

Fill rate in a periodic review order-up-to policy under auto-correlated normally distributed, possibly negative, demand

**Stephen M. Disney¹, Gerard J.C. Gaalman²,
Carl Philip T. Hedenstierna³ and Takamichi Hosoda⁴**

1. Logistics Systems Dynamics Group, Cardiff Business School, Cardiff University,
Aberconway Building, Colum Drive, Cardiff, CF10 3EU, UNITED KINGDOM.

DisneySM@cardiff.ac.uk. Tel: +44 (0) 2920 876310. Corresponding Author

2. Department of Operations, Faculty of Economics and Business, University of Groningen, P.O. Box 800, 9700 AV
Groningen, THE NETHERLANDS.

g.j.c.gaalman@rug.nl

3. Logistics Systems Dynamics Group, Cardiff Business School, Cardiff University,
Aberconway Building, Colum Drive, Cardiff, CF10 3EU, UNITED KINGDOM.

cpthed@gmail.com

4. Graduate School of International Management, Aoyama Gakuin University, 4-4-25, Shibuya,
Shibuya-ku, Tokyo, 150-8366, JAPAN.

hosoda@gsim.aoyama.ac.jp

Abstract

We investigate the inventory service metric known as the fill rate—the proportion of demand that is immediately fulfilled from inventory. The task of finding analytical solutions for general cases is complicated by a range of factors including; correlation in demand, double counting of backlogs, and proper treatment of negative demand. In the literature, two approximate approaches are often proposed. Our contribution is to present a new fill rate measure for normally distributed, auto-correlated, and possibly negative demand. We treat negative demand as returns. Our approach also accounts for accumulated backlogs. The problem reduces to identifying the minimum of correlated normally distributed bivariate random variables. There exists an exact solution, but it has no closed form. However, the solution is amenable to numerical techniques, and we present a custom Microsoft Excel function for practical use. Numerical investigations reveal that the new fill rate is more robust than previous measures. Existing fill rate measures are likely to cause excessive inventory investment, especially when fill rate targets are modest, a strongly positive or negative autocorrelation in demand is present, or negative demands exist. Our fill rate calculation ensures that the target fill rate is achieved without excessive inventory investments.

Key words: Fill rate, Order-up-to policy, ARMA (1,1) demand, Negative demand.

Highlights

- Presents a fill rate for normally distributed, auto-correlated, possibly negative demand
- Solution is based on finding the minimum of bivariate correlated normal random variables
- A comparison with two popular fill rate measures is conducted
- An Excel Add-in allows practitioners easy access to the theoretical results

1. Introduction

The fill rate is a popular measure of inventory service in high volume industries as it directly measures the customer's experience of demand fulfilment. The fill rate is defined as *the proportion of demand fulfilled directly from inventory* (Silver, Pyke and Peterson, 1998: p245; Sobel, 2004; Axsäter, 2000: p57). However, this simple definition hides technical details that are often overlooked. In particular there are issues with; double counting of backlogs, lead times, autocorrelation in demand, cross-correlation between net stock and demand, negative demand, and the distribution of demand and net stock. This paper presents a procedure for identifying the true fill rate obtained in the presence of these complicating factors.

1.1. Contribution

Our contribution is the exact expression for the long run fill rate under auto-correlated, possibly negative demand. It is important to have an exact expression as errors can cause excessive inventory investments or over-optimistic fill rate guidance. Indeed, when demand is negatively or strongly positively auto-correlated excessive fill rates are achieved indicating that an opportunity to reduce safety stocks exists. We extend the definition of the fill rate to be compatible with negative demand. It is a generalisation of the common fill rate definition and will produce identical results for non-negative demand.

Existing fill rate measures provide nonsensical results in the presence of negative demand—either fill rates of over 100% or below 0%. Additionally, simulation results can differ significantly from theoretical guidance. Our proposed approach is mathematically correct and numerically accurate, and gives logical and consistent results. The solution reduces to the identification of the distribution

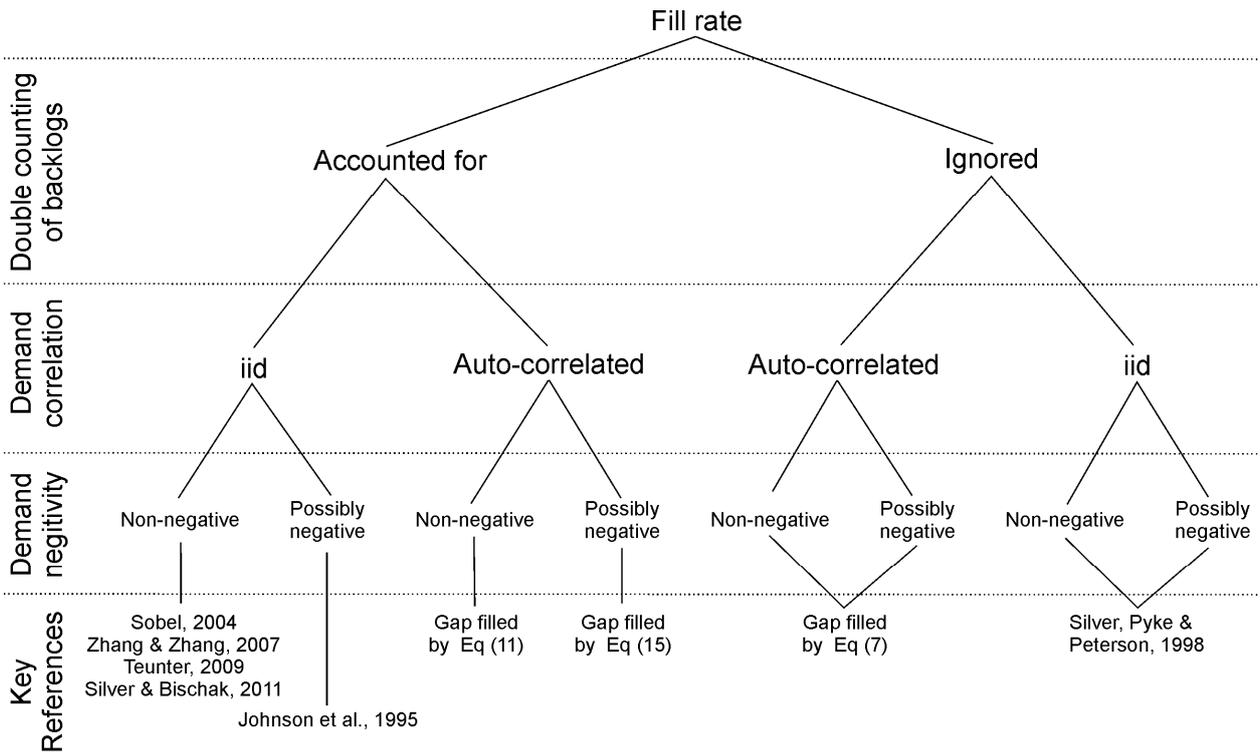


Figure 1. Our contribution to the fill rate literature

of the minimum of two normally distributed correlated random variables. This distribution has an exact solution, but no closed form solution exists. However the problem is amenable to numerical methods. For practical work we provide an Excel Add-in for calculating the true fill rate. We highlight the research gaps and our contribution to the field in Figure 1.

1.2. Motivation

Demand patterns can be both auto-correlated and possibly negative. For example, Figure 2 illustrates a consumer electronics product with a demand that is approximately normally distributed but is not independent and identically distributed (i.i.d.) as there are clear rising and falling trends. It has weekly demand with a mean demand of 146.6 and a standard deviation of 82.7. It also contains two negative demands. Negative demand in a period indicates that the returns from customers are larger than those

delivered. The fitted normal distribution in the density plot has a mean of 150.2 and a standard deviation of 76.7. This was determined by minimising the squared error in the density plot after removing the two outliers that were more than three standard deviations from the mean.

Returns can be significant, particularly in industries such as books, consumer electronics and fashion retailing. We have also noticed that when a large batch of raw materials is checked out of stores and only partially used in production during a period, the remaining raw materials can be returned to the stores in a following period. This procedure can result in a negative demand being recorded in the latter period. Stock adjustments to correct accumulated recording errors can also result in negative demand. Johnson et al. (1995) provide further justification for negative demands.

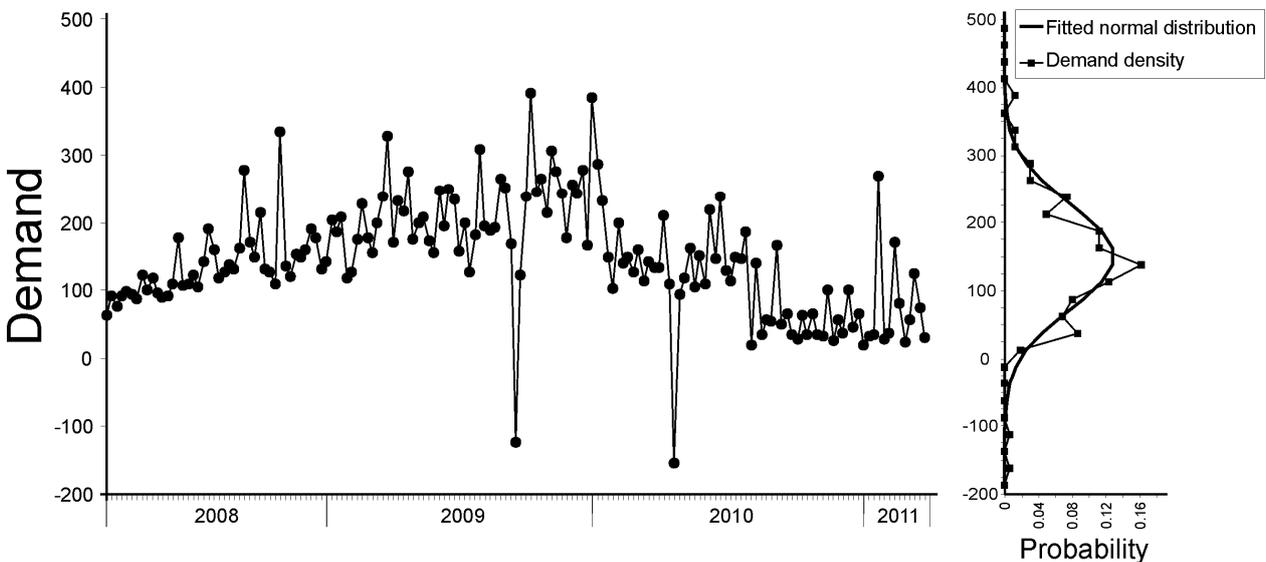


Figure 2. A real-life demand pattern with returns from the consumer electronics industry

Practical fill rate targets are most likely to be above 50%. However, it is mathematically plausible fill rate targets could be anywhere between 0–100%. Saprà, Troung and Zhang (2010) discuss the

inventory withholding strategies of fashion and luxury manufacturers and retailers. This sector has experimented with limiting supply and creating waiting lists to generate a sense of scarcity and exclusivity that may over time increase both the demand volume and the sale price that can be commanded. Here low fill rates are purposely targeted.

1.3. Summary of results

We explore the fill rate in a setting with normally distributed, auto-regressive moving average demands (Box and Jenkins, 1976). We assume that inventory is managed by a linear, discrete time, order-up-to (OUT) replenishment policy and that lead times are arbitrary but constant. We develop our measure analytically and verify its performance via simulation. This reveals that our fill rate is more robust than previous ones, giving accurate predictions over the whole range of fill rates, for any proportion of negative demand, for both i.i.d. and auto-correlated demands. Numerical investigations reveal that our approach is particularly useful when the probability of negative demands is large and fill rates near 100% or 0% are required.

1.4 Structure of the paper

The structure of this paper is as follows. Section 2 reviews literature and highlights the research gap. Section 3 reviews background knowledge of the normal distribution and the distribution of the minimum of two correlated normally distributed random variables. Section 4 considers two fill rate measures from the literature and adapts them for auto-correlated demand. Section 5, the main contribution of the paper, presents a new fill rate measure that is able to cope with normally distributed, possibly negative, correlated demand. In section 6 we illustrate the use of our new fill rate measure for the case of first order auto-regressive moving average (ARMA(1,1)) demand. Section 7 compares

the performance of three fill rate measures, first analytically and then numerically. Section 8 concludes and reflects upon managerial implications.

2. Literature review

We are interested in the fill rate for a single item at a single echelon in a supply chain. This is sometimes referred to as item fill rate, volume fill rate, unit fill rate, or the immediate fill rate (Guijarro Tarradellas, Cardós and Babiloni, 2012). It is different to the order fill rate, which applies to the proportion of fulfilled customer orders that may consist of multiple products (Larsen and Thorstenson, 2014). Schneider (1981), Johnson et al. (1995), Silver and Bischak (2011) and Guijarro Tarradellas, Cardós and Babiloni (2012) provide literature reviews of the fill rate.

The first fill rate measure in the literature is likely to be Hadley and Whitin (1963: p217), although it was not called the fill rate therein. Schneider (1981) reviewed two fill rate measures. The traditional fill rate measure that is common in most text books and a corrected fill rate measure that prevents the double counting of backlogs. Johnson et al. (1995) discussed the double counting issue in fill rate expressions in periodic inventory systems. They identify further issues with stochastic lead-times and order crossovers and discuss the issue of normally distributed, possibly negative, i.i.d. demands. They also considered the link between periodic and continuous review systems. Silver, Pyke and Peterson (1998) and Silver and Bischak (2011) report of another early fill rate derivation by de Kok (2002).

Sobel (2004) studied a fill rate for general demand distributions, assuming positive and i.i.d. demand. Citing Sobel directly "When demand is normally distributed, the new expressions yield an exact formula ... that can be calculated using only the standard normal distribution and density functions". Zhang and Zhang (2007) extended Sobel's approach to arbitrary review periods. Teunter (2009)

presented a shorter derivation of a fill rate measure. Instead of calculating the expected fraction of demand over an infinite horizon, due to the renewal property, he showed a fill rate calculation in a single arbitrary period is sufficient. Assuming positive demand, the expected demand satisfied immediately can be written as the expected inventory on-hand after the order arrived in a period minus the expected inventory on-hand at the end of this period. Kwon, Kim and Baek (2006) considered serial inventory systems where demand is assumed to be normally distributed. The authors discussed the traditional fill rate and the work of Johnson et al. (1995), but recommended the use of the expression derived by Sobel (2004) for normally distributed demand. They also considered a multi-stage supply chain. Silver and Bischak (2011) provide an intuitive derivation of the fill rate expression under normally distributed i.i.d. demand. They avoid double counting by restricting the amount of backlog if the total demand is larger than the order up to level.

In summary, it is usual in the fill rate literature to assume i.i.d. stochastic demand, which is commonly modelled as a Poisson, Erlang, normal, gamma or binomial distribution. Most research in this field develops fill rate expressions assuming the demand is positive. Normally distributed demand is used to find attractive expressions under the assumption that negative demand has a negligible influence. We adapt Sobel's expression for normal demand to find an approximation of the fill rate in the correlated demand case. We also adapt the traditional fill rate for correlated demand. However our main contribution is the exact fill rate for auto-correlated, normally distributed, possibly negative demand case. We obtain this by a new approach based on the distribution of the minimum of bivariate, correlated normal random variables.

3. Preliminary matter

To investigate the influence of auto-correlated demand on the fill rate we have made a number of assumptions. We assume that demand is normally distributed and the inventory control system is described by linear difference equations. As such, all system variables will be normally distributed and can take on real values between $-\infty$ and $+\infty$. Thus, it is useful to define certain relationships associated with the normal distribution. The probability density function (pdf) of the standard normal distribution of a variable x is

$$\varphi[x] = \exp[-x^2/2] / \sqrt{2\pi}. \quad (1)$$

The cumulative distribution function (cdf) of the standard normal distribution is given by,

$$\Phi[x] = \int_{-\infty}^x \varphi[z] dz = (1 + \operatorname{erf}[x/\sqrt{2}]) / 2. \quad (2)$$

The standard normal loss function is given by,

$$L[x] = \int_x^{\infty} \varphi[z](z-x) dz = \varphi[x] - x(1 - \Phi[x]). \quad (3)$$

Cain (1994) and Nadarajah and Kotz (2010) build upon Basu and Ghosh (1978) and Nagaraja and Mohan (1982) to provide the following expression for the pdf of the minimum of two normally distributed, correlated random variables,

$$\xi_{\min}[x] = \psi[x] + \zeta[x], \quad (4)$$

where $\psi[x] = \frac{1}{\sigma_1} \varphi\left[\frac{x-\mu_1}{\sigma_1}\right] \Phi\left[\frac{\rho(x-\mu_1)}{\sigma_1\sqrt{1-\rho^2}} - \frac{x-\mu_2}{\sigma_2\sqrt{1-\rho^2}}\right]$ and $\zeta[x] = \frac{1}{\sigma_2} \varphi\left[\frac{x-\mu_2}{\sigma_2}\right] \Phi\left[\frac{\rho(x-\mu_2)}{\sigma_2\sqrt{1-\rho^2}} - \frac{x-\mu_1}{\sigma_1\sqrt{1-\rho^2}}\right]$, with

$\{\mu_1, \sigma_1\}$ and $\{\mu_2, \sigma_2\}$ being the mean and standard deviation of x_1 and x_2 respectively. Equation

(4) also contains the Pearson correlation co-efficient ρ , $-1 \leq (\rho = \operatorname{cov}(x_1, x_2) / (\sigma_1\sigma_2)) \leq 1$ which

captures the correlation between the normally distributed random variables x_1 and x_2 . The maximum

operator, $(x)^+ = \max[x, 0]$ and the expectation operator, $E[x] = \bar{x}$ are also used.

4. Existing fill rate measures

We consider two fill rate measures from the literature. First, the traditional fill rate measure which is known to be an approximation as it ignores the double counting problem, Cachon and Terwiesch (2006: p391). Second, the fill rate measure from Sobel (2004), which is exact when demand is positive. We first provide, for both cases, the fill rate expressions under i.i.d. normally distributed demand. Then we show how one might adjust these two measures to accommodate auto-correlated normally distributed demand processes.

4.1 The traditional fill rate measure

Cachon and Terwiesch (2006: p198, p257) outline the common approach to calculating the fill rate as

$$\beta_T = 1 - \frac{E[(-ns_t)^+]}{E[d_t]}. \quad (5)$$

Like Johnson et al. (1995), we term it the ‘traditional’ fill rate and use the subscript T to denote this measure. In (5) ns_t is the net stock in time period t (net stock is the inventory on-hand minus backlogged demand) and d_t is the demand. Equation (5) computes the expected inventory short, rather than the expected unfulfilled demand and as such this measure is an approximation (Hadley and Whitin, 1963; Johnson et al., 1995). The shortage is accumulative and backlogs can persist in the system for more than one period, leading to a double counting problem when the lead time is positive. In periods where the backlog exceeds demand, some of the backlogged quantity must have incurred in a previous period. In such periods, the unfulfilled demand is only the current demand, not the current backlog. Due to this double counting, the expected backlog overestimates the missed demand. This causes (5) to become a lower bound of the true fill rate. Practically this means that the

safety stock guidance is too high, recommending an excessive investment in safety stock. Nevertheless the traditional fill rate is reasonably accurate when the fill rate is near 100% and when the probability of negative demands is negligible. However, when the achieved fill rate is more modest, the errors can become large and, in extreme cases, this measure can become negative—a nonsensical result. Attempts to adjust this measure for the double counting of backlogs can be found in the literature (Schneider, 1981; Johnson et al., 1995; Sobel, 2004; Silver and Bischak, 2011).

We assume that the linear order-up-to (OUT) policy with minimum mean square error (MMSE) forecasting is present. Under i.i.d. demand drawn from a normal distribution with mean μ_d and standard deviation σ_d , the net stock, ns_t , is normally distributed with expected value of μ_{ns} and standard deviation $\sigma_{ns} = \sqrt{1+T_p}\sigma_d$. Here T_p is the replenishment lead time. In this case, the standard normal loss function for the expected backlog holds. The traditional fill rate then becomes

$$\beta_T = 1 - \sqrt{1+T_p}\sigma_d L\left[\frac{\mu_{ns}}{\sqrt{1+T_p}\sigma_d}\right] / \mu_d \quad (6)$$

The net stock and demand are linked together via $ns_t = ns_{t-1} - d_t + o_{t-(T_p+1)}$, the inventory balance equation. Here o_t is the order placed at time t . Notice that the traditional fill rate is only valid for i.i.d. demand. To account for auto-correlated demand (or other forecasting and/or replenishment policies), a natural extension of (5) would be,

$$\beta_T^* = 1 - \sigma_{ns} L[\mu_{ns} / \sigma_{ns}] / \mu_d. \quad (7)$$

Equation (7) is an important contribution which we will investigate further in Section 7. Notice that (7) is different from (6). Equation (6) assumes i.i.d. demand and sets the variance of the net stock to $\sigma_d^2(1+T_p)$. Equation (7) acknowledges that the variance of the net stock levels, σ_{ns}^2 could take a different form. In fact, with ARIMA type demand, MMSE forecasting, and the linear OUT policy,

the variance of the net stock is given by $\sigma_{ns}^2 = \sigma_\varepsilon^2 \sum_{n=0}^{T_p} \left(\sum_{t=0}^n \tilde{d}_t \right)^2$, Gaalman and Disney (2009). Here

σ_ε^2 is the variance of the error—the white noise driving the demand process—and \tilde{d}_t is the auto-covariance function of demand at lag t . \tilde{d}_t is also equivalent to the impulse response of the demand process, Gaalman and Disney (2012). Equation (7) takes the influence of the correlated demand on the variance of the net stock levels into account and we use the *star* in β_T^* to draw attention to this.

4.2. Sobel's fill rate

Sobel (2004) considered that the fulfilled demand in a period is given by

$$f_t = \min \left\{ d_t, (ns_t + d_t)^+ \right\}, \quad (8)$$

where $ns_t + d_t$ is the net stock after the orders placed $T_p + 1$ periods ago have been received, but before the demand has been satisfied. The term $(ns_t + d_t)^+$ reflects the on-hand inventory available after the order arrived to satisfy demand in a given period. Then the fulfilled demand is simply the minimum of the period's demand and our ability to satisfy it. This approach avoids the double-counting problem. Sobel then defines the fill rate as

$$\beta_S = E[f_t] / E[d_t]. \quad (9)$$

The subscript S (not to be confused with the order-up-to level S) denotes the Sobel fill rate, which is exact when demand is i.i.d. and always positive. Sobel (2004) derives an expression for the fill rate from the cumulative distribution of the demand over the lead time and review period minus the cumulative distribution of demand over the lead time, both with an upper limit of the order up to level, $S = \mu_{ns} + \mu_d (T_p + 1)$. S is a constant order up to level. This expression is applied to gamma and

normally distributed demand. As we consider the normal demand we show Sobel's fill rate expression for normal i.i.d. demand (translated into the notation of this paper),

$$\beta_S = \mu_d^{-1} \int_0^S \left(\Phi \left[\frac{(x - S + \mu_{ns+d})}{\sigma_{ns+d}} \right] - \Phi \left[\frac{(x - S + \mu_{ns})}{\sigma_{ns}} \right] \right) dx. \quad (10)$$

The OUT policy will, under i.i.d. demand, produce an expected net stock μ_{ns} and standard deviation of the net stock $\sigma_{ns} = \sqrt{1 + T_p} \sigma_d$. The expected value and standard deviation of $(ns_t + d_t)$ satisfy $\mu_{ns+d} = \mu_{ns} + \mu_d$ and $\sigma_{ns+d} = \sigma_d \sqrt{T_p}$. In (10) the first cdf also equals the inventory on-hand after the order has been arrived at the beginning of the cycle, and the second cdf equals the inventory at the end of the cycle, Teunter (2009). β_S is the complement of the proportion of unfulfilled demand to the mean demand. The expected unfulfilled demand is the difference between the expected shortage during the lead time and review period and the expected shortage during the lead time. Sobel (2004) also provides a lengthy expression for β_S , which is based on standard normal pdf and cdf functions. Using just the standard normal loss function we derived a compact fill rate expression,

$$\beta_S^* = \mu_d^{-1} \left(\frac{\sigma_{ns+d} \left(L \left[-\mu_{ns+d} (\sigma_{ns+d})^{-1} \right] - L \left[\mu_d T_p (\sigma_{ns+d})^{-1} \right] \right)}{\sigma_{ns} \left(L \left[-\mu_{ns} (\sigma_{ns})^{-1} \right] - L \left[\mu_d (T_p + 1) (\sigma_{ns})^{-1} \right] \right)} \right). \quad (11)$$

This expression can be used for i.i.d. as well as correlated demand when the standard deviations of $\{\sigma_{ns}, \sigma_{ns+d}\}$ are appropriately updated. In (11) we have again used the *star* notation in β_S^* to highlight that the consequences of non-i.i.d. demand are accounted for. We investigate the performance of (10) and (11) in Section 7. As (10) and (11) deems negative demand to be fulfilled in extreme cases, $\beta_S < 0\%$ can occur, see case 1 in Table 2.

5. Fill rates with auto-correlated normally distributed, possibly negative, demand

Johnson et al. (1995) and Guijarro Tarradellas, Cardós and Babiloni (2012) argue that the condition for positive demand during a cycle must be explicitly taken into account to correctly determine the fill rate. We relax the assumption of non-negative demand, by letting negative demands denote net returns from customers. Since negative demands should not count towards the fulfilled demand, we define the fulfilled demand, f_t^- as,

$$f_t^- = \min\left\{(d_t)^+, (d_t + ns_t)^+\right\} = \left(\min\{d_t, d_t + ns_t\}\right)^+ . \quad (12)$$

In (12), the demand that can be satisfied in a single period is $(d_t)^+$, which becomes zero for negative demand. The term $(ns_t + d_t)^+$ deals with the double counting issue. This implies that if the net stock at the end of the period is positive, then all demand must have been satisfied. If demand was negative (due to returns), then the fulfilled demand is zero. If the net stock was negative at the end of the period, then the fulfilled demand in the period is equal to the positive part of the sum of the demand and the net stock at the end of the period. We use the superscript in f_t^- to make clear that we have accounted for negative demand.

Fulfilled demand, f_t^-		Net stock at the end of the period	
		$ns_t \geq 0$	$ns_t < 0$
Demand during the period	$d_t > 0$	$f_t^- = d_t = (d_t)^+$	$f_t^- = (ns_t + d_t)^+$
	$d_t \leq 0$	$f_t^- = 0 = (d_t)^+$	$f_t^- = 0 = (d_t)^+$

Table 1. Logic table to determine the fulfilled demand

The logic behind (12) can be verified with Table 1, describing fulfilled demand for all possible combinations of positive/negative demand and positive/negative net stock levels. For example, if the

net stock position at the end of the period was -2 and demand was 10, then 8 units of demand were satisfied immediately. However if the net stock level was -12 at the end of the period and demand was 10, then none of the current demand could have been satisfied.

The fill rate definition needs to be stated carefully when demand can be negative. We let the fill rate, β^* , reflect *the proportion of immediately satisfied demand to the demand that can be satisfied*. This is consistent with the established fill rate definition, as it is only possible to satisfy positive demand.

With this definition, the fill rate is

$$\beta^* = \frac{E[f_t^-]}{E[(d_t)^+]}. \quad (13)$$

We consider a linear inventory system and stationary ARMA demand. This implies that all variables are linear, unbounded and normally distributed. In the linear OUT policy the orders are solely a function of the inventory and the state of the demand process at time t and the distributions of ns_t and $(d_t + ns_t)$ are time-invariant.

Recall from Section 3 the pdf of the minimum of two normally distributed, correlated random variables. Let the minimum in (12) be $x = \min\{d_t, d_t + ns_t\}$. Due to the normal distribution of both variables, $-\infty < x < \infty$. Also, the positive component of (12) equals $f_t^- = x^+$. Thus, the expected value of f_t^- satisfies

$$E[f_t^-] = \int_0^\infty x(\xi_{\min}[x])dx, \quad (14)$$

and β^* can be expressed as

$$\beta^* = \frac{\int_0^\infty x(\xi_{\min}[x])dx}{\int_0^\infty x\sigma_d^{-1}\varphi[(x-\mu_d)/\sigma_d]dx} = \frac{\int_0^\infty x(\xi_{\min}[x])dx}{\sigma_d L[\mu_d/\sigma_d] + \mu_d}. \quad (15)$$

Equation (15)—the main result of this paper—can be obtained using the following approach: In the numerator, the pdf of $\min\{d_t, d_t + ns_t\}$ is given by $\xi_{\min}[x]$. Integrating $x\xi_{\min}[x]$ over positive x captures the expected fulfilled demand, $\int_0^\infty x(\xi_{\min}[x])dx = E\left[(\min\{d_t, d_t + ns_t\})^+\right] = E[f_t^-]$. The expected value of the positive demand $\sigma_d L[\mu_d/\sigma_d] + \mu_d = \int_0^\infty x\sigma_d^{-1}\varphi[(x-\mu_d)/\sigma_d]dx = E[(d_t)^+]$ is given in the denominator.

For general correlated demand there does not appear to be an easily obtainable solution to the integral in the numerator of (15), see Basu and Ghosh (1978). Thus, we need numerical techniques to calculate β^* . This is easy to do using mathematical software, such as Matlab or Mathematica. We have also developed a Microsoft Excel Add-in (the source code is provided in Appendix A) for practical use in the absence of specialist software.

The Pearson correlation coefficient in (15) is represented by $\rho = \text{cov}(ns_t + d_t, d_t)/(\sigma_{ns+d}\sigma_d)$. The covariance is easily obtained from the product of the demand and the (net stock + demand) impulse responses of the system. The impulse response of i.i.d. demand is,

$$d_0 = 1 \text{ otherwise } \forall t, d_t = 0 \quad (16)$$

and for the OUT policy, $ns_0 + d_0 = 0$, implying that $\rho = 0$. This may not be the case for non-i.i.d. demand or if a different forecasting and replenishment system is used. In the next section, we will show how to compute $ns_t + d_t$ and ρ for an ARMA(1,1) demand process with MMSE forecasts.

6. ARMA(1,1) demand and the correlation between net stock and demand

Moving from the general situation to a specific case, we consider the linear OUT policy reacting to ARMA(1,1) demand. This allows us to illustrate how to evaluate the three fill rate measures. ARMA(1,1) demand has been found to represent long life cycle products, such as home care products (Disney et al. 2006), fuel, food products and machine tools (Nahmias, 1993). The mean centred ARMA(1,1) demand (Box and Jenkins, 1976) is described as,

$$d_t = \mu_d + \phi(d_{t-1} - \mu_d) - \theta\varepsilon_{t-1} + \varepsilon_t, \quad (17)$$

where $\varepsilon_t \in N(0, \sigma_\varepsilon^2)$ is an i.i.d. normally distributed random variable with a mean of zero and a variance of σ_ε^2 . The mean demand is μ_d , $-1 < \phi < 1$ is the auto-regressive parameter and $-1 < \theta < 1$ is the moving average parameter. When $\theta = \phi$ an i.i.d. white noise demand pattern is produced.

To preserve normality of the system variables, we assume the existence of a linear OUT replenishment system, allowing one to obtain the mean and variance of the system state states. Thus, negative inventory levels represent backlogs, negative demand indicates net customer returns, negative orders represent returns to suppliers, no capacity limits exist, and what was ordered is received after a constant and known lead time. The system operates in discrete time, and all system variables take continuous values. For example, inventory is observed, fill rates are measured and orders are placed on integer moments of time, but orders and inventory can take on any value on the real number line. This works well for products that are sold by volume, weight or length, but there will be some quantization issues when units (single items or boxes / batches of items) are sold. However, when the average demand becomes sufficiently large compared to the batch size this quantization error becomes insignificant.

The sequence of events is as follows: During the period, previously ordered goods are received and demand is satisfied. At the end of each period inventory is observed, fill rates are measured and new production (replenishment) orders are calculated, see Figure 3.

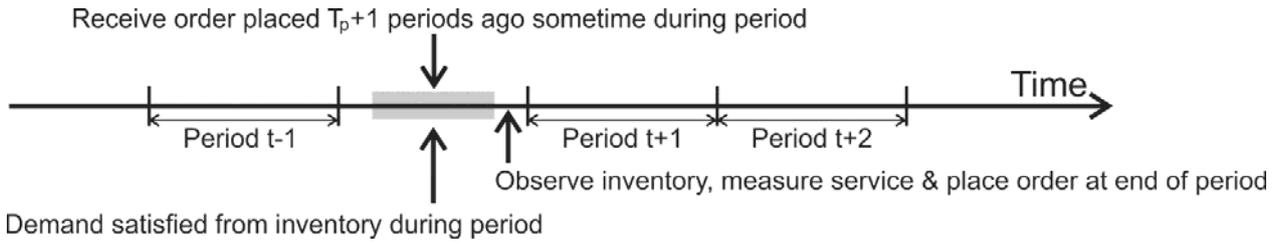


Figure 3. The sequence of events in the OUT policy

Box and Jenkins (1976) show that the impulse response of the ARMA(1,1) demand process is given by,

$$d_t = \phi^{t-1}(\phi - \theta), \quad (18)$$

from which the stationary variance of the demand (in steady state over an infinite time horizon) can be easily be obtained,

$$\sigma_d^2 = \sigma_\varepsilon^2 \sum_{t=0}^{\infty} (\phi^{t-1}(\phi - \theta))^2 = \sigma_\varepsilon^2 (1 + (\phi - \theta)^2 (1 - \phi)^{-1}). \quad (19)$$

The linear OUT policy generates replenishment orders at time t ,

$$o_t = \hat{d}_{t,t+T_p+1} + \left(\mu_{ns} + \sum_{i=1}^{T_p} \hat{d}_{t,t+i} - ns_t - \sum_{i=1}^{T_p} o_{t-i} \right), \quad (20)$$

where, as before, T_p is a nonnegative integer is the replenishment lead-time and μ_{ns} is the target net stock—the expected value of the net stock.

The OUT replenishment policy requires two forecasts (Hosoda and Disney, 2009). One of these forecasts is a prediction of demand over lead-time, $\sum_{i=1}^{T_p} \hat{d}_{t,t+i}$, made at time t ,

$$\sum_{i=1}^{T_p} \hat{d}_{t,t+i} = (1 - \phi^{T_p})(1 - \phi)^{-1} (\phi(d_t - \mu_d) - \theta \varepsilon_t) + \mu_d T_p. \quad (21)$$

The other forecast is a prediction of the demand in the period after the lead-time, $\hat{d}_{t,t+T_p+1}$, made at time t ,

$$\hat{d}_{t,t+T_p+1} = \phi^{T_p} (\phi(d_t - \mu_d) - \theta \varepsilon_t) + \mu_d. \quad (22)$$

Note that the order-up-to level, $S = \hat{d}_{t,t+T_p+1} + \sum_{i=1}^{T_p} \hat{d}_{t,t+i} + \mu_{ns}$, is now a function of the dynamic forecasts, which vary with the state of the demand process (d_t, ε_t) . However, the mean inventory, μ_{ns} , is still a constant as the forecast errors over the lead-time and review period are i.i.d. and time-invariant. Finally, the net stock balance equation completes the OUT policy,

$$ns_t = ns_{t-1} - d_t + o_{t-(T_p+1)}, \quad (23)$$

where ns_t is the net stock at time t and $o_{t-(T_p+1)}$ is the order placed $T_p + 1$ periods ago. The '+1' is the sequence of events delay, which is always present in discrete time systems. The impulse response of the net inventory levels (Gaalman and Disney, 2009 and 2012) is given by,

$$ns_t = \begin{cases} (\phi - \theta)(\phi^t - 1)(1 - \phi)^{-1} - 1 & 0 \leq t \leq T_p, \\ 0 & \text{otherwise,} \end{cases} \quad (24)$$

which we may use to find the stationary variance of the inventory levels,

$$\begin{aligned} \sigma_{ns}^2 &= \sigma_\varepsilon^2 \sum_{t=0}^{T_p} ((\phi - \theta)(\phi^t - 1)(1 - \phi)^{-1} - 1)^2 \\ &= \sigma_\varepsilon^2 \left(\frac{T_p (\theta - 1)^2}{(\phi - 1)^2} + \frac{2\phi(1 + \phi^{T_p}(\theta - 1)(\phi - \theta))}{(\phi - 1)^3} + \frac{(\theta - \phi)^2 \phi^{2+2T_p} + \theta\phi(\theta(2 + \phi) - 2 - 4\phi) - 1}{(\phi - 1)^3 (1 + \phi)} \right). \end{aligned} \quad (25)$$

Notice (25) is non-decreasing in T_p . From (18) and (24) it is clear that,

$$ns_t + d_t = \begin{cases} (\phi - \theta)(\phi^t - 1)(1 - \phi)^{-1} - 1 + \phi^{t-1}(\phi - \theta) & 0 \leq t \leq T_p, \\ \phi^{t-1}(\phi - \theta) & \text{otherwise.} \end{cases} \quad (26)$$

From (26), as $ns_0 + d_0 = 0$, the stationary (long run) variance of $ns_t + d_t$ can be calculated as,

$$\begin{aligned} \sigma_{ns+d}^2 &= \sigma_\varepsilon^2 \left(\sum_{t=1}^{T_p} \left((\phi - \theta)(\phi^t - 1)(1 - \phi)^{-1} - 1 + \phi^{t-1}(\phi - \theta) \right)^2 + \sum_{t=T_p+1}^{\infty} \left(\phi^{t-1}(\phi - \theta) \right)^2 \right) \\ &= \sigma_\varepsilon^2 \left(\frac{T_p(\theta - 1)^2}{(\phi - 1)^2} + \frac{2(\theta - 1)\phi^{T_p}(\phi - \theta)}{(\phi - 1)^3} - \frac{(\theta - \phi)^2 \phi^{2T_p}}{\phi^2 - 1} + \frac{(\theta - \phi)(\theta - 2 + \phi(2\theta - 1) + \phi^{2T_p}(\theta - \phi))}{(\phi - 1)^3(1 + \phi)} \right), \end{aligned} \quad (27)$$

which has two components. The first is due to the net stock and is a sum of random components over the lead-time and review period. The second component, due to the demand, has random components over the whole time horizon. The Pearson correlation coefficient, ρ , can be obtained using the impulse responses (18) and (26) as well as the square root of the variances (19) and (27).

$$(\sigma_{ns+d} \sigma_d) \rho = \sum_{t=0}^{\infty} (ns_t + d_t)(d_t) = \frac{(\theta - \phi)^2 \phi^{2T_p}}{1 - \phi^2} + \frac{(\phi - \theta)(\phi^{T_p} - 1)(1 - \phi^{T_p+1} + \theta(\phi^{T_p} - \phi))}{(\phi - 1)^2(1 + \phi)}. \quad (28)$$

Depending of the parameter values of the demand process, the correlation coefficient may influence the fill rate, see Figure 4. We notice two (possibly three) curves with zero correlation. One of these curves is the i.i.d. case where $\phi = \theta$. The correlation is positive and increasing when $\theta > \phi$. When $\phi > \theta$ the correlation is negative, first decreasing and then increasing to another curve of zero correlation. Further to the right of the second curve, the correlation is positive and increasing. As the lead-time increases, this interval becomes smaller. There is an also odd-even lead-time effect near $\{\phi, \theta\} \approx -1$, resulting in an additional curve with zero correlation for even lead times. The industrially prevalent ARMA(1,1) coefficients (Disney et al. 2006), $\{\phi > 0.5, \theta > 0.5\}$ exhibit a negative correlation between d_t and $(d_t + ns_t)$.

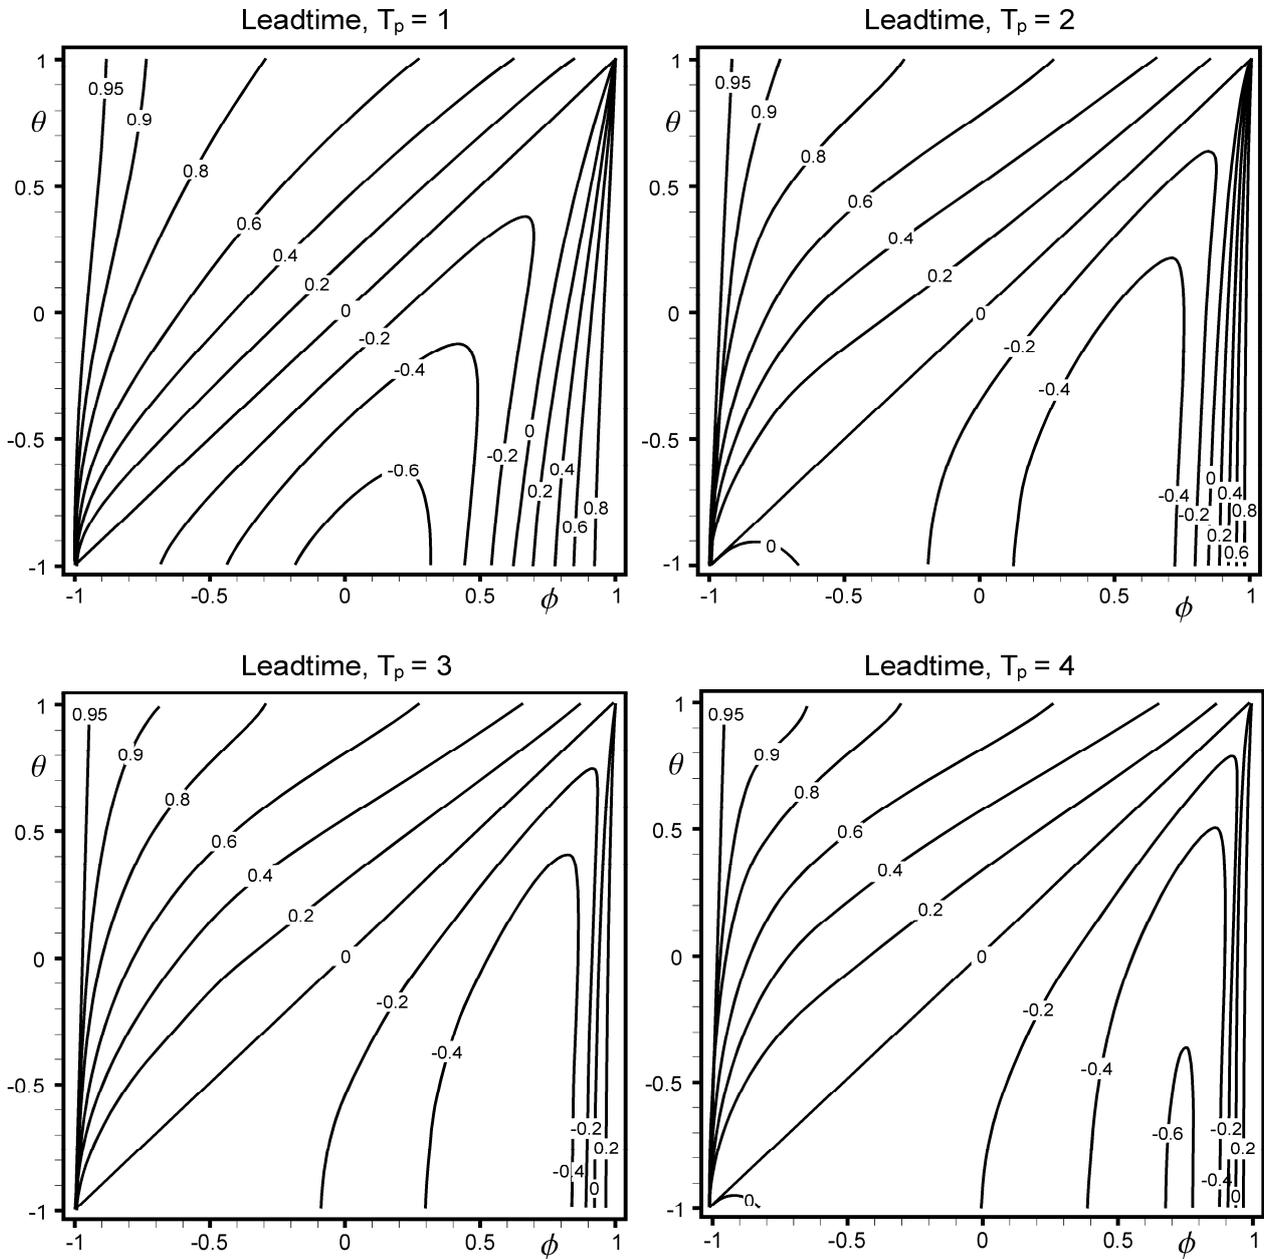


Figure 4. The Pearson correlation coefficient for the OUT policy under ARMA(1,1) demand

7. Investigation of fill rate measures

In this section we investigate the performance of the three fill rate measures for auto-correlated demand. First we consider the analytical performance of the measures for both i.i.d. and auto-

correlated demand. Then we verify the performance of each measure by comparing the analytical fill rate results to simulation outputs.

7.1. Analytical performance of the three fill rate measures

Let us investigate the performance of the various fill rate measures. Throughout section 7 we assume the lead-time $T_p = 1$, and we have scaled the variance of the white noise process to ensure unit demand variance unless otherwise stated. This allows for a fair comparison of the fill rate measures, as different ARMA parameters produce demand with different variances (and hence, would have different probabilities of negative demands).

First, consider the case of i.i.d. demand, see Figure 5. Here, the upper three plots detail the fill rate with different μ_d values, and the lower three plots highlight the difference between the true β^* fill rate, and the approximations, $\{\beta_T, \beta_S\}$. As we are now considering i.i.d. demand there is no star in the superscripts of $\{\beta_T, \beta_S\}$. Here, the x-axis is μ_{ns} , the average net stock, highlighting the influence of the safety stock on the fill rate. Negative μ_{ns} values are allowed and these do not necessarily imply a negative order-up-to level, S . Recall, $S = \mu_{ns} + \mu_d (T_p + 1)$ implying that $S > 0$ if $\mu_{ns} > -\mu_d (T_p + 1)$. We see that the traditional fill rate, β_T , is indeed a lower bound whose accuracy improves as the probability of a backlog carrying over from one period to the next reduces when the fill rate approaches 100%. β_S does well when the probability of negative demand is low (when $\mu_d \geq 3$), but it experiences noticeable errors when $\mu_d = 1$ (where 15.8% of periods have negative demand), even producing some negative fill rates. Notably β_S drops below β_T in some settings. This is a consequence of the negative demand as the following reasoning shows: When demand is non-

negative $\beta_T \leq \beta^*$ (due to double counting of backlogs) and $\beta_S = \beta^*$ implying $\beta_T \leq \beta_S$. However relaxing the assumption of non-negative demand leads to cases where these relationships no longer hold. Thus, negative demand is a necessary, but not sufficient condition for $\beta_T > \beta^*$, $\beta_S < \beta^*$ and $\beta_T > \beta_S$ (see also Table 2 in Section 7.2). An alternative explanation can be given as follows. The traditional fill rate measure can be expressed as

$$\beta_T = \mu_d^{-1} \left(\sigma_{ns+d} L \left[-\mu_{ns+d} \sigma_{ns+d}^{-1} \right] - \sigma_{ns} L \left[-\mu_{ns} \sigma_{ns}^{-1} \right] - \sigma_{ns+d} L \left[\mu_{ns+d} \sigma_{ns+d}^{-1} \right] \right). \quad (29)$$

From (11) and (29) the difference $\beta_S - \beta_T$ can be obtained,

$$\beta_S - \beta_T = \mu_d^{-1} \left(\sigma_{ns} L \left[\mu_d (1+T_p) \sigma_{ns}^{-1} \right] - \sigma_{ns+d} L \left[\mu_d T_p \sigma_{ns+d}^{-1} \right] + \sigma_{ns+d} L \left[\mu_{ns+d} \sigma_{ns+d}^{-1} \right] \right). \quad (30)$$

(30) can be both positive and negative and can be used to determine necessary and sufficient conditions for $\beta_S > \beta_T$. The first two terms are from the upper bound of S in Sobel's integration, the last is the correction for double counting of backlogs. The β_T^* fill rate, despite its simplicity and known issues with low fill rates, performs well for high fill rates.

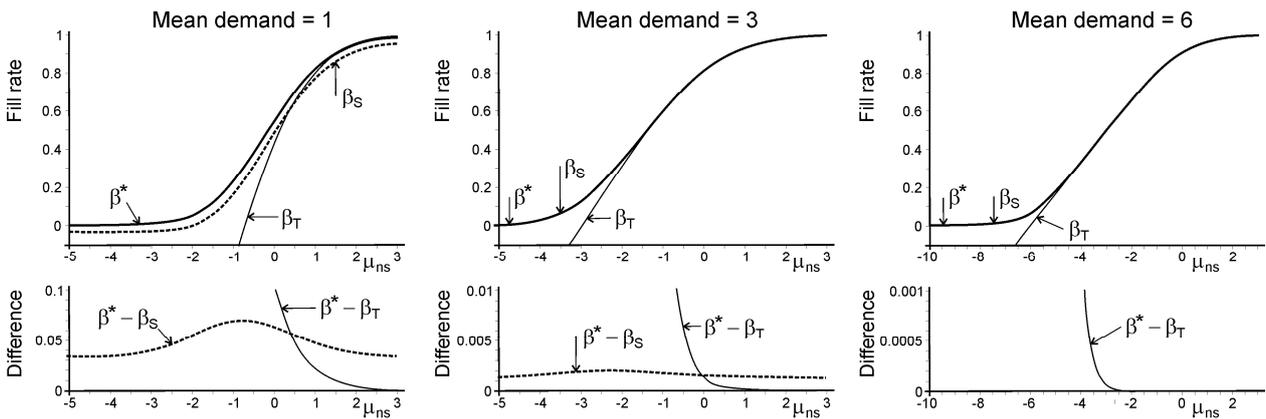


Figure 5. OUT policy fill rates with unit variance normally distributed i.i.d. demands

Figure 6 illustrates the case of ARMA(1,1) demand at $\theta = \{-0.9, 0, 0.9\}$, with the mean demand $\mu_d = \{1, 3\}$, a safety stock $\mu_{ns} = 1$ and a varying ϕ . When $\mu_d = 3$, the probability of a period having negative demand is 0.135%, and the three fill rate measures (see (7), (11) and (15)) that change with the standard deviation of the net stock are so close that they cannot be distinguished from each other in the first row of figures. There is, however, a slight difference for the two approximations $\{\beta_T^*, \beta_S^*\}$.

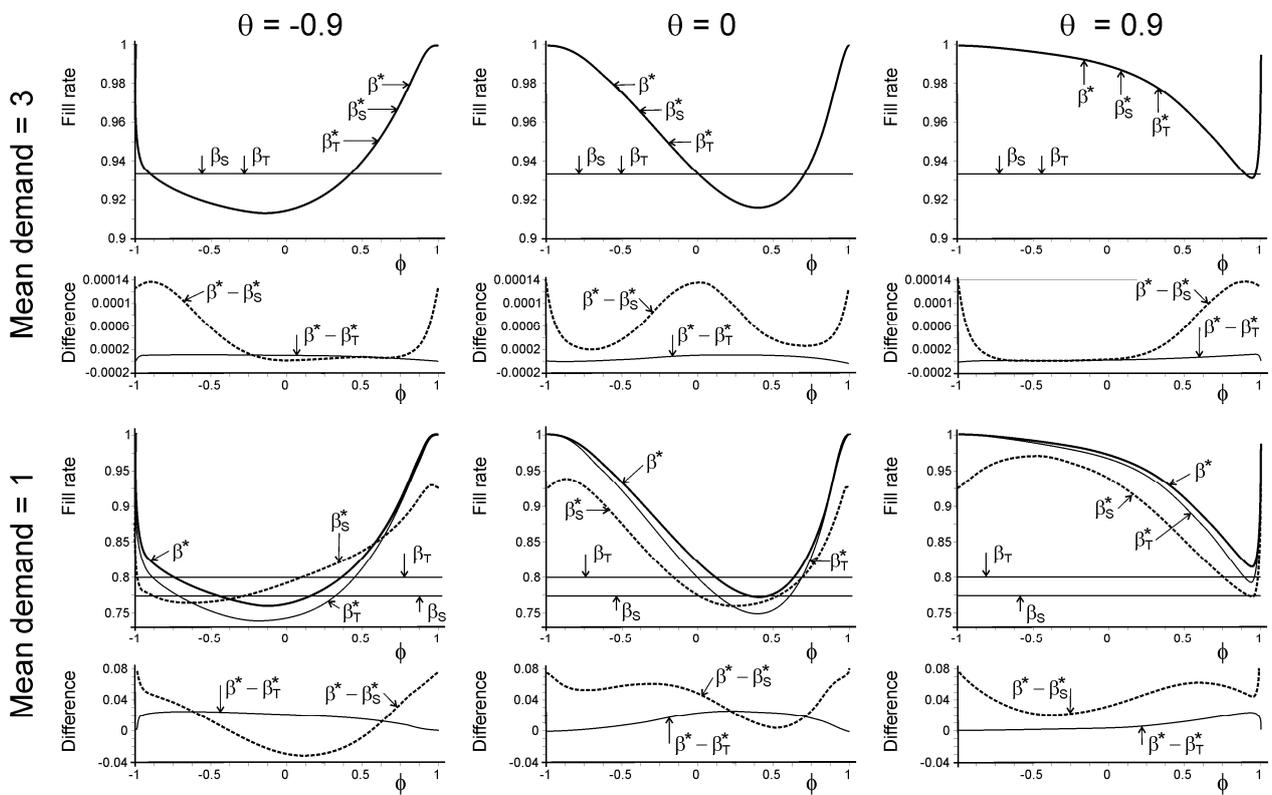


Figure 6. OUT policy fill rates with unit variance normally distributed ARMA(1,1) demands

Sobel's fill rate has the largest difference. The two fill rate measures based on incorrectly assuming i.i.d. demand, $\{\beta_T, \beta_S\}$, result in two indistinguishable horizontal lines invariant to the ARMA(1,1) parameters, as they do not account for the autocorrelation in demand. This shows the importance of

accounting for correlated demand in the fill rate measure. We can see that as $\phi \rightarrow \pm 1$, the β^* fill rate approaches 100% and that there appears to be a minimum fill rate in ϕ . Despite the probability of negative demand remaining constant and that the influence of the demand correlation has been factored into β_T^* and β_S^* , both are further influenced by the cross-correlation between ns_t and $(ns_t + d_t)$. The differences $\{\beta^* - \beta_T, \beta^* - \beta_S\}$ have not been plotted in the graphs in the second and fourth rows, as they would dominate the figure. When $\mu_d = 1$, the effect of the negative demand on the differences $\{\beta^* - \beta_T, \beta^* - \beta_S\}$ is larger than when $\mu_d = 3$. The β_S^* measure often has the largest difference with the true fill rate β^* over the whole range of ϕ in the Figure 6.

Figure 7 illustrates the effect of the demand parameters on the fill rate by showing a contour plot of β^* in the ARMA plane. Here, $\mu_d = \{1, 3\}$, $\mu_{ns} = 1$ and σ_ε^2 is scaled to ensure $\sigma_d^2 = 1$, so that we have a constant probability of negative demand. Interestingly, there are instances of 100% fill rate for $(\phi \approx -1, \theta > 0)$. Despite returns inflating inventory levels (and one would expect, increasing fill rates), we see that when the probability of negative demand is larger, the fill rate is lower. Furthermore, the i.i.d. case has some of the lowest fill rates in the whole ARMA(1,1) parameter plane. Indeed, from Figure 4 and Figure 7 we conclude that ARMA(1,1) demands with a positive Pearson correlation coefficient generally have high fill rates. Practically this implies that when the demand autocorrelation is ignored, there is a likely over-investment in inventory.

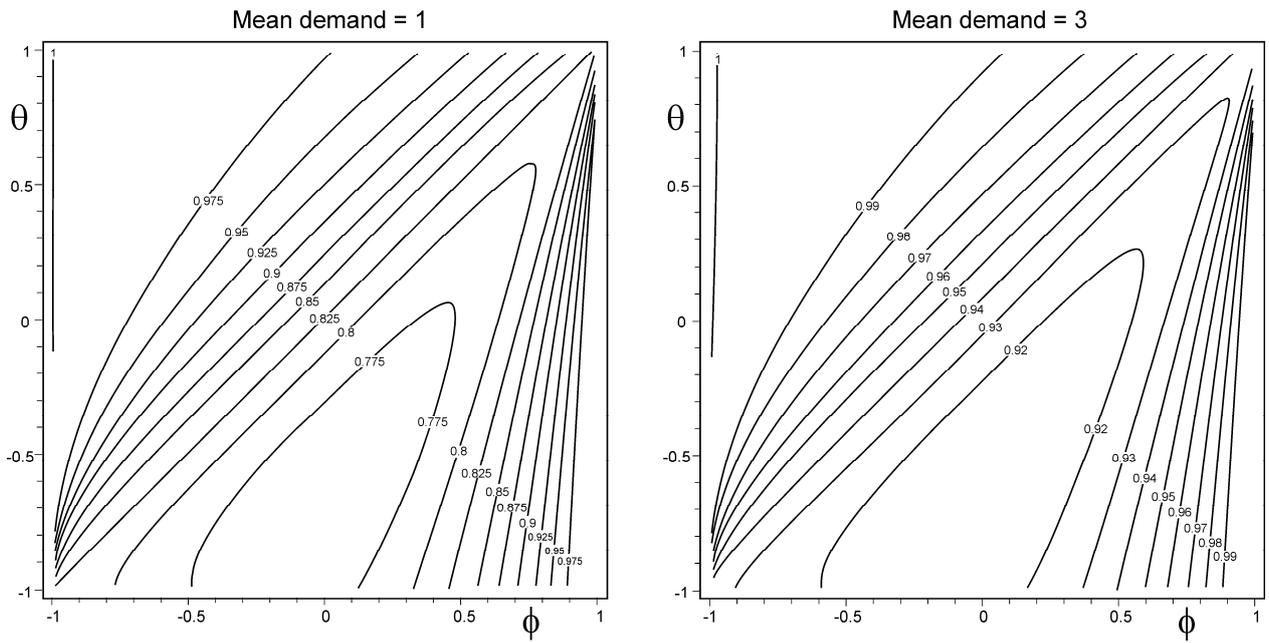


Figure 7. β^* fill rates maintained by the OUT policy with unit variance ARMA(1,1) demand with a safety stock of $\mu_{ns} = 1$.

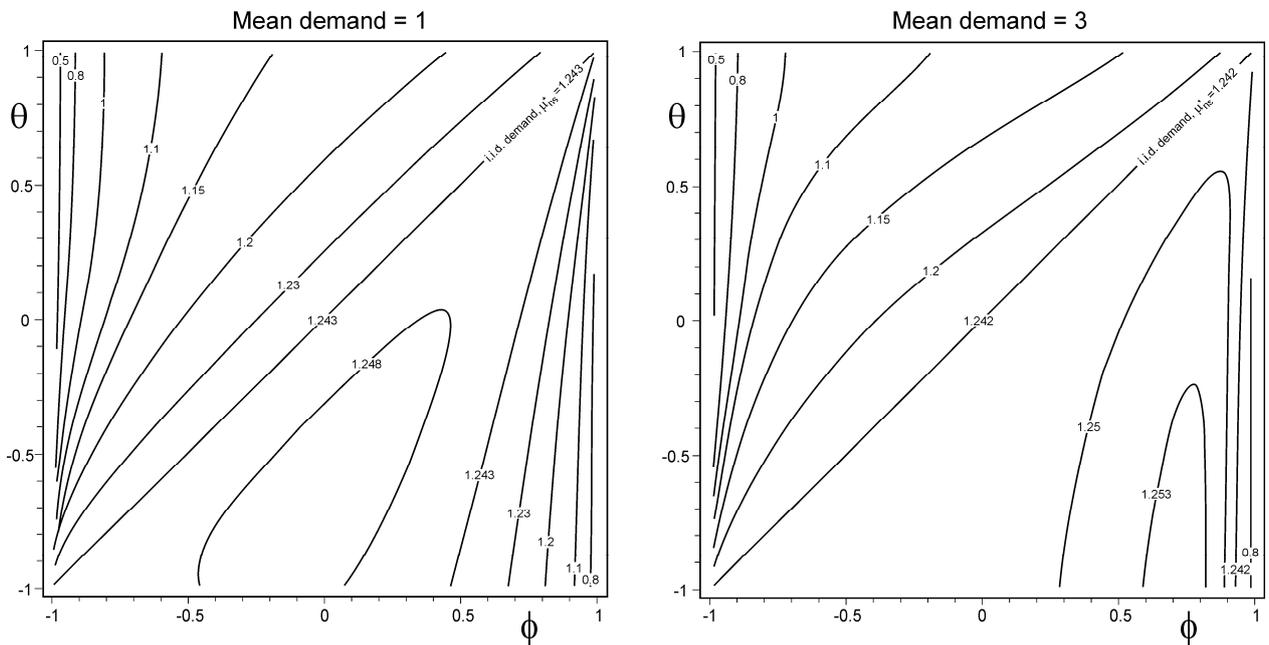


Figure 8. Safety stock requirements (μ_{ns}^*) to achieve 95% fill rate.

Figure 8 illustrates safety stock requirements to achieve 95% fill rate under ARMA(1,1) demand where σ_ε^2 is scaled to ensure $\sigma_{ns}^2 = 1$ as this normalizes a different influence on the fill rate. We see that for i.i.d. demand, when $\phi = \theta$, the required safety stock is $\mu_{ns}^* = \{1.243, 1.242\}$ for $\mu_d = \{1, 3\}$. As revealed by Figure 8, safety stock requirements under ARMA demand that have been determined using guidance derived from the i.i.d. demand formula result in excessive inventory holding if demand is negative correlated ($\theta > \phi$) or strongly positively correlated ($\phi \approx 1$).

7.2. Numerical verification via simulation

In this section we present results from a simulation study to verify our analytical results and compare the performance of the three fill rate measures. We simulated the linear OUT policy reacting to scaled ARMA(1,1) demand patterns with unit variance for 10,000 periods and replicated our study 1,000 times; the average of these 1,000 replications are in Table 1. The parameter settings were chosen as they were interesting parameter sets in their own right, or because they produced interesting results. These numerical results confirm that β^* measures the fill rate correctly under correlated, normally distributed, possible negative demand, whereas the other established measures cannot consistently achieve this. For example, we see that with a significant chance of negative demand (see Test 11), $\{\beta_T^*, \beta_S^*\} > 1$, indicating impossible fill rates above 100%. For very low fill rates close to zero (see Test 1), $\{\beta_T^*, \beta_S^*\} < 0$; another impossible result. Tests 21 and 19 investigate an ARMA(1,1) demand process close to the Integrated Moving Average demand pattern which would be optimally forecasted by exponential smoothing (Box and Jenkins, 1976). Here, we can see that a high fill rate is achieved and the $\{\beta_T^*, \beta_S^*\}$ measures perform quite well. The traditional fill rate, despite the double counting issue, is at least consistent with its theoretical and simulation results. This is not so for Sobel's fill

rate which can give inconsistent theoretical and simulation results, even for i.i.d. demand—see Tests 1, 5, 9, 11, 13, 18, 19, 22.

The β^* fill rate expression seems to be particularly useful when the probability of negative demand is high as 6 of the 8 tests with $\mu_d = 1$ (where some 15.8% of periods have negative demands) result in β_T^* and β_S^* with significant errors – see tests 1, 4, 5, 9, 13, 18, 19 and 22. If the consequences of the demand correlation on the distribution of d_t and $(ns_t + d_t)$ are properly taken into account then β_T^* and β_S^* perform well. Together these insights imply β^* should be used when negative demand is frequent when the linear OUT policy assumptions are adopted.

8. Concluding remarks

8.1. Theoretical contributions

Motivated by a real life observation of demand we challenged the assumptions of i.i.d. positive demand commonly used in the fill rate literature. We have explored the consequences of two fill rate measures from the literature under auto-correlated normally distributed demand. We have also presented a new fill rate measure based on the distribution of the minimum of two correlated normally distributed random variables. We compared our new fill rate measure to the two existing measures. When the mean demand is large in comparison to the standard deviation (i.e., negative demand is negligible), all fill rate measures work reasonably well when operating near 100% fill rate. The impact of the demand autocorrelation can be easily accounted for by simply updating the variance expression in the existing solution approaches. However when the probability of negative demand becomes larger we recommend that our exact fill rate measure is used.

Test	μ_d	μ_{ns}	ϕ	θ	Trad. fill rate, β_T^*		Sobel fill rate, β_S^*		Proposed fill rate, β^*	
					Simulation (Eq 5)	Theory (Eq 7)	Simulation (Eq 9)	Theory (Eq 11)	Simulation (Eq 13)	Theory (Eq 15)
1	1	-2	0	0	-1.049922	-1.05025	-0.02522	0	0.053738	0.053713
2	3	-2	0	0	0.3166355	0.316582	0.344353	0.344227	0.344454	0.344423
3	3	-2	0.9	0	0.331786	0.331512	0.352432	0.353047	0.352789	0.353084
4	1	0	0.7	0	0.438129	0.43808	0.488322	0.487507	0.527709	0.527607
5	1	0	0	0	0.4357787	0.43581	0.511815	0.486065	0.549456	0.54943
6	2	-0.5	0.7	0	0.576125	0.576524	0.583828	0.582773	0.585581	0.585569
7	3	-1	0.7	0	0.600554	0.600709	0.601851	0.60172	0.60175	0.601789
8	2	-0.2	0.3	-0.9	0.64733	0.647384	0.647759	0.648514	0.649216	0.649219
9	1	0.5	0	0	0.651209	0.650911	0.677545	0.647157	0.702499	0.70228
10	2	0	0.7	0	0.718812	0.719042	0.720782	0.719511	0.721672	0.721176
11	-2	3	0	0	1.0042853	1.004312	1.00107	-0.03359	0.739968	0.737554
12	2	0	-0.5	0	0.806879	0.806862	0.808587	0.806865	0.809445	0.809431
13	1	1	0	0	0.80062	0.800359	0.807976	0.775789	0.822962	0.82277
14	2	1	0.5	0.1	0.87673	0.876684	0.876535	0.875411	0.877312	0.877285
15	3	1	0.7	0.5	0.9240384	0.923995	0.923994	0.923899	0.924053	0.924
16	3	1	0	0	0.9334444	0.933453	0.933422	0.933329	0.933459	0.933464
17	3	1	0.5	-0.9	0.9383143	0.93822	0.938292	0.938221	0.938327	0.938228
18	1	2	0	0	0.949637	0.949745	0.950261	0.917067	0.953847	0.953925
19	1	1	0.99	0.7	0.973914	0.973854	0.972867	0.901089	0.977039	0.977172
20	3	1	0.9	-0.5	0.9881204	0.988115	0.987578	0.988077	0.988118	0.988117
21	3	1	0.99	0.7	0.991278	0.991284	0.991214	0.991171	0.991275	0.991287
22	1	3	0	0	0.9913637	0.991377	0.991691	0.958323	0.992036	0.992046
23	3	5	0	0	0.9999749	0.999976	0.999871	0.99985	0.999975	0.999976
24	3	1	-0.98	0.99	1	1	1	0.999901	1	1

Table 2. Numerical verification of the three fill rate measures

8.2. Managerial implications

Demand autocorrelation can have both a positive and negative influence on fill rate. When demand is negatively, or strongly positively correlated and safety stocks have been set using guidance based on i.i.d. demand, the fill rates actually achieved is higher than expected (see Figure 8), implying an overinvestment in inventory. When demand is weakly positively correlated, and safety stocks have been set using i.i.d. demand guidance, the fill rate decreases. In cases where there could be negative demand, irrespective of the autocorrelation in demand, we recommend that our new fill rate β^* is

used, either in the form of (13) for time series evaluation or in the form of (15) in target setting/analytical work. This measure avoids over investment in inventory and ensures target fill rates are met.

Practically, most enterprise software systems keep good records over time of demand and receipts, but historical records of inventory are seldom kept. This has to be calculated via the inventory balance equation. The same issue likely exists for the variable $(ns_t + d_t)$. This means that new reporting and record keeping mechanisms may be needed. The only other factor we need to accommodate with the new fill rate measure is to calculate the Pearson correlation co-efficient between the variables $\{d_t, ns_t + d_t\}$. This can be found in Excel as the CORREL function and used as an input into our Excel Add-in in Appendix A for numerical work. To find the optimal safety stock requirements in an inventory management policy, the *Goal Seek* function in the *What-If Analysis* menu of Excel can be used. The fill rate cell can be set to a specific target by changing μ_{ns+d} . Note that the safety stock is then given by $\mu_{ns} = \mu_{ns+d} - \mu_d$.

8.3. Further work

Future work could investigate the performance of the expressions in Silver and Bischak (2011) and Johnson et al. (1995). We could explore this new fill rate measure for other replenishment policies such as the proportional OUT policy (Disney and Towill, 2003), or the full-state policy (Gaalman, 2006). The consequences of non-MMSE forecasting methods may also be practically important and worthy of exploration (Li, Disney and Gaalman, 2014). Investigations on the inverse of our new fill rate measure could also be undertaken, perhaps along the lines of the analysis in Cardós and Babiloni (2011). Finally, the link between fill rates and availability (p1) could be further explored in the case of auto-correlated demand.

11. Acknowledgements

We gratefully acknowledge the helpful insights and contribution of two anonymous referees, Professors Aris A. Syntetos of Cardiff University, Stefan Minner of Munich University, and Kai A. Rosling of Linnaeus University for providing useful comments and feedback to help improve early versions of this work. We are also grateful to Professors Diane Bischak and Edward Silver from the University of Calgary for help identifying proper sources to some of the literature.

12. References

- Axsäter, S., 2000. *Inventory Control*. Springer Science+Business Media, New York, USA.
- Anon, 2012. Romberg's method. *Wikipedia*. www.wikipedia.org/wiki/Rombergs%27s_method. Accessed 26/8/12.
- Basu, A.P. and Ghosh, J.K., 1978. Identifiability of the multinormal and other distributions under competing risks model. *Journal of Multivariate Analysis*, 8, 413–429.
- Box, G.E.P. and Jenkins, G.M., 1976. *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco, USA.
- Cachon, G. and Terwiesch, C., 2006. *Matching Supply with Demand: An Introduction to Operations Management*. McGraw-Hill Irwin, Boston, USA.
- Cain, M., 1994. The moment-generating function of the minimum of bi-variate normal random variables. *The American Statistician*, 48 (2), 124–125.
- Cardós, M., and Babiloni, E., 2011. Exact and approximate calculation of the cycle service level in period review inventory policies. *International Journal of Production Economics*, 131 63–68.
- de Kok, T., 2002. Analysis of one product/one location inventory control models. Unpublished course notes at Technische Universiteit Eindhoven [circa 2002].

Disney, S.M., Gaalman, G., Hedenstierna, C.P.T. and Hosoda, T., (2015), "Fill rate in a periodic review order-up-to policy under auto-correlated normally distributed, possibly negative, demand", *International Journal of Production Economics*, 170, 501–512. DOI: 10.1016/j.ijpe.2015.07.019.

- Disney, S.M., Farasyn, I., Lambrecht, M., Towill, D.R. and Van de Velde, W., 2006. Taming the bullwhip effect whilst watching customer service in a single echelon of a supply chain. *European Journal of Operational Research*, 173, 151–172.
- Disney, S.M. and Towill, D.R., 2003. On the bullwhip and inventory variance produced by an ordering policy. *OMEGA: The International Journal of Management Science*, 31 (3), 157–167.
- Gaalman, G., 2006. Bullwhip reduction for the ARMA demand model: The proportional OUT policy, verses the full-state feedback policy. *Automatica*, 42 (8), 1283-1290.
- Gaalman, G. and Disney, S.M., 2009. On bullwhip in a family of order-up-to policies with ARMA(2,2) demand and arbitrary lead-times. *International Journal of Production Economics*, 121 (2), 454–463.
- Gaalman, G. and Disney, S.M., 2012. Bullwhip behavior in the Order-Up-To policy with ARIMA demand. 4th World Production and Operations Management Conference, Amsterdam, 1st – 5th July, 11 pages.
- Guijarro Tarradellas, E., Cardós, M. and Babiloni, E., 2012. On the exact calculation of the fill rate under discrete demand patterns. *European Journal of Operational Research*, 218 (2), 442–444.
- Hadley, G. and Whitin, T.M., 1963. *Analysis of Inventory Systems*. Prentice-Hall, Englewood Cliffs, N.J., USA.
- Hosoda, T. and Disney, S.M., 2009. Impact of market demand mis-specification on a two-level supply chain. *International Journal of Production Economics*, 121 (2), 739–751.
- Johnson, M.E., Lee, H.L., Davis, T. and Hall, D., 1995. Expressions for item fill rates in periodic inventory systems. *Naval Research Logistics*, 42 (1), 57–80.
- Kwon, I.-H., Kim, S.-S. and Baek, J.-G., 2006. A simulation based heuristic for serial inventory systems under fill-rate constraints. *International Journal of Advanced Manufacturing Technology*, 31 (3–4), 297–304.

Disney, S.M., Gaalman, G., Hedenstierna, C.P.T. and Hosoda, T., (2015), "Fill rate in a periodic review order-up-to policy under auto-correlated normally distributed, possibly negative, demand", *International Journal of Production Economics*, 170, 501–512. DOI: 10.1016/j.ijpe.2015.07.019.

Larsen, C. and Thorstenson, A., 2014. The order and volume fill rates in inventory control systems.

International Journal of Production Economics, 147, 13–19.

Li, Q., Disney, S.M. and Gaalman, G., 2014. Avoiding the bullwhip effect using damped trend forecasting and the order-up-to replenishment policy. *International Journal of Production Economics*, 149, 3–16.

Nadarajah, S. and Kotz, S., 2010. Exact distribution of the max/min of two Gaussian random variables. *IEEE Transactions on Very Large Scale Integration Systems*, 16 (2), 210–212.

Nagaraja, H.N. and Mohan, N.R., 1982. On the independence of system life distribution and cause of failure. *Scandinavian Actuarial Journal*, 1982, 188–198.

Nahmias, S., 1993. *Production and Operations Analysis*. Irwin Publishers, Boston, MA, USA.

Sapra, A., Troung, V.-A., Zhang, R.Q., 2010. How much demand should be fulfilled. *Operations Research*, 58 (3), 719–733.

Schneider, H., 1981. Effect of service-levels on order-points or order-levels in inventory models, *International Journal of Production Research*, 19 (6), 615–631.

Silver, E.A, Pyke, D.F. and Peterson, R., 1998. *Inventory Management and Production Planning and Scheduling*. John Wiley and Sons, New York.

Silver, E.A. and Bischak, D.P., 2011. The exact fill rate in a periodic review base stock system under normally distributed demand. *OMEGA*, 39, 346–349.

Sobel, M.J., 2004. Fill rates of single stage and multistage supply systems. *Manufacturing and Service Operations Management*, 6 (1), 41–52.

Teunter, R., 2009. Note on the fill rate of single-stage general periodic review inventory systems. *Operations Research Letters*, 37, 67–68.

Weisstein, E.W., 2012. Normal Distribution. From MathWorld - A Wolfram Web Resource. <http://mathworld.wolfram.com/NormalDistribution.html>. Accessed 7th January 2012.

Disney, S.M., Gaalman, G., Hedenstierna, C.P.T. and Hosoda, T., (2015), "Fill rate in a periodic review order-up-to policy under auto-correlated normally distributed, possibly negative, demand", *International Journal of Production Economics*, 170, 501–512. DOI: 10.1016/j.ijpe.2015.07.019.

Zhang, J. and Zhang, J., 2007. Fill rate of single stage general periodic review inventory systems. *Operations Research Letters*, 35, 503–509.

Appendix A. An Excel Add-in for determining the fill rate

An Excel Add-in that uses the moments of the minimum of bivariate normal random variables (Cain, 1994) and Romberg's method (Anon, 2012) to numerically estimate the definite integral in (15) between

$$a \leq f(y) \leq b; \left\{ a = \left(\mu_{f(y)} - 6\sigma_{f(y)} \right)^+, b = \left(\mu_{f(y)} + 6\sigma_{f(y)} \right)^+ \right\} \quad (A1)$$

is given below in Table A1. Romberg's method was chosen due to its stability and accuracy. In big O notation, the error for estimate $R(n, m)$ is $O\left(\left(\frac{b-a}{2^n}\right)^{2m+2}\right)$ where $\{a, b\}$ is given in (A1) and $\{n, m\} = 10$ in the VBA code in Table A1. When the code below is cut and pasted into an Excel function module, the expression '=Fillrate($\mu_{ns+d}, \sigma_{ns+d}, \mu_d, \sigma_d, \rho$)' is available in Excel.

```

VBA code required to determine the fill rate

Option Explicit

Function fy(mu1 As Double, sigma1 As Double, mu2 As Double, sigma2 As Double, rho As Double, y As Double)
Dim f1part, f2part As Double

f1part = ((-y - mu2) / sigma2) + rho * ((y - mu1) / sigma1) / ((1 - rho ^ 2) ^ 0.5)
f2part = ((-y - mu1) / sigma1) + rho * ((y - mu2) / sigma2) / ((1 - rho ^ 2) ^ 0.5)
fy = (1 / sigma1) * WorksheetFunction.NormDist((y - mu1) / sigma1, 0, 1, False) * WorksheetFunction.NormDist(f1part, 0, 1, True) +
(1 / sigma2) * WorksheetFunction.NormDist((y - mu2) / sigma2, 0, 1, False) * WorksheetFunction.NormDist(f2part, 0, 1, True)

End Function

Function Fillrate(mu1 As Double, sigma1 As Double, mu2 As Double, sigma2 As Double, rho As Double)
Dim R(10, 10), h, f, a, b, m1, m2, theta, y, var, s, d As Double
Dim n, m, k As Integer

theta = (sigma2 ^ 2 - 2 * rho * sigma1 * sigma2 + sigma1 ^ 2) ^ 0.5
m1 = mu1 * WorksheetFunction.NormDist((mu2 - mu1) / theta, 0, 1, True) + mu2 * WorksheetFunction.NormDist((mu1 - mu2) / theta,
0, 1, True) - theta * WorksheetFunction.NormDist((mu2 - mu1) / theta, 0, 1, False)
m2 = (sigma1 ^ 2 + mu1 ^ 2) * WorksheetFunction.NormDist((mu2 - mu1) / theta, 0, 1, True) + (sigma2 ^ 2 + mu2 ^ 2) *
WorksheetFunction.NormDist((mu1 - mu2) / theta, 0, 1, True) - (mu1 + mu2) * theta * WorksheetFunction.NormDist((mu2 - mu1) /
theta, 0, 1, False)
var = m2 - m1 ^ 2

If m1 - 6 * var ^ 0.5 < 0 Then
    a = 0
Else
    a = m1 - 6 * var ^ 0.5
End If

If m1 + 6 * var ^ 0.5 < 0 Then
    b = 0
Else
    b = m1 + 6 * var ^ 0.5
End If

For n = 0 To 10
    h = (b - a) / 2 ^ n
    If n = 0 Then
        R(0, 0) = 0.5 * (b - a) * (fy(mu1, sigma1, mu2, sigma2, rho, a) * a + fy(mu1, sigma1, mu2, sigma2, rho, b) * b)
    Else
        For m = 0 To n
            If m = 0 Then
                s = 0
                For k = 1 To 2 ^ (n - 1)
                    s = s + fy(mu1, sigma1, mu2, sigma2, rho, a + (2 * k - 1) * h) * (a + (2 * k - 1) * h)
                Next k
                R(n, m) = 0.5 * R(n - 1, 0) + h * s
            Else
                R(n, m) = R(n, m - 1) + (1 / (4 ^ m - 1)) * (R(n, m - 1) - R(n - 1, m - 1))
            End If
        Next m
    End If
Next n

d = 0.5 * (mu2 + Exp(-(mu2 ^ 2 / (2 * sigma2 ^ 2))) * sigma2 * 0.797884560802865 + mu2 * (2 *
Application.WorksheetFunction.NormDist((mu2 / sigma2), 0, 1, True) - 1))

If Abs(R(9, 9) - R(10, 10)) > 0.000000001 Then
    Fillrate = "The integral has not converged to within 0.000000001"
Else
    Fillrate = R(10, 10) / d
End If

End Function

```

Table A1. VBA code for the fill rate with correlated normally distributed demands