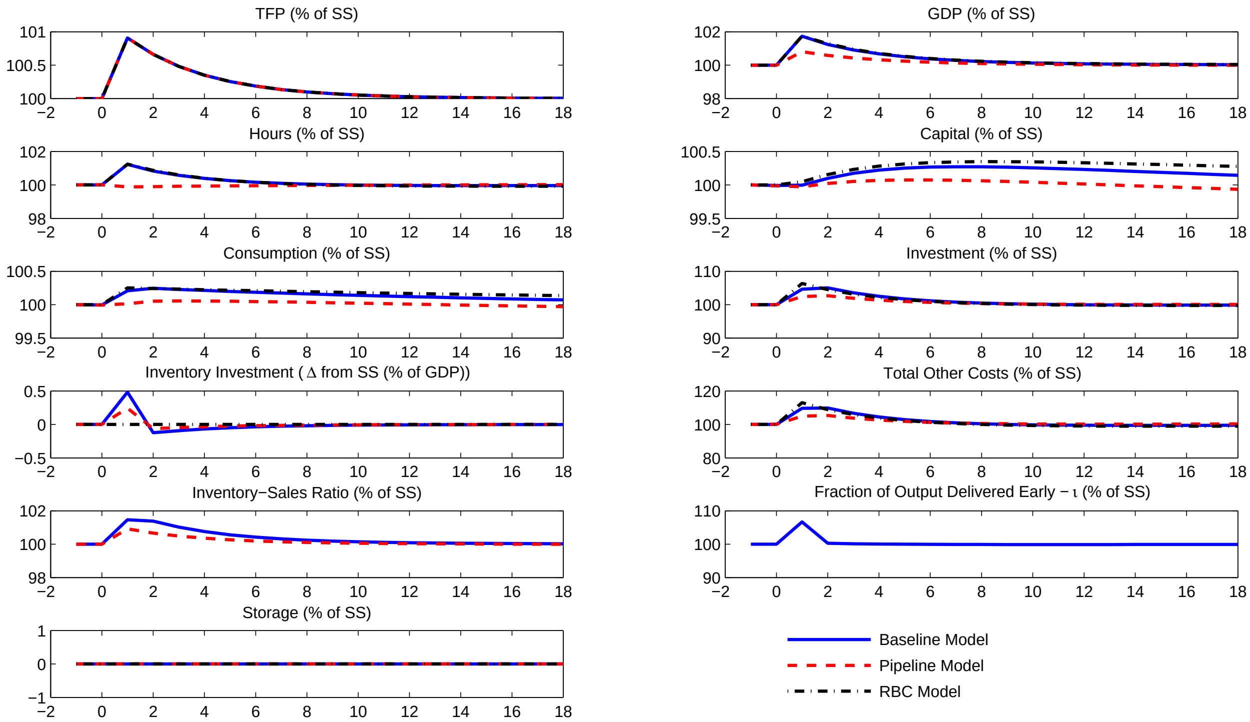

brought forward ($\iota^* > 0$) and, therefore, if the agent wishes to defer consumption using inventories, the first response is to reduce ι and leave more goods in the pipeline.

While I can generate a greater role for the storage motive by increasing the costs of adjusting capital ($\chi^{Alt} = 401$), increasing the cost of early delivery ($w^{Alt} = 5$), and/or reducing the cost of storage ($\upsilon^{Alt} = 0$), storage is still not prevalent in the model responses. Another possibility is to make investment irreversible; if investment cannot be converted back to consumption goods, then it will be less desirable in providing intertemporal storage.[30] However, as I my focus is on the inventories-in-motion concept, I do not try to increase the use of the storage motive in this paper.

Finally, the procyclical response of the inventory-sales ratio in the impulse response functions seems to go against the stylised facts listed above. However, the response of the inventory control variables (and so the inventory-sales ratio) depends, as already discussed, not only on TFP but also on the other state variables; the impulse response functions begin at the steady-state while the response of inventories will differ if the economy starts away from the steady-state. Therefore, I now wish to examine the behaviour of this correlation between GDP and the inventory-sales ratio, as well as

[30]However, in the baseline model most of the capital adjustments that are equivalent to "consuming capital" are in fact achieved by simply investing lower (but still positive) amounts and letting depreciation run capital down.

Table 8: Other Correlations of Interest

	$\rho(\frac{D}{Sales}$, GDP)	$\frac{\sigma_{sales}}{\sigma_{GDP}}$	$\frac{\sigma_{salesHP}}{\sigma_{GDPHP}}$	ρ(GDP,ΔStocks)	ρ(MPK,ΔStocks)
Baseline Model	-0.09	0.97	0.94	0.19	0.27
Pipeline Model	0.17	1.00	1.00	0.12	0.19
RBC Model	NA	1.00	1.00	NA	NA

other stylised facts, by using the whole 1,000 quarter simulation.

In Table 8, I confirm that, overall, the correlation between the inventory-sales ratio ($\frac{D}{Sales}$) and GDP is negative. In fact, if I reduce the cost of early delivery, this correlation falls even more. In the Pipeline Model, this correlation is positive suggesting the control of inventory-in-motion may be an important for matching (at least qualitatively) the data. I also confirm that sales are less volatile than output whether we use the raw series or, even more so, if we HP-filter the individual data first. The Pipeline model also fails to match this fact although both inventory models can generate the positive correlation between inventory investment and GDP.

Finally, I examine the relationship between the real interest rate, represented by the marginal product of capital (MPK) and inventories. Within my framework, inventory investment is positively correlated with the real interest rate. Most models of inventory behaviour predict a negative relationship between the real interest rate and inventory investment, but the empirical evidence fails to support such a relationship (see, for example, Maccini, Moore, and Schaller (2004)). Hence, there may be a theoretical explanation, complementary to the econometric reasoning put forward by Maccini, Moore, and Schaller (2004), for why the negative relationship suggested by typical firm-level analyses is difficult to find in the data.

7 Inventory Management and the Great Moderation

My first application is to investigate the 'good luck' and better inventory management hypotheses as potential explanations for the Great Moderation.[31] The 'good luck' hypothesis argues that the greater macroeconomic stability is due simply to the absence of large shocks (such as the 1970s oil shocks) in the period since 1983; for example, Stock and Watson (2002) tend to favour the 'good luck' hypothesis attributing

[31] Others, not discussed here, include 'good policy', compositional shifts from highly-volatile production to less-volatile service sectors, and improvements in the institutional framework (Acemoglu, Johnson, Robinson, and Thaicharoen 2003)

70%-80% of the volatility decline to good luck. Improved inventory management is the main explanation put forward by McConnell and Perez-Quiros (2000) and Kahn, McConnell, and Perez-Quiros (2002). The idea that improved inventory management has contributed to lower GDP volatility is pertinent amongst policy-makers (see, for example, Bean (2003) and Bernanke (2004)). The arguement is that improvements in inventory management techniques, made possible by advances in information and communications technology, are the source of this lower volatility.

Figure 10 shows the 10-year rolling variance of GDP growth in the US between 1960 and 2007, as well as the contributions to the variance from the goods output (further split into sales and inventories), services output, and structures output.[32] This graph supports the conclusions of McConnell and Perez-Quiros (2000) that the cause of the lower volatility was improved inventory management techniques. They traced the decline in volatility in the early 1980s to a fall in the volatility of goods output which was driven by the reduced use of inventories in that sector.[33] The figure also points to the role of declining covariance in the Great Moderation.

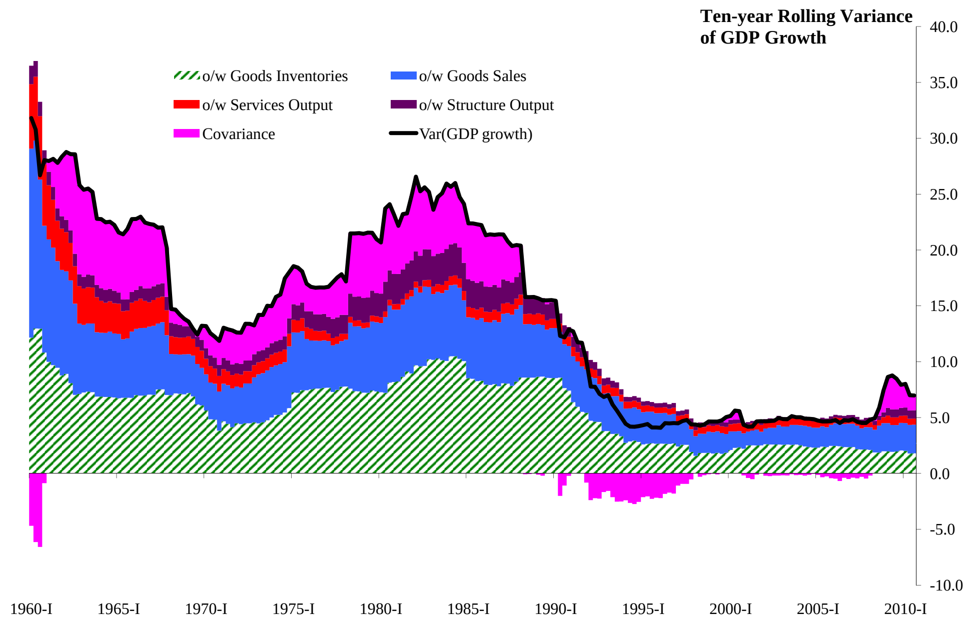

Figure 10: Ten-year rolling-variance of quarterly U.S. GDP growth and the contributions by types of product.

[32]This graph uses the decomposition:

$$GDP_t = sales_t^{gds} + \Delta inventories_t^{gds} + y_t^{services} + y_t^{structures}$$

[33]Their argument that goods sales are unaffected is contested by Ahmed, Levin, and Wilson (2004) who find it was a decline in both sales and production volatility. Moreover, their analysis concludes that the main component of the goods inventories decline is durable goods.

If we split the 48 years between 1960 and 2007 into the two periods identified by McConnell and Perez-Quiros; a more volatile period from 1960-1983 and the Great Moderation era from 1984-2007. The variance of GDP growth during the Great Moderation is about one-third of the variance from the earlier period. Moreover, while there have been large declines in the main components of GDP between the two periods, the declines are not as large as those in GDP volatility and the difference is captured by a key role for a decline in the covariance of the components of GDP; the covariance terms have swung from amplifying fluctuations in the components of GDP, to reducing them.

The Great Moderation Period also affected inventories in other ways. The procyclicality of the inventory-sales ratio increased in the post-1983 period; the correlation between the HP-filtered log I-S ratio and HP-filtered log GDP is -0.46 in the volatile period, but increases to -0.15 in the period of the Great Moderation. There has also been a coincident decline in both GDP volatility and the inventory-sales ratio. Both the real and nominal ratio of non-farm inventories to sales of goods and structures has declined by about 10% from 2.53 in 1983 to 2.27 in 2007. The capital-output ratio has not, on average, changed between the two periods.

The Experiments

In order to explore how well the improved inventory management and 'good luck' hypotheses perform as two possible explanations, I carry out three experiments using the Baseline Inventories Model (unless otherwise stated, all the calibrated parameters are those in Table 3): **Experiment 1 - 'Improved Inventory Management'**
The management of inventory encapsulates not only the monitoring and control of existing inventories but includes controlling their optimal level through the ordering of new stock at optimal times and the analysis of sales data as well. The development of better IT (both hardware and software) have made it easier for firms to improve the management of demand and inventory (McCarthy and Zakrajsek 2007). Such technologies include barcoding and scanners, Radio Frequency Identification Tags (RFID), Electronic Data Interchange (EDI), and Collaborative Planning, Forecasting and Replenishment (CPFR). These developments make inventories easier to monitor and control including while in motion. I model better inventory management techniques as lower costs of distribution, and then examine whether the model generates lower macroeconomic volatility that is comparable to the decline experienced during the 'Great Moderation'. Particularly, I reduce the cost of delivery such that the change in the steady-state inventory to sales ratio matches the actual decline from 1983 to 2007

Table 9: Great Moderation Experiments I: Contributions to the Variance of GDP Growth

	GDP		C		I		ΔStocks		Covariance	
	Var	%	Contrib	%	Contrib	%	Contrib	%	Contrib	%
Baseline Model	6.14	100	0.07	1	3.94	64	1.55	25	0.58	9
Improved Inventory	6.28	100	0.07	1	3.37	54	0.76	12	2.08	33
Good Luck	2.15	100	0.03	1	1.35	63	0.54	25	0.23	10
Both	2.21	100	0.03	1	1.17	53	0.27	12	0.75	34

(10%); the cost of active management of inventories (w) falls from 0.1 in the baseline model, to 0.0019 in Experiment 1.

Experiment 2 - 'Good Luck'

To model 'good luck', I calibrate the change in the parameters of the TFP shocks to match the decline in volatility of GDP. While it is more directly going to generate lower volatility of GDP growth, the interesting test of this experiment is whether the inventory series change in accordance with the observed data. In particular, I reduce the volatility of the TFP shocks such that I match the relative decline in variance of GDP growth between the pre- and post-1984 samples (variance of GDP growth post-1984 is $\frac{1}{3}$ the variance of the pre-1984 period). To do this, I reduce σ_z from 0.01 to 0.006.

Experiment 3 - 'Both together'

Experiment 3 combines the two explanations and explores the outcomes if both the "good luck" and better inventory management techniques play a role.

Results

The results of the experiments are presented in Table 9, 10 and 11. My calibration is successful in ensuring that lower costs of inventory lead to a 10% decline in the inventory-sales ratio, and that the reduced shock volatility leads to volatility of GDP growth that is one third of the pre-Moderation level. The test of these experiments is then whether the lower costs reduces the variance of GDP growth and whether the 'good luck' hypothesis' generates the changes in behaviour of inventories.

The results indicate that it is unlikely that declining costs of active inventory control could have generated the Great Moderation within my model environment. Lower costs of distribution actually increase the volatility of GDP growth (see Table 9) and make the inventory-sales ratio more counter-cyclical (Table 11). As was argued earlier, this is driven by the behaviour of labour hours; labour input is more volatile as the representative agent is more willing to take advantage of high (low) productivity

Table 10: Great Moderation Experiments II: Key ratios and Shares of GDP

	$\frac{K}{GDP}$	$\frac{D}{Sales}$	$\frac{\Delta D}{GDP}$	$\frac{C}{GDP}$	$\frac{I}{GDP}$
Baseline Model	3.0	1.0	-0.0	0.8	0.2
Improved Inventories	3.0	0.9	-0.0	0.8	0.2
Good Luck	3.0	1.0	-0.0	0.8	0.2
Both	3.0	0.9	-0.0	0.8	0.2

Table 11: Great Moderation Experiments III: Other Correlations of Interest

	$\rho(\frac{D}{Sales},$ GDP$)$	$\frac{\sigma_{sales}}{\sigma_{GDP}}$	$\frac{\sigma_{sales^{HP}}}{\sigma_{GDP^{HP}}}$	ρ(GDP,ΔStocks)	ρ(MPK,ΔStocks)
Baseline Model	-0.09	0.97	0.94	0.19	0.27
Improved Inventories	-0.23	0.93	0.89	0.29	0.40
Good Luck	-0.07	0.97	0.94	0.18	0.27
Both	-0.17	0.94	0.89	0.28	0.41

by working more (less) and using inventories to smooth consumption. Moreover, this adjustment comes at the cost of increasing the covariance between the main elements of GDP, rather than decreasing it.

On the other hand, the 'good luck' experiment matches the relative decline in GDP variance. This is unsurprising given that it was calibrated to achieve this result. But it cannot match other changes in the data. The 'good luck' hypothesis successfully reduces the counter-cyclicality of this ratio (slightly) but does not lower the inventory-sales ratio. Covariance is little changed by the reduced volatility of shocks.

Thus, it is not surprising that my final experiment, a combination of the two, is more successful at matching most of the relevant data. I would, therefore, conclude that the both stories are required to explain the recent behaviour of inventories and GDP volatility, although only the 'good luck' hypothesis contributes to the reduced variance of GDP growth meaning that, conditional on my model, inventory management played at best a supporting role in the Great Moderation.

This analysis is similar to that carried out by Iacoviello, Schiantarelli, and Schuh (2009), Khan and Thomas (2007b), and McCarthy and Zakrajsek (2007). Although each paper uses different approaches (for example, calibration versus Bayesian estimation) and different motives for inventories, the conclusions are robust across models; inventory management techniques appear to have played only a minor role in the Great Moderation. However, even allowing for both improved inventory management techniques and 'good luck', we miss a large part of the Great Moderation story. The combined experiment cannot match the decline in the covariance between components of demand, sectors of the economy, or types of product in the economy.

8 Conclusion

In this paper, I focus on the concept of inventories-in-motion - that is, where a firm tries to optimally control inventories that arise naturally between the production and the consumption of the goods. The first contribution is to calibrate and solve a model of distribution inventories at the monthly horizon which, when aggregated to the quarterly frequency, is able to successfully match (at least qualitatively) a number of key facts about the macroeconomic behaviour of inventories:

- Inventory adjustment is a small component of GDP growth, but it contributes a great deal to its volatility;
- Sales are less volatile than production;
- Production and inventory investment are procyclical;
- The inventory-sales ratio is counter-cyclical.

I then use this model to assess the impact on business cycle volatility of changes in the technology used to manage distribution inventories. In particular, I explore whether, according to my model, these improvements in inventory management can explain the decline in macroeconomic volatility in the last 30 years. Mapping the salient features of the improvements in inventory management into the parameters of my model, I find that although the inventory management changes are useful to match aspects of the changes in inventory behaviour over the period, they play no role in the reduction of the variance of GDP growth. In my model, the "good luck" hypothesis is a more likely explanation for the Great Moderation decline in volatility of GDP growth. However, the "good luck" hypothesis alone fails to match other developments in the aggregate data. These other developments are more closely matched by the inventory-management explanation. I therefore conclude that the two explanations have played a role in the behaviour of GDP since the mid-1980s.

The results that I have presented in this paper only allow for productivity shocks as a source of business cycle variation, and prices are flexible. As already discussed, I believe that an extension to an environment with sticky prices and demand shocks is a natural extension of the this work. Moreover, a shift from supply to demand shocks after the end of the 1970s may help to generate the the observed falling covariance during the Great Moderation. I am also extending the model to explicitly consider input inventories and work-in-progress.

Nonetheless, although it is conceptually very simple, I believe the inventories-in-motion concept captures a new and important approach to inventories over the business cycle which has previously been ignored by macroeconomists.

References

AADLAND, D. (2001): "High frequency real business cycles," *Journal of Monetary Economics*, 48(2), 271–292.

AADLAND, D., AND K. X. D. HUANG (2004): "Consistent high-frequency calibration," *Journal of Economic Dynamics and Control*, 28(11), 2277–2295.

ABRAMAVITZ, M. (1950): *Inventories and Business Cycles*. National Bureau of Economic Research, Cambridge, Mass.

ACEMOGLU, D., S. JOHNSON, J. ROBINSON, AND Y. THAICHAROEN (2003): "Institutional Causes, Macroeconomic Symptoms: Volatility, Crises and Growth," *Journal of Monetary Economics*, 50, pp. 49–123.

AHMED, S., A. LEVIN, AND B.-A. WILSON (2004): "Recent U.S. Macro Stability: Good Policies, Good Practices, or Good Luck?," *Review of Economics and Statistics*, 86(3), 824–832.

ALESSANDRIA, G., J. KABOSKI, AND V. MIDRIGAN (2008): "Inventories, Lumpy Trade, and Large Devaluations," NBER Working Papers 13790, National Bureau of Economic Research, Inc.

ARROW, K. J., T. HARRIS, AND J. MARSCHAK (1951): "Optimal Inventory Policy," *Econometrica*, 19, 250–272.

BARNICHON, R. (2007): "Productivity, Aggregate Demand and Unemployment Fluctuations," CEP Discussion Papers dp0819, Centre for Economic Performance, LSE.

BEAN, C. (2003): "Inflation Targeting: The UK Experience," Speech at the Annual Congress of the German Economic Association, University Zurich-Irchel in Switzerland 1 October 2003, Bank of England.

BELLMAN, R. (1956): "On the Theory of Dynamic Programming-A Warehousing Problem," *Management Science*, 2(3), 272–275.

BENATI, L. (2004): "Evolving post-World War II UK economic performance," *Journal of Money, Credit, and Banking*, 36(4), 691–717.

Bernanke, B. S. (2004): "The Great Moderation," Remarks by Governor Ben S. Bernanke at the meetings of the Eastern Economic Association, Washington, DC February 20, 2004, Federal Reserve Board of Governors.

Bils, M., and J. A. Kahn (2000): "What Inventory Behavior Tells Us about Business Cycles," *American Economic Review*, 90(3), 458–481.

Blanchard, O., and J. Simon (2001): "The Long and Large Decline in U.S. Output Volatility.," *Brookings Papers on Economic Activity*, 1, pp. 135–174.

Blanchard, O. J. (1983): "The Production and Inventory Behavior of the American Automobile Industry," *Journal of Political Economy*, 91(3), 365–400.

Blinder, A. S. (1981): "Retail Inventory Behavior and Business Fluctuations," *Brookings Papers on Economic Activity*, 12(1981-2), 443–505.

Blinder, A. S. (1986): "Can the Production Smoothing Model of Inventory Behavior Be Saved?," *The Quarterly Journal of Economics*, 101(3), 431–53.

Blinder, A. S., and L. J. Maccini (1991a): "The Resurgence of Inventory Research: What Have We Learned?," *Journal of Economic Surveys*, 5(4), 291–328.

——— (1991b): "Taking Stock: A Critical Assessment of Recent Research on Inventories," *Journal of Economic Perspectives*, 5(1), 73–96.

Browning, M., L. P. Hansen, and J. J. Heckman (1999): "Micro data and general equilibrium models," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford.

Bureau of Economic Analysis (2008): "Concepts and Methods of the U.S. National Income and Product Accounts," Discussion paper, Bureau of Economic Analysis.

Burnside, C., M. Eichenbaum, and S. Rebelo (1993): "Labor Hoarding and the Business Cycle," *Journal of Political Economy*, 101(2), 245–73.

Burstein, A. T., J. C. Neves, and S. Rebelo (2003): "Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations," *Journal of Monetary Economics*, 50(6), 1189–1214.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2000): "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?," *Econometrica*, 68(5), 1151–1180.

Christiano, L. J. (1985): "A method for estimating the timing interval in a linear econometric model, with an application to Taylor's model of staggered contracts," *Journal of Economic Dynamics and Control*, 9(4), 363–404.

Christiano, L. J., M. Eichenbaum, and D. Marshall (1991): "The Permanent Income Hypothesis Revisited," *Econometrica*, 59(2), 397–423.

Christiano, L. J., and J. D. M. Fisher (2000): "Algorithms for solving dynamic models with occasionally binding constraints," *Journal of Economic Dynamics and Control*, 24(8), 1179–1232.

Christiano, L. J., and T. J. Fitzgerald (1989): "The magnitude of the speculative motive for holding inventories in a real business cycle model," Discussion Paper / Institute for Empirical Macroeconomics 10, Federal Reserve Bank of Minneapolis.

Copacino, W. (1988): "Managing Inventory in Motion," in *Supply Chain Management: The Basics and Beyond*, ed. by W. Copacino. CRC Press, 1997.

Den Haan, W. (2007): "Parameterized Expectations," Lecture notes, University of Amsterdam.

Den Haan, W., and A. Marcet (1990): "Solving the Stochastic Growth Model by Parameterizing Expectations," *Journal of Business and Economic Statistics*, 8, 31–34.

——— (1994): "Accuracy in Simulations," *Review of Economic Studies*, 61, 3–17.

Forrester, J. (1961): *Industrial Dynamics*. MIT Press, Cambridge, MA.

Gali, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review*, 89(1), 249–271.

Geweke, J. F. (1978): "Temporal Aggregation in the Multiple Regression Model," *Econometrica*, 46(3), 643–61.

Gilmore, D. (2008): "Supply Chain News: The 10 Best Quotes from CSCMP 2008," Supply Chain Digest 23 October 2008.

Granger, C. W. J., and P. L. Siklos (1995): "Systematic sampling, temporal aggregation, seasonal adjustment, and cointegration theory and evidence," *Journal of Econometrics*, 66(1-2), 357–369.

Heaton, J. (1993): "The Interaction between Time-Nonseparable Preferences and Time Aggregation," *Econometrica*, 61(2), 353–85.

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

Iacoviello, M., F. Schiantarelli, and S. Schuh (2009): "Input and Output Inventories in General Equilibrium," Boston College Working Papers in Economics 658, Boston College Department of Economics.

Judd, K. (1998): *Numerical methods in economics.* MIT Press, Cambridge, Mass.

Kahn, J., M. McConnell, and G. Perez-Quiros (2002): "On the Causes of the Increased Stability of the U.S. Economy," *FRB New York Policy Review*, May, 183–202.

Kahn, J. A. (1987): "Inventories and the Volatility of Production," *American Economic Review*, 77(4), 667–79.

Khan, A., and J. K. Thomas (2007a): "Explaining Inventories: A Business Cycle Assessment Of The Stockout Avoidance and (S,s) Motives," *Macroeconomic Dynamics*, 11(05), 638–664.

——— (2007b): "Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies," *American Economic Review*, 97(4), 1165–1188.

King, R. G., and S. T. Rebelo (1999): "Resuscitating real business cycles," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford.

Kydland, F. E., and E. C. Prescott (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50(6), 1345–70.

Lippi, M., and L. Reichlin (1991): "Trend-Cycle Decompositions and Measures of Persistence: Does Time Aggregation Matter?," *Economic Journal*, 101(405), 314–23.

Maccini, L. J., B. J. Moore, and H. Schaller (2004): "The Interest Rate, Learning, and Inventory Investment," *American Economic Review*, 94(5), 1303–1327.

Maccini, L. J., and A. Pagan (2006): "Inventories, Fluctuations and Business Cycles," NCER Working Paper Series 4, National Centre for Econometric Research.

MaCurdy, T. E. (1983): "A Simple Scheme for Estimating an Intertemporal Model of Labor Supply and Consumption in the Presence of Taxes and Uncertainty," *International Economic Review*, 24(2), 265–89.

Marcellino, M. (1999): "Some Consequences of Temporal Aggregation in Empirical Analysis," *Journal of Business & Economic Statistics*, 17(1), 129–36.

Marcet, A. (1991): "Time Aggregation of Economic Time Series," in *Rational Expectations Econometrics*, ed. by T. J. Sargent, and L. P. Hansen, pp. 237–281. Westview Press.

McCarthy, J., and E. Zakrajsek (2007): "Inventory Dynamics and Business Cycles: What Has Changed?," *Journal of Money, Credit and Banking*, 39(2-3), 591–613.

McConnell, M., and G. Perez-Quiros (2000): "Output Fluctuations in the United States: What Has Changed Since the 1980s?," *American Economic Review*, 90(5), 1464–1476.

Metzler, L. A. (1941): "The Nature and Stability of Inventory Cycles," *The Review of Economics and Statistics*, 23(3), 113–129.

Mills, E. S. (1957): "Expectations and Undesired Inventory," *Management Science*, 4(1), 105–109.

Miron, J. A., and S. P. Zeldes (1988): "Seasonality, Cost Shocks, and the Production Smoothing Models of Inventories," *Econometrica*, 56(4), 877–908.

Ramey, V. A. (1989): "Inventories as Factors of Production and Economic Fluctuations," *American Economic Review*, 79(3), 338–54.

——— (1991): "Nonconvex Costs and the Behavior of Inventories," *Journal of Political Economy*, 99(2), 306–34.

Ravn, M. O., and E. Mazzenga (2004): "International business cycles: the quantitative role of transportation costs," *Journal of International Money and Finance*, 23(4), 645–671.

Ravn, M. O., and H. Uhlig (2001): "On Adjusting the HP-Filter for the Frequency of Observations," CEPR Discussion Papers 2858, C.E.P.R. Discussion Papers.

Richardson, H. (1995): "Control Your Costs then Cut Them," *Transportation and Distribution*.

Rossana, R. J., and J. J. Seater (1992): “Aggregation, Unit Roots and the Time Series Structure on Manufacturing Real Wages,” *International Economic Review*, 33(1), 159–79.

Shibayama, K. (2008): “Inventory Cycles,” Studies in Economics 0804, Department of Economics, University of Kent.

Sims, C. A. (1971): “Discrete Approximations to Continuous Time Distributed Lags in Econometrics,” *Econometrica*, 39(3), 545–63.

Sterman, J. (1989): “Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiemnt,” *Management Science*, 35(3), 321–339.

Stock, J., and M. Watson (2002): “Has the Business Cycle Changed? Evidence and Explanations,” Discussion Paper 9127, National Bureau of Economic Research.

Summers, L. H. (1981): “Comments on Blinder’s Retail Inventory Behavior and Business Fluctuations,” *Brookings Papers on Economic Activity*, 12(1981-2), 513–517.

Wen, Y. (2005): “Understanding the Inventory Cycle,” *Journal of Monetary Economics*, 52(8), 1533–1555.

Wright, B. D., and J. C. Williams (1982): “The Economic Role of Commodity Storage,” *Economic Journal*, 92(367), 596–614.

——— (1984): “The Welfare Effects of the Introduction of Storage,” *The Quarterly Journal of Economics*, 99(1), 169–92.

A Full Solution Of 3-Period Model

The full Kuhn-Tucker optimisation problem is:

$$\begin{aligned}
\max_{\{c_1,c_2,c_3;\iota_1,\iota_2,\iota_3\}} \mathcal{L} = & \ \mathbb{E}\big[\,U(c_1)+\beta U(c_2)+\beta^2 U(c_3)\,\big]-\mathbb{E}[\,\kappa_1\iota_1\,]-\mathbb{E}[\,\kappa_2\iota_2\,]-\mathbb{E}[\,\kappa_3\iota_3\,] \\
& -\mathbb{E}[\,\lambda_1(c_1-[\iota_1-J(\iota_1)]\,a_M-f_1)\,] \\
& -\mathbb{E}[\,\lambda_2(c_2-[\iota_2-J(\iota_2)]\,a_2-(1-\iota_1)a_M)\,] \\
& -\mathbb{E}[\,\lambda_3(c_3-[\iota_3-J(\iota_3)]\,a_M-(1-\iota_2)a_2)\,]
\end{aligned}$$

The necessary FOCs for an equilibrium are:

$$\begin{aligned}
\mathbb{E}[\,U_c'(c_1)\,]-\mathbb{E}[\,\lambda_1\,] &= 0 \\
\mathbb{E}[\,\beta U_c'(c_2)\,]-\mathbb{E}[\,\lambda_2\,] &= 0 \\
\mathbb{E}\big[\,\beta^2 U_c'(c_3)\,\big]-\mathbb{E}[\,\lambda_3\,] &= 0 \\
\mathbb{E}[\,\lambda_1(1-J'(\iota_1)).a_M\,]-\mathbb{E}[\,\lambda_2 a_M\,]-\mathbb{E}[\,\kappa_1\,] &= 0 \\
\mathbb{E}[\,\lambda_2(1-J'(\iota_2)).a_2\,]-\mathbb{E}[\,\lambda_3 a_2\,]-\mathbb{E}[\,\kappa_2\,] &= 0 \\
\mathbb{E}[\,\lambda_3(1-J'(\iota_3)).a_M\,]-\mathbb{E}[\,\kappa_3\,] &= 0 \\
c_1-[\iota_1-J(\iota_1)].a_M-f_1 &= 0 \\
c_2-[\iota_2-J(\iota_2)].a_2-(1-\iota_1).a_M &= 0 \\
c_3-[\iota_3-J(\iota_3)].a_M-(1-\iota_2).a_2 &= 0 \\
\kappa_1\geq 0 &\qquad \iota_1\geq 0 \\
\kappa_2\geq 0 &\qquad \iota_2\geq 0 \\
\kappa_3\geq 0 &\qquad \iota_3\geq 0
\end{aligned}$$

As described in section 3.1 above, the model is solved recursively starting with periods 2 and 3. In period 3, the agent brings forward as much inventory as possible such that the marginal cost of the inventories is 1; c_3 is given by

$$[\iota_3-J(\iota_3)]\,a_M-(1-\iota_2)a_2$$

.

The decision of the agent in period 2 is either one in which agents wish to bring forward consumption to today ($\kappa_2=0,\iota_2>0$), or one in which agents do not wish to bring forward consumption ($\kappa_2>0,\iota_2=0$). The agent will choose to bring forward consumption into period 2 if the marginal utility of consumption in period 2 would otherwise be higher than that in period 3; as ι_2 increases, c_2 increases ($U_c'(c_2)$ decreases) so long as $J'(\iota_2) < 1$ and at the same time c_3 falls which increases $U_c'(c_3)$. The equations for the equilibrium choice in period 2 and 3, assuming $\iota_2 > 0$ are:

$$\begin{aligned}
\beta U_c'(c_2)(1 - J'(\iota_2))\, a_2 - \beta^2 U_c'(c_3)\, a_2 &= 0 \\
1 - J'(\iota_3) &= 0 \\
c_2 - [\iota_2 - J(\iota_2)]\, a_2 - (1 - \iota_1) a_M &= 0 \\
c_3 - [\iota_3 - J(\iota_3)]\, a_M - (1 - \iota_2) a_2 &= 0
\end{aligned}$$

on the other hand, if $\iota_2 = 0$, I solve for c_2, c_3 and ι_3 using:

$$\begin{aligned}
1 - J'(\iota_3) &= 0 \\
c_2 - (1 - \iota_1) a_M &= 0 \\
c_3 - [\iota_3 - J(\iota_3)]\, a_M - a_2 &= 0
\end{aligned}$$

Moving to the first period, the solution depends on the expectation of the marginal utility of consumption in period 2; if consumption is expected to be higher in period 2 (marginal utility is expected to be lower), then the agent will wish to bring forward consumption by setting ι_1 0. The first period solution is a pair c_1 and $\iota_1 \geq 0$ that solves:

$$\begin{aligned}
U_c'(c_1)(1 - J'(\iota_1))\, a_M &= \mathbb{E}[\,\beta U_c'(c_2)\,] a_M + \mathbb{E}[\,\kappa_1\,] \\
c_1 - [\iota_1 - J(\iota_1)]\, a_M - f_1 &= 0 \\
\kappa_1 \geq 0 &\qquad \iota_1 \geq 0
\end{aligned}$$

If $\kappa_1 > 0$ this is a trivial solution of:

$$\begin{aligned}
c_1 &= f_1 \\
\iota_1 &= 0
\end{aligned}$$

However, if $\iota_1 \geq 0$, then $\kappa_1 = 0$ and the optimal choice is a pair c_1 and ι_1 which solves:

$$\begin{aligned}
U_c'(c_1)(1 - J'(\iota_1))\, a_M &= \mathbb{E}[\,\beta U_c'(c_2)\,] a_M \\
c_1 &= [\iota_1 - J(\iota_1)]\, a_M + f_1
\end{aligned}$$

As outlined above, I construct a polynomial approximation of $\Omega \equiv \mathbb{E}[\,\beta U_c'(c_2)\,]$ as follows:

1. Construct a vector of N discrete values between 0 and 1 to represent possible values for ι_1;
2. Assume that $a_2 = a_L$:
   (a) calculate the vector of choices of c_2 for each possible ι_1 choice;
   (b) generate the vector of marginal utilities
3. Repeat for $a_2 = a_M$ and $a_2 = a_H$ to yield, in total, 3 vectors of N marginal utilities corresponding to different choices of ι_1;
4. Use p_L, $p_M = 1 - p_L - p_H$ and p_H to appropriately construct the expected marginal utility for each possible choice of ι_1;
5. Finally, construct a 5th-order polynomial approximation of the relationship between ι_1 and Ω, and replace Ω in the FOCs with this approximation.

B DSGE Model Optimisation

The optimisation problem in period t is then given by:

$$\begin{aligned}
\max_{\{k_{\tau+1},c_\tau,n_\tau\}} \mathbb{U} &= \mathbb{E}_t\left[\sum_{\tau=t}^{\infty}\beta^{\tau}U\left(c_\tau,1-n_\tau\right)\right] \\
s.t.\ \ y_\tau &= a_\tau\left(n_\tau\right)^{\alpha}k_\tau^{1-\alpha} \\
c_\tau+s_{\tau+1}+k_{\tau+1}-(1-\delta)k_\tau &= (\iota_\tau-J\left(\iota_\tau\right))y_\tau+(1-\upsilon)s_\tau+f_\tau-\frac{\varkappa}{2}\left(\frac{k_{\tau+1}-(1-\delta)k_\tau}{k_\tau}\right)^2 \\
f_{\tau+1} &= (1-\iota_\tau)a_\tau\left(n_\tau\right)^{\alpha}k_\tau^{1-\alpha} \\
s_{\tau+1} &\geq 0 \\
\iota_\tau &\geq 0
\end{aligned}$$

which we can treat as a Kuhn-Tucker problem in period t:

$$\max_{\{c_\tau,n_\tau,k_{\tau+1},\iota_\tau,s_{\tau+1}\}}\mathcal{L}=\mathbb{E}_t\sum_{\tau=t}^{\infty}\beta^{\tau-t}.U\left(c_\tau,1-n_\tau\right)-\mathbb{E}_t\sum_{\tau=t}^{\infty}\mu_\tau\left[s_{\tau+1}\right]-\mathbb{E}_t\sum_{\tau=t}^{\infty}\kappa_\tau\left[\iota_\tau\right]$$
$$-\mathbb{E}_t\sum_{\tau=t}^{\infty}\lambda_\tau\left[c_\tau+k_{\tau+1}-(1-\delta)k_\tau-(\iota_\tau-J\left(\iota_\tau\right))a_\tau\left(n_\tau\right)^{\alpha}k_\tau^{1-\alpha}-(1-\upsilon)s_\tau-(1-\iota_{\tau-1})a_{\tau-1}\left(n_{\tau-1}\right)^{\alpha}k_{\tau-1}^{1-\alpha}+\frac{\varkappa}{2}\left(\frac{k_{\tau+1}-(1-\delta)k_\tau}{k_\tau}\right)^2\right]$$

$$
\begin{aligned}
\frac{d\mathcal{L}}{dc_\tau} &: \mathbb{E}_t\left[\beta^{\tau-t}U_c'\left(c_\tau\right)\right] = \mathbb{E}_t\left[\lambda_\tau\right] \quad \forall\tau \\
\frac{d\mathcal{L}}{dn_\tau} &: \mathbb{E}_t\left[\beta^{\tau-t}N_n'\left(1-n_\tau\right)\right] = \mathbb{E}_t\left[\lambda_\tau(\iota_\tau - J\left(\iota_\tau\right))\alpha a_\tau\left(n_\tau\right)^{\alpha-1}k_\tau^{1-\alpha}\right] + \mathbb{E}_t\left[\lambda_{\tau+1}(1-\iota_\tau)\alpha a_\tau\left(n_\tau\right)^{\alpha-1}k_\tau^{1-\alpha}\right] \quad \forall\tau \\
\frac{d\mathcal{L}}{dk_{\tau+1}} &: E_t\left[\lambda_{\tau+1}\left(\left(\iota_{\tau+1} - J\left(\iota_{\tau+1}\right)\right)\left(1-\alpha\right)a_{\tau+1}\left(n_{\tau+1}\right)^\alpha k_{\tau+1}^{-\alpha} + (1-\delta) + \frac{\varkappa}{k_{\tau+1}}\left(\frac{k_{\tau+2}}{k_{\tau+1}}\right)\left(\frac{k_{\tau+2}-(1-\delta)k_{\tau+1}}{k_{\tau+1}}\right)\right)\right] \\
&\quad + E_t\left[\lambda_{\tau+2}(1-\iota_{\tau+1})\left(1-\alpha\right)a_{\tau+1}\left(n_{\tau+1}\right)^\alpha k_{\tau+1}^{-\alpha}\right] = \mathbb{E}_t\left[\lambda_\tau\left(1+\frac{\varkappa}{k_\tau}\left(\frac{k_{\tau+1}-(1-\delta)k_\tau}{k_\tau}\right)\right)\right] \quad \forall\tau \\
\frac{d\mathcal{L}}{d\iota_\tau} &: \mathbb{E}_t\left[\lambda_\tau\left(1-J'\left(\iota_\tau\right)\right)y_\tau\right] = \mathbb{E}_t\left[\lambda_{\tau+1}y_\tau\right] - \mathbb{E}_t\left[\kappa_\tau\right] \quad \forall\tau \\
\frac{d\mathcal{L}}{ds_{\tau+1}} &: \mathbb{E}_t\left[\lambda_\tau\right] - \mathbb{E}_t\left[\mu_\tau\right] = E_t\left[\lambda_{\tau+1}\left(1-\upsilon\right)\right] \quad \forall\tau \\
\frac{d\mathcal{L}}{d\lambda_\tau} &: \mathbb{E}_t\left[c_\tau + s_{\tau+1} + k_{\tau+1} - (1-\delta)k_\tau + \frac{\varkappa}{2}\left(\frac{k_{\tau+1}-(1-\delta)k_\tau}{k_\tau}\right)^2\right] = E_t\left[(\iota_\tau - J\left(\iota_\tau\right))y_\tau + (1-\upsilon)s_\tau + f_\tau\right] \quad \forall\tau \\
\frac{d\mathcal{L}}{d\kappa_\tau} &: \mathbb{E}_t\left[\kappa_\tau\iota_\tau\right] = 0 \quad \forall\tau \\
\frac{d\mathcal{L}}{d\mu_\tau} &: \mathbb{E}_t\left[\mu_\tau s_{\tau+1}\right] = 0 \quad \forall\tau
\end{aligned}
$$

Therefore the equations defining the equilibrium are (using $\tau = t$):

$$\begin{aligned}
U_c'(c_t)\left(1+\frac{\varkappa}{k_t}\left(\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)\right) - E_t\left[\beta^2.U_c'(c_{t+2}).(1-\iota_{t+1}).(1-\alpha)\,a_{t+1}(n_{t+1})^{\alpha}k_{t+1}^{-\alpha}\right] & \\
= E_t\left[\beta U_c'(c_{t+1})\left((\iota_{t+1}-J(\iota_{t+1}))(1-\alpha)\,a_{t+1}(n_{t+1})^{\alpha}k_{t+1}^{-\alpha}+(1-\delta)+\frac{\varkappa}{k_{t+1}}\left(\frac{k_{t+2}}{k_{t+1}}\right)\left(\frac{k_{t+2}-(1-\delta)k_{t+1}}{k_{t+1}}\right)\right)\right] & \\
N_n'(1-n_t) &= U_c'(c_t)(\iota_t-J(\iota_t))\alpha a_t(n_t)^{\alpha-1}k_t^{1-\alpha}+\mathbb{E}_t\left[\beta U_c'(c_{t+1})(1-\iota_t)\alpha a_t(n_t)^{\alpha-1}k_t^{1-\alpha}\right] \\
U_c'(c_t)(1-J'(\iota_t))\,y_t &= \mathbb{E}_t\left[\beta U_c'(c_{t+1}).y_t\right]-\kappa_t \\
U_c'(c_t)-\mu_t &= E_t\left[\beta U_c'(c_{t+1})(1-\upsilon)\right] \\
c_t+s_{t+1}+k_{t+1}-(1-\delta)k_t &= (\iota_t-J(\iota_t))y_t+(1-\upsilon)s_t+f_t-\frac{\varkappa}{2}\left(\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)^2 \\
\ln a_t &= \rho\ln a_{t-1}+\varepsilon_t \\
\kappa_t &\geq 0, \quad \iota_t\geq 0 \\
\mu_t &\geq 0, \quad s_{t+1}\geq 0
\end{aligned}$$

C Numerical Solution Of DSGE Model

Basic Approach

As discussed in the text, I solve the model using the parameterized expectations algorithm (PEA). PEA can be implemented in either a stochastic or non-stochastic fashion. Stochastic PEA works with simulated data and uses the realised values of the target variable as a measure of the expectation. This approach, therefore, depends on outcomes of the exogenous shocks used in the algorithm and it is in this sense that it is stochastic. I use the the non-stochastic method which involves calculating the conditional expectation explicitly using Gaussian quadrature at each point of a selected grid of values for the state variables. The main advantages of this non-stochastic approach are, firstly, that by using the actual conditional expectation, rather than the stochastic realisation, sampling noise is eliminated and a linear, rather than non-linear, regression can be used to estimate the coefficients. Secondly, by carefully choosing the nodes used (in particular by making use of Chebyshev nodes and Chebyshev polynomials), we gain efficiency. Den Haan (2007) discusses these issues in greater depth.

As discussed in the text, I approximate the three expectations that appear in the necessary first order conditions of the model using exponentiated polynomials in the three state variables given by Ω, Θ, and Ψ (see equations (26), (27) and (28) above). This means that, for example using equation (28), I estimate the log of the expectation with polynomial of order $L = (l_z + 1) \times (l_k + 1) \times (l_d + 1)$ in the logs of the state variables, with coefficients given by ψ:

$$\Psi_n \left(a_t, k_t, D_{t,}; \psi\right) \approx \exp\left(P_L\left(\ln(a_t), \ln(k_t), \ln(D_t); \psi\right)\right)$$

where P_L is a L-th order polynomial. As I use Chebyshev polynomials, the total number of coefficients depends on the multiplication of the order of the basis function for each state variable. Therefore, if I use simple 1st order basis functions for each of the state variables ($l_z = l_k = l_d = 1$), the approximant will have eight coefficients to estimate.[34]

Within-Period Solution

For a given period and set of expectation parameters (ω^0,θ^0 and ψ^0), solving the PEA problem with two potentially binding constraints uses the fact that both non-

[34] Namely, the regressors are a constant, $\ln(a_t)$, $\ln(k_t)$, $\ln(D_t)$, $\ln(a_t) \times \ln(k_t)$, $\ln(a_t) \times \ln(D_t)$, $\ln(D_t) \times \ln(k_t)$, and $\ln(a_t) \times \ln(D_t) \times \ln(k_t)$.

negativity constraints cannot bind at the same time (although they may both not bind in a given period). The expectations parameters can be used to give approximations of the expectation terms (equations (26)-(27)) and then the within-period solution is as follows:

1. Assume that the agent does not wish to store any goods but rather wishes to bring forward consumption; $s_{t+1} = 0, \mu_t > 0,\ \kappa_t = 0$, and $\iota_t \geq 0:$

   (a) The intratemporal Euler equation for labour allocation (equation (5)) uses 1 approximation and is given by:

   $$N_n' (1 - n_t) = U_c' (c_t) (\iota_t - J (\iota_t)) \alpha a_t (n_t)^{\alpha - 1} k_t^{1-\alpha} + (1 - \iota_t) \alpha a_t (n_t)^{\alpha - 1} k_t^{1-\alpha} \beta \Psi (a_t, k_t, D_t; \psi)$$

   (b) The intertemporal Euler equation for investment (equation (4)) uses 2 approximations to yield:

   $$U_c' (c_t) \left(1 + \frac{\varkappa}{k_t} \left(\frac{k_{t+1} - (1-\delta) k_t}{k_t}\right)\right) = \beta \Omega (a_t, k_t, D_t; \omega) + \beta^2 \Theta (a_t, k_t, D_t; \theta)$$

   (c) Using $\kappa_t = 0$, and (6), and the approximating function for expected marginal utility, we get the relevant Euler equation as:

   $$U_c' (c_t) (1 - J' (\iota_t)) = \beta \Psi (a_t, k_t, D_t; \psi)$$

   (d) In this case, the budget constraint (equation (8)) can be written as:

   $$c_t + k_{t+1} - (1-\delta) k_t + \frac{\varkappa}{2} \left(\frac{k_{t+1} - (1-\delta) k_t}{k_t}\right)^2 = (\iota_t - J (\iota_t)) y_t + D_t$$

   (e) The 4 equations in (a)-(d) solve the 4 choice variables $(c_t, n_t, k_{t+1}, \iota_t)$ and:

   $$\begin{aligned} y_t &= a_t (n_t)^{\alpha} k_t^{1-\alpha} \\ D_{t+1} &= f_{t+1} = (1 - \iota_t) y_t \end{aligned}$$

   (f) Calculate κ_t from (6); if $\kappa_1 \leq 0$, skip to step 4; else, move to step 2.

2. If $\kappa_t \geq 0$ then set $\iota_t = 0$ ($\implies f_{t+1} = y_t$); and check whether the agent wishes to store any extra goods:

(a) The intratemporal Euler equation for labour allocation uses 1 approximation and is given by:

$$N_n'(1-n_t)=\alpha a_t(n_t)^{\alpha-1}k_t^{1-\alpha}\beta\Psi(a_t,k_t,D_t;\psi)$$

(b) The intertemporal Euler equation for investment again uses 2 approximations to yield:

$$U_c'(c_t).\left(1+\frac{\varkappa}{k_t}\left(\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)\right)=\beta.\Omega(a_t,k_t,D_t;\omega)+\beta^2.\Theta(a_t,k_t,D_t;\theta)$$

(c) Using $\iota_t=0,\mu_t=0$ and (7), and the approximating function for expected marginal utility, we get the relevant Euler equation as

$$U_c'(c_t)=\beta(1-\upsilon)\Psi(a_t,k_t,D_t;\psi)$$

(d) In this case, the budget constraint (equation (8)) can be written as:

$$c_t+s_{t+1}+k_{t+1}-(1-\delta)k_t+\frac{\varkappa}{2}\left(\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)^2=D_t$$

(e) The 4 equations in (a)-(d) solve the 4 choice variables $(c_t,n_t,k_{t+1},s_{t+1})$ and:

$$\begin{aligned}y_t&=a_t(n_t)^\alpha k_t^{1-\alpha}\\D_{t+1}&=y_t+s_{t+1}\end{aligned}$$

(f) Calculate μ_t from (7); if $\mu_t\leq 0$, skip to step 4; else, move to step 3.

3. Set $\iota_t=s_{t+1}=0$;

(a) The intratemporal Euler equation for labour allocation uses 1 approximation and is given by:

$$N_n'(1-n_t)=\alpha a_t(n_t)^{\alpha-1}k_t^{1-\alpha}\beta\Psi(a_t,k_t,D_t;\psi)$$

(b) The intertemporal Euler equation for investment again uses 2 approximations to yield:

$$U_c'(c_t).\left(1+\frac{\varkappa}{k_t}\left(\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)\right)=\beta\Omega(a_t,k_t,D_t;\omega)+\beta^2.\Theta(a_t,k_t,D_t;\theta)$$

(c) In this case, the budget constraint (equation (8)) can be written as:

$$c_t + k_{t+1} - (1-\delta)k_t + \frac{\varkappa}{2}\left(\frac{k_{t+1} - (1-\delta)k_t}{k_t}\right)^2 = D_t$$

(d) The 3 equations in (a)-(c) solve the 3 choice variables (c_t, n_t, k_{t+1}) and:

$$\begin{aligned} y_t &= a_t\,(n_t)^{\alpha}\,k_t^{1-\alpha} \\ D_{t+1} &= y_t \end{aligned}$$

4. The model is solved for period t; repeat for process for next period.

Ensuring 'Good' Approximations

It is not sufficient simply to have an approximation for the expectation. In order to resemble the rational expectations solution, the approximation should lead to a set of beliefs that are consistent with the approximation. Therefore the necessary algorithm to implement the Non-Stochastic PEA solution to my model is:

1. I create a discrete three-dimensional grid of the state space. To do this, I define bounds within which to restrict the grid in each direction and choose q_z, q_k, and q_d as the number of points in each direction of the grid (for TFP, capital and stocks respectively). Within each direction, the nodes are given by Chebyshev nodes; this means that more points toward the bounds are used and this improves the accuracy of the function approximation (Judd 1998).

2. Using an initial estimate for the coefficients of the approximations (ω^0, θ^0, and ψ^0), I solve the model at each grid point using the steps outlined above. Once I have the model solved for each grid point, I can also compute the conditional expectation in equations (26, 27 and 28) using Gauss-Hermite quadrature.

3. I now have both state variables, and the corresponding conditional expectations based on the initial expectation function, for each grid point. I now fit the exponentiated polynomial of the logarithm of the state variables on the logarithm of the three expectations separately using a linear regression to obtain new coefficients given by $\omega^{new}, \theta^{new}$, and ψ^{new}.

4. The parameter vector is updated in the direction of the newly estimated vector:

$$\omega^1 = (1 - ¥)\omega^0 + ¥\omega^{new}$$

where ¥ determines the amount of weight placed on the new estimates.

5. I repeat the procedure until the difference between the old and new estimates is below a chosen tolerance level (I use 0.00001).

Some Specific Difficulties And Solution Decisions

Relative to the standard application of PEA, my expectation functions add one extra level of difficulty. Namely, it is problematic calculating the actual value of the $\Theta(a_t, k_t, D_t; \theta)$ approximation as it contains the expectation of a variable from 2 periods ahead (c_{t+2}) :

$$\Theta(a_t, k_t, D_t; \theta) \approx \mathbb{E}_t \left[U_c'(c_{t+2}) .(1-\iota_{t+1}).(1-\alpha) a_{t+1} (n_{t+1})^{\alpha} .k_{t+1}^{-\alpha} \right]$$

This would involve repeating the full Gauss-Hermite quadrature loop a second time within each loop of the first Gauss-Hermite quadrature. To avoid this extra computational burden, I make use of the expected marginal utility of consumption approximation given by equation (28). Thus, once I have solved for the optimal decision in period t which gives k_{t+1} and D_{t+1}, it is easy to calculate the possible TFP shocks which yield a_{t+1} values and so the range of possible state variables for period $t+1$. Given the states, I repeat the solution to get the optimal decision under each possible shock and calculate the expectation using Gauss-Hermite quadrature as:

$$\begin{aligned} & \mathbb{E}_t \left[U_c'(c_{t+2}) (1-\iota_{t+1}) (1-\alpha) a_{t+1} (n_{t+1})^{\alpha} k_{t+1}^{-\alpha} \right] \\ \approx & \sum_{a_{t+1}^i} \left[\Psi\left(a_{t+1}^i, k_{t+1}, D_{t+1}; \psi\right) \left((1-\iota_{t+1}) (1-\alpha) a_{t+1}^i (n_{t+1})^{\alpha} k_{t+1}^{-\alpha} \right) \right] \end{aligned} \quad (29)$$

The baseline results are reported for a grid that contains a total of 75 points ($q_z = 5$, $q_k = 5$ and $q_d = 3$). I have also used a finer grid (245 points) with no material impact on the equilibrium. The Gauss-Hermite quadrature is performed using 5 nodes. The tolerance level is set to 10^{-5}. I have experimented with Chebyshev polynomials of varying degrees; the baseline results use $l_z = l_k = l_d = 1$.

D Standard RBC Model: Necessary First Order Conditions for Equilibrium

The necessary FOCs for an equilibrium in period t are:

$$
\begin{aligned}
& U_c'(c_t)\left(1+\frac{\varkappa}{k_t}\left(\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)\right)\ldots \\
\ldots &= E_t\left[\beta U_c'(c_{t+1})\left((1-\alpha)\,a_{t+1}(n_{t+1})^{\alpha}k_{t+1}^{-\alpha}+(1-\delta)+\frac{\varkappa}{k_{t+1}}\left(\frac{k_{t+2}}{k_{t+1}}\right)\left(\frac{k_{t+2}-(1-\delta)k_{t+1}}{k_{t+1}}\right)\right)\right] \\
N_n'(1-n_t) &= U_c'(c_t)\,\alpha a_t(n_t)^{\alpha-1}k_t^{1-\alpha} \\
c_t+k_{t+1}-(1-\delta)k_t &= y_t-\frac{\varkappa}{2}\left(\frac{k_{t+1}-(1-\delta)k_t}{k_t}\right)^2 \\
\ln a_t &= \rho\ln a_{t-1}+\varepsilon_t
\end{aligned}
$$

Numerically, this model is easier to solve as there is only one expectation which we need to parameterise, and there are only 2 state variables:

$$\Omega_n^R\left(a_t,k_t;\omega^R\right)=E_t\left[U_c'(c_{t+1})\left((1-\alpha)\,a_{t+1}(n_{t+1})^{\alpha}k_{t+1}^{-\alpha}+(1-\delta)+\frac{\varkappa}{k_{t+1}}\left(\frac{k_{t+2}}{k_{t+1}}\right)\left(\frac{k_{t+2}-(1-\delta)k_{t+1}}{k_{t+1}}\right)\right)\right]$$