Int. J. Production Economics 149 (2014) 47–55

Contents lists available at ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

Time-phased safety stocks planning and its financial impacts: Empirical evidence based on European econometric data

Martin Stößlein [a,*], John Jack Kanet [b], Mike Gorman [b], Stefan Minner [a]

[a] Logistics and Supply Chain Management, School of Management, Technische Universität München, Germany
[b] Department of MIS, OM, and Decision Sciences, School of Business Administration, University of Dayton, OH, USA

ARTICLE INFO

Article history:
Received 15 April 2012
Accepted 26 March 2013
Available online 25 April 2013

Keywords:
Time-phased safety stocks planning
Econometric analysis
Heteroscedasticity tests
Empirical study

ABSTRACT

This paper explores the rationale for planning time-phased safety stocks. We assert that a single safety stock vector for the entire planning horizon (typically based on stationary demand forecast errors and stationary replenishment lead times) may be insufficient for hedging against uncertainties. We argue that planning time-phased safety stocks is prudent when faced with non-stationary demand and/or non-stationary supply. We scrutinize particularly whenever non-stationarity is due to heteroscedastic demand and resulting heteroscedastic demand forecast errors. Consequently, an empirical evidence on a wide basis is provided that such errors for manufactured products are highly heteroscedastic. To test the phenomenon and to estimate its impact at stock keeping unit level, we have conducted an econometric analysis using the EUROSTAT data from 1985 onwards. Specifically, we analyze new industrial orders across various industries and types of goods manufactured in the five largest European economies by using EViews 7.0. To demonstrate which inventory savings can accrue when safety stock levels are deliberately planned to vary in accordance with the observed heteroscedasticity, we estimate potential safety stock savings reusing the same data sets. Our findings indicate that one realization of non-stationarity, i.e., heteroscedastic demand, is indeed pervasive in the European industry. Thus, recognition of this demand nature may add to effective inventory management policies: reducing unnecessary safety stocks, improving service, or both relative to a single-valued safety stock regimen.

1. Introduction

European companies invest and hold inventory worth billions of Euros (e.g., in Germany on average 14% of revenue; Harting, 2005), with around one fifth of inventory assets accounting for safety stocks. Yet, despite this significance, it can be observed in inventory practice that safety stock targets are typically set at the beginning of an operating year to a single appropriate level to be applied for future time periods (after estimating risks).

We argue that a time-invariant safety stock level may not be sufficient to respond to all the variabilities a company is exposed to nowadays. We propose that time-phased targets for each planning period could enhance inventory performance, when deliberately considering non-stationary demand and supply processes. Indeed, various sources of systematically increasing or decreasing variabilities can be identified. On the supply side, capacities change, e.g., factories schedule holidays during the summer, resulting to variabilities in lead times. Transportation lead times may fluctuate more in winter (due to sudden icy roads) than in summer. On the demand side, components and spare-parts providers often face irregular customer orders. Technological advances or reputation crises can persistently destroy 'healthy' demand patterns. Distribution of demand may change during a product's life cycle—in the initial stage, demand is often more variable than later on. Also, unpredictable economic changes can trigger demand changes (e.g., customers may prefer smaller packages during times of low GDP growth rates). Exchange rates can also affect the price of products and, in turn, lead to demand variability.

In practice, however, safety stock levels are commonly calculated on the assumption of stationary processes, that is the size and timing of the increase or decrease can be predicted accurately. It is observed that (time invariant) safety stocks are significantly higher during the operating year than the single-valued target level agreed on at the beginning; additionally, there is usually no system to alert inventory managers to this situation (Mahadevan, 2009, p. 363). Certainly, when replenishment lead time and demand are stationary, constant safety stocks may then suffice.

In this article, we focus on the case of non-stationary demand. We explicitly aim to provide evidence that demand forecast errors systematically vary over time (i.e., heteroscedastic), adding to the

* Corresponding author. Tel.: +49 89 289 28204; fax: +49 89 289 28210.
E-mail addresses: martin.stoesslein@tum.de (M. Stößlein), jkanet1@udayton.edu (J.J. Kanet), mgorman1@udayton.edu (M. Gorman), stefan.minner@tum.de (S. Minner).

0925-5273/$ - see front matter
http://dx.doi.org/10.1016/j.ijpe.2013.03.023

motivation to implement a time-phased safety stock policy. Three specific questions are studied here:

(1) To what degree is non-stationarity in demand forecast errors due to heteroscedastic demand variance in manufacturing industries in the major European economies?
(2) What are the right 'doses' of safety stocks reflecting heteroscedastic demand variance?
(3) Which European industries could have the largest cost-saving potentials switching to a time-phased safety stock policy?

Our empirical study could bring us nearer to our vision of using econometrical findings for inventory planning and control purposes. More specifically, policies for different inventory types may be refined if the econometric knowledge about the non-stationary nature of supply and demand processes can be incorporated.

The rest of the article is organized as follows. Section 2 outlines relevant literature and software applications. Section 3 describes the selected sample data and adjustments. Next, we present our empirical analysis and test results in Section 4. Managerial implications are illustrated with potential safety stock savings by both using an example (Section 5) and industry data (Section 6). Finally, we conclude with a brief discussion and promising directions for future research.

2. Literature review

Safety stock research has been an area of intensive inquiry since Graves (1988) identified a huge research gap, and has been revisited since then, with considerable efforts invested into inventory modeling (e.g. Silver et al., 1998; Grubbström and Tang, 1999; Minner, 2000; Zipkin, 2000; Minner, 1997). However, as we see it, the body of inventory theory literature for treating demand and lead times as random variables is tiny, despite this aspect being identified as a promising and fruitful research topic a decade ago (see Wagner, 2002; Silver, 2008).

Notable exceptions in the literature pick up non-stationary processes. Hillestad and Carrillo (1980) investigate time-varying demand processes in spare-parts planning and repair processes in the US Air Force—(war) scenarios are typically highly dynamic. Speaking of single-item inventories, Song and Zipkin (1993) present optimal policies for non-stationary demand (that follows a Markov-modulated Poisson process) using value-iteration algorithms. Similar demand processes are simulated by Chen and Song (2001) and Abhyankar and Graves (2001). Moreover, the forecast system of a non-stationary ARMA(1,1) demand process and an inventory system is integrated in order to explain causes for the bullwhip effect (Gaalman and Disney, 2006, 2009). Fully observed non-stationary demand and partially observed non-stationary demand (e.g. when the economic environment is unknown) (e.g., Bayraktar and Ludkovski, 2010) can be distinguished. Graves (1999) suggests an adaptive base-stock policy for the same type of inventory system (i.e. single-item), where non-stationary demand follows an integrated moving average process of order (0,1,1). Safety stock required is much greater for stationary demand. While the relationship between safety stock requirements and replenishment lead-time in the stationary case is concave, the relationship in the non-stationary case is convex.

Various contributions explore non-stationary demand in safety stock placement models. For example, Graves and Willems (2008) focus on limited-life (electronic) products that alter quickly from one life-cycle phase to another (while the replenishment lead time is assumed to be deterministic). They assume a periodic review of base-stock replenishment regimen. A constant-service-time policy (that assumes stationary locations for safety stocks) is compared with a dynamic policy (service times are optimized in each period). If the non-stationary demand bound is a concave function, then the optimization problem is equivalent to that for the stationary demand case (for which highly efficient algorithms exist).

Time aspects have recently been integrated into service levels. Helber et al. (2013) propose a delta service level that reflects both the size of backorders and customers' waiting time in the context of dynamic lot sizing. Little's law is used to deduct an average expected waiting time that can be interpreted as a linear representation of the delta service level.

Aiming at improving production scheduling, Vargas (2009) extends the well-known Wagner–Whitin (1958) algorithm to the case when periodic demand density is given and time-varying. Safety stock requirements are implicitly included in planned order quantities whereby the objective is to minimize the sum of expected setup, backorder, and inventory holding costs. The stochastic dynamic programming problem corresponds to solving a shortest path problem. Taking random demand and a target beta service level into account, Tempelmeier and Herpers (2011) propose a heuristic for the dynamic multi-item capacitated lot sizing problem—the stochastic version is built on Maes and van Wassenhove (1988).

Formerly, due to high computational demands, non-stationary inventory policies have not been used in many real-life applications (Silver et al. 1998). However, recent advances in the information technology have motivated some software vendors to refine inventory algorithms to reflect more realistic planning environments.

SAP provides an '*Enhanced Safety Stock Planning*' feature in APO (Advanced Planning and Optimization), which automatically adapts safety stocks based on target service levels using requirement forecasts and actual and planned replenishment lead times. If planned safety stocks in one period exceed those in a previous period, a 'virtual requirement' is generated representing the quantity difference; if lower, requirements are partially covered with excess safety stock from the previous period. However, APO suggests that virtual requirements should be regular and not highly fluctuating (see Hoppe, 2006, p. 311).

With its multi-stage inventory planning (MIPO) module, SmartOps (recently acquired by SAP) offers a periodic review base-stock replenishment policy responding to non-stationary end customer demands. Starting at the end customer node(s), demand and orders are propagated to back up the supply chain so that service level requirements at each node in the network are satisfied. Each location receives a week-specific suggestion for safety stock for a particular product line. The algorithm even breaks inventory down by purpose such as cycle stock, safety stock, and prebuild stock. However, end customer demand is not realized until the due date, so that optimal (in terms of total inventory cost) allocation of safety stocks in the network cannot be modeled. MIPO has led to remarkable improvements in inventory and customer service at Deere & Company (Troyer et al., 2005).

Logility's Voyager Solution Suite (which acquired Optiant Software in 2010) provides time-phased safety stock plans that are fed into SAP APO. The module is based on Graves and Willems' model (2000) (for details see Neale and Willems, 2009). Logility's website claims: 'you can automatically calculate time-phased inventory plans—instead of stockpiling excess products.' Operating improvements have been produced in the management of Microsoft's Entertainment and Devices (E&D) Division (e.g., Xbox).

Inventory performance has been more and more econometrically analyzed—with the majority of research linked to inventory turnover in publicly listed US retailers. So factors are identified explaining inventory productivity and the extent to which they influence corporate performance. Recently, Alan et al. (2012) stress that inventory performance strongly predicts future stock returns. To enhance the value of interpretation, text-book inventory measures have also