STOCHASTIC MODELS FOR INTERMITTENT DEMANDS FORECASTING AND STOCK CONTROL

W. Sandmann, O. Bober
University of Bamberg, Germany

Corresponding author: W. Sandmann
University of Bamberg, Dep. Information Systems and Applied Computer Science
96045 Bamberg, Feldkirchenstraße 21, Germany, werner.sandmann@uni-bamberg.de

Abstract. Demand forecasting with regard to stock control is a central issue of inventory management. Serious difficulties arise for intermittent demands, that is, if there are slow-moving items demanded only sporadically. Prevalent methods then usually perform poorly as they do not properly take the stochastic nature of intermittent demand patterns into account. They often rely on theoretically unfounded heuristic assumptions and apply inappropriate deterministic smoothing techniques. We overcome these weaknesses by means of systematically built and validated stochastic models that properly fit to real (industrial) data. Initially, no assumptions are made but statistical methods are invoked for model fitting. Reasonable model classes are found by summary statistics and correlation analysis. Specific models are obtained by parameter estimation and validated by goodness-of-fit tests. Finally, based on the stochastic models, stock control strategies are proposed to facilitate service levels guarantees in terms of probability bounds for being out of stock.

1 Introduction

Any organization or company that offers, sells and delivers items to others has to take care about proper inventory management. Success and efficiency substantially depend on the ability to provide and deliver demanded items within reasonable time. Stock control is crucial and the inventory policy manages how many units of an item must be in stock subject to certain constraints. Inventory capacities are limited and inventory costs should be as low as possible but at the same time a desired level of item availability should be assured. Typically, in order to be well prepared, stock control relies on forecasting future demands by means of time series analysis based on past demand patterns. Comprehensive treatments of time series analysis and forecasting can be found in, e.g., [1, 2, 4, 5, 8]. For the broader scope of inventory management we refer the reader to [14, 18, 19].

In practice, the most common forecasting technique is simple exponential smoothing (SES), that is forecasts are made by means of a weighted sum of past observations in that based on given time series data $x_1,\ldots,x_t$ a forecast $\hat{x}_{t+1}$ for the next data point x_{t+1} is computed recursively by $\hat{x}_1 = x_1$ and $\hat{x}_{t+1} = \alpha x_t + (1-\alpha)\hat{x}_t$ for $t > 0$ where $\alpha \in (0,1)$ is a smoothing constant that needs to be chosen appropriately. Unfortunately, SES does not provide satisfactory forecasts for intermittent demands, i.e. in the case of so-called slow-moving items or low-demand items that are only demanded sporadically. Instead, Croston's method [6] is most widely applied to intermittent demands.

Croston separates the time intervals between successive demands (interdemand times) and the number of units that are demanded (demand sizes). He argues that the time periods (measured in days, weeks, or months) between successive demands as well as the demand sizes are independent and identically distributed (iid) random variables, which means that intermittent demands essentially appear at random without identifiable trends, seasonality, or the like. He heuristically assumes the geometric distribution for the interdemand times and the normal distribution for the demand sizes. If a demand occurs, separate forecasts for both the interdemand time and the demand size are updated according to SES using the same smoothing constant for both forecasts and the current demand per period forecast is obtained by the ratio of these two forecasts. Obviously, the critical issue of choosing an appropriate smoothing constant remains open.

A couple of drawbacks such as biased forecasts or potential violations of the independence assumptions have been reported in the literature and many corrections and modifications, repectively, have been proposed, e.g. [3, 11, 12, 15, 16, 17]. However, though specifically targeted to intermittent demands and often more accurate than SES, in some cases Croston's method and its various modifications do not provide proper forecasts and not even outperform pure SES. After all, there is still no satisfactory approach to deal with stock control for slow-moving items based on forecasting intermittent demands. We argue that the problems essentially stem from the application of deterministic smoothing techniques to random patterns. Stochastic models appear to be more appropriate and promising for tackling the peculiarities of intermittent demands forecasting and stock control.

2 Stochastic Modeling Approach

In order to improve forecasting and stock control for slow-moving items we first have to figure out the weaknesses of existing methods and the requirements for overcoming these weaknesses. The mixture of assuming stochastic behavior and applying SES as a forecasting technique is inappropriate. The modifications of Croston's original version, adopting the independence assumptions as well as the assumption of geometrically distributed interdemand times, are mainly concerned with nonnormal (but still continuous) distributed demand sizes and modified forecasts. In particular, they still produce deterministic point forecasts though the demand pattern is essentially random. Mathematically, they work with deterministic realizations rather than with stochastic processes which are supposed to be the data generating mechanism. It has been recently shown by Shenstone & Hyndman [13] that the application of these deterministic forecasting techniques cannot be consistent with stochastic models, because any underlying stochastic model must be non-stationary and defined on a continuous sample space with negative values, which both does not match to the real properties of intermittent demand patterns.

We believe that starting with a forecasting technique and building an according underlying model is exactly the reversed order of what is required. In particular, the major problem lies in the inappropriate forecasting technique rather than in approaching intermittent demands via stochastic models. Additionally, we point out that in the previously cited literature specific probability distributions are heuristically assumed and - if at all - checked against artificial simulated data. It is often just a matter of luck whether or not the assumptions well fit to real intermittent demand data. Consequently, we argue that one should first build an adequate stochastic model based on real data, then validate its goodness of fit and finally derive forecasts and stock control strategies to meet certain requirements such as service level guarantees.

2.1 Stochastic Time Series Models

A time series is an ordered sequence $x_1,\dots,x_n$, interpreted as a realization of a stochastic process $(X_t)_{t\in T}$. As we are concerned with discrete time points (months), we shall assume that the index set T is a subset of the nonnegative integers. Note that, though often neglected in the literature, there is an important difference between a time series and its "generating" stochastic process. Other than a stochastic process, which is an ensemble of time series, a single time series is just one sequence of deterministic data.

Time series properties are characterized by corresponding properties of stochastic processes. We briefly present those that are most important with regard to stochastic time series models.

Definition 1 (Moment Functions of Stochastic Processes) *For a stochastic process $(X_t)_{t\in T}$ its*

- *mean function is defined by* $\mu(t) := E[X_t], t \in T$,
- *variance function is defined by* $\sigma^2(t) := Var[X_t], t \in T$,
- *autocovariance function is defined by* $\gamma(s,t) := Cov(X_s,X_t) = E[(X_s-\mu(s))(X_t-\mu(t))], \quad s,t \in T$.

Note that $\gamma(t,t) = \sigma^2(t)$ for all $t \in T$.

Definition 2 (Strict Stationarity) *A stochastic process $(X_t)_{t\in T}$ is called strictly stationary, if its finite dimensional distributions are time invariant. That is, for all $n,t_1,\dots,t_n,s$ the random vectors $(X_{t_1},X_{t_2},\dots,X_{t_n})$ and $(X_{t_1+s},X_{t_2+s},\dots,X_{t_n+s})$ have the same distribution.*

Definition 3 (Weak Stationarity) *A stochastic process $(X_t)_{t\in T}$ is called weakly stationary, if it has constant mean and variance function and its autocovariance function does not dependent on specific time points but only on the time difference, the so-called lag τ. That is, for all $s,t \in T, \tau := t-s$:*

$$\mu(t) = \mu < \infty, \quad \sigma^2(t) = \sigma^2 < \infty, \quad \gamma(s,t) = \gamma(s+\tau,t+\tau).$$

Then the autocovariance function for lag τ is defined as $\gamma(\tau) = \gamma(t-s) := \gamma(s,t)$.

It follows for the autocovariance function of a weakly stationary stochastic process, that for all τ :

$$\gamma(0) = \sigma^2, \quad \gamma(\tau) = \gamma(-\tau), \quad |\gamma(\tau)| \le \gamma(0).$$

Definition 4 (Autocorrelation Function) *The autocorrelation function of a weakly stationary stochastic process $(X_t)_{t\in T}$ is defined by*

$$\rho(\tau) := \frac{\gamma(\tau)}{\gamma(0)} = \frac{\gamma(\tau)}{\sigma^2}.$$

It follows for the autocorrelation function of a weakly stationary stochastic process, that for all τ :

$$\rho(0) = 1, \quad \rho(\tau) = \rho(-\tau), \quad |\rho(\tau)| \le 1.$$

Now, being equipped with the most important properties of stochastic processes, we introduce the most important stochastic processes with regard to time series analysis and in particular with regard to our modeling approach.

Definition 5 (White Noise) *A white noise is a sequence $(Z_t)_{t\in T}$ of independent and identically distributed (iid) random variables. If all these random variables are normally distributed with expectation $\mu = 0$ and variance $\sigma^2 < \infty$, then $(Z_t)_{t\in T}$ is called a Gaussian white noise.*

Obviously, a white noise is a strictly stationary stochastic process. In time series analysis, white noise is used for constructing more complex stochastic processes. In the following, let $(Z_t)_{t\in T}$ denote a white noise with expectation $\mu = 0$ and variance $\sigma^2 < \infty$.

Definition 6 (Autoregressive Process) *An autoregressive process of order p, denoted by AR[p], is a stochastic process $(X_t)_{t\in T}$ defined by*

$$X_t = \alpha_1 X_{t-1} + \ldots + \alpha_p X_{t-p} + Z_t, \quad t \in T,$$

where $\alpha_1, \ldots, \alpha_p$ are constant coefficients.

Definition 7 (Moving Average Process) *A moving average process of order q, denoted by MA[q], is a stochastic process $(X_t)_{t\in T}$ defined by*

$$X_t = \beta_0 Z_t + \beta_1 Z_{t-1} + \beta_2 Z_{t-2} + \ldots + \beta_q Z_{t-q}, \quad t \in T,$$

where $\beta_0, \ldots, \beta_q$ are constant coefficients. The random variables $Z_t, t \in T$ constituting the underlying white noise are usually normalized such that $\beta_0 = 1$.

Definition 8 (ARMA-Process) *An autoregressive moving average process of order (p,q), denoted by ARMA[p,q], is a stochastic process $(X_t)_{t\in T}$ defined by*

$$X_t = \alpha_1 X_{t-1} + \ldots + \alpha_p X_{t-p} + Z_t + \beta_1 Z_{t-1} + \ldots + \beta_q Z_{t-q}, \quad t \in T,$$

where $\alpha_1, \ldots, \alpha_p$ and $\beta_1, \ldots, \beta_q$ are constant coefficients.

Hence, ARMA processes are composed of AR processes and MA processes and both AR processes and MA processes are specific ARMA processes. An AR$[p]$ process is an ARMA$[p,0]$ process and an MA$[q]$ process is an ARMA$[0,q]$ process. Note that a white noise also fits this framework in that it is an ARMA$[0,0]$ process.

3 Building and Validating Stochastic Models

We have developed a systematic procedure for model fitting to real data by courtesy of *Siemens AG – Healthcare Sector Customer Services Material Logistics* (aka *Siemens Medical Solutions*), Erlangen, Germany. Initially, no independence assumption are made but after computing comprehensive statistics, the independence of data is checked. Dependent on the outcome of suitable tests, either the time series corresponding to interdemand times and demand sizes are fitted to autoregressive moving average (ARMA) processes or fitted to adequate probability distributions.

The essential steps of our modeling procedure and their order in an automated algorithmic application starting with the raw data are outlined below. We emphasize that this procedure is flexible and well accessible to practitioners. Many steps are supported by statistical software packages, which is important for being viable as a part of real inventory management within industrial companies.

1. Study summary statistics and the correlation structure of interdemand times and demand sizes
2. Test the independence hypothesis: Ljung-Box test
3. If independent, fit data to appropriate probability distribution
   (a) Select potentially appropriate distribution families based on summary statistics
   (b) For each candidate distribution, obtain parameters by maximum likelihood estimation
   (c) Validate goodness-of-fit by visualization: graphical plots
   (d) Validate goodness-of-fit by statistical tests: χ^2, Anderson-Darling, Kolmogorov-Smirnov
4. If not independent, fit data to ARMA model
   (a) Select potentially appropriate class of ARMA based on (partial) autocorrelation functions
   (b) For each candidate class, obtain parameters by
      (b1) least squares estimation for the AR part
      (b2) numerical iteration for the MA part
   (c) Validate goodness-of-fit by residual analysis

This procedure has been applied to the demand patterns of 54 different slow-moving items, each recorded from September 1994 to May 2008. In addition to model fitting for each of these items, one goal was to identify similarities in order to build an aggregated model that integrates as much slow-moving items as possible (inventory of Siemens Medical Solutions takes care bout altogether about 8500 slow-moving items). In the following we describe some more details of the steps and outline the main findings that we obtained in this manner.

Note that we could have formulated our model fitting procedure without explicitly distinguishing between independent and dependent data by just fitting to an ARMA process and keeping in mind that an ARMA$[0,0]$ process is a white noise which means that the data is independent. However, this would be overly generalized. We make the distinction with regard to the specific fitting methodologies, which are much easier for independent data.

3.1 Summary Statistics and Correlation Structure

Summary statistics are usually not common in time series analysis when trends, seasonality, dependencies or correlations are present. Nevertheless, they should be computed as a first step in any statistical data analysis because they almost always give useful insights, in particular in the case of intermittent demands where completely random patterns in the sense of iid data or white noise are very likely. In addition, the correlation structure is of major important in time series analysis. Therefore, we consider a variety of statistical measures that give us a first quantitative impression of the data and its correlation structure. More precisely, we compute the following empirical measures from the time series data $x_1,\dots,x_n$.

- *Empirical Mean*
$$\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i,$$
- *Empirical Variance*
$$s^2 = \frac{1}{n-1}\sum_{i=1}^{n}(x_i-\bar{x})^2,$$
- *Empirical Standard Deviation*
$$s = \sqrt{s^2},$$
- *Empirical Coefficient of Variation*
$$c = \frac{s}{\bar{x}},$$
- *Empirical Standard Error*
$$\sigma_n = \frac{s}{\sqrt{n}},$$
- *Empirical Skewness*
$$\frac{\frac{1}{n}\sum_{i=1}^{n}(x_i-\bar{x})^3}{\sqrt{\left(\frac{1}{n}\sum_{i=1}^{n}(x_i-\bar{x})^2\right)^3}},$$
- *Empirical Kurtosis*
$$\frac{\frac{1}{n}\sum_{i=1}^{n}(x_i-\bar{x})^4}{\left(\frac{1}{n}\sum_{i=1}^{n}(x_i-\bar{x})^2\right)^2}-3,$$
- *Empirical Covariance*
$$s_{xy} = \frac{1}{n-1}\sum_{i=1}^{n}(x_i-\bar{x})(y_i-\bar{y}),$$
- *Empirical Coefficient of Correlation*
$$r_{xy} = \frac{s_{xy}}{s_x s_y}.$$

Note that the latter two measures can be computed with regard to different time series as well as within a single time series. One can shift the time series by a lag τ and interpret the resulting pairs $(x_i,y_i) = (x_i,x_{i+\tau})$ in the same way as data points from different time series. Concerned with a single time series one also speaks of autocorrelation as for stochastic processes. In particular, when considering such autocorrelations for different lags, one can get useful insights on the strength of potentially present dependencies, which then yields guidelines for the choice of appropriate ARMA processes.

3.2 Independence Test

After computing comprehensive summary statistics and studying the correlation structure of the given data, our next step is to test the hypothesis of independence. More specifically, we want to examine whether it is appropriate to assume that successive interdemand times and demand sizes are independent. In the terminology of stochastic time series models, this independence hypothesis corresponds to the hypothesis that the data generating stochastic processes are white noise.

Hence, we are concerned with the independence of successive data from one time series (not with the independence of two or more data sets). In order to test it, the Ljung-Box test [9] can be applied.

With the Ljung-Box test all autocorrelation coefficients are considered simultaneously. The independence hypothesis H_0 and its alternative H_1 are defined as

$$H_0: \quad \rho(s) = 0, \text{ for all } s \in \mathbb{N}, \ \ s > 0,$$

$$H_1: \quad \{s \mid \rho(s) \neq 0, s \in \mathbb{N}, s > 0\} \neq \emptyset$$

and the test statistic is computed by

$$T(\hat{\rho}) = n(n+1) \sum_{j=1}^{h} \frac{\hat{\rho}_j^2}{n-j}.$$

where n is the sample size, $\hat{\rho}_j$ is the empirical coefficient of autocorrelation on lag j and h is the number of lags. Then for a given significance level α the critical section is defined by

$$T(\hat{\rho}) > \chi^2_{1-\alpha,h}$$

where $\chi^2_{1-\alpha,h}$ is the $1-\alpha$-quantile of the chi square distribution with h degrees of freedom.

The Ljung-Box test has been specified as an S-Plus program and applied to the intermittent demand patterns, that is to the interdemand times and the demand sizes of all slow-moving items. One of our main findings is that for almost all intermittent demand patterns the independence assumption is valid. More specifically, according to the Ljung-Box test on a statistical significance level of 0.05, the independence hypothesis was only rejected for two out of the 54 items considered, and the run tests that we additionally performed did not indicate any dependence.

3.3 Fitting Independent Data

If the data is assumed to be independent, our modeling procedure proceeds by fitting the data to a probability distribution. In order to find appropriate probability distributions that well represent the measured data, we performed a parametric fit. Some candidate distributions were selected and for each of these distributions a maximum likelihood estimation (MLE) was carried out to find those parameters that fit best to the data. The goodness of fit was tested by the χ^2 test, the Kolmogorov-Smirnov (KS) test, and the Anderson-Darling (AD) test.

In a nutshell, given a parametric distribution and measurement data, MLE determines the parameter values of the distribution that are most likely with respect to the data. Goodness-of-fit tests check the assumption that the data is according to a specific distribution and compute test statistics that indicate how well justified this assumption actually is. Though a statistical test is mainly designed for rejecting a hypothesis that is not well justified, it can also provide some sort of ranking, which is important in cases where several distributions are not rejected and may be appropriate. Roughly speaking, the goodness of fit is typically evaluated by the rule that the smaller the test statistic, the better the fit. Also p-values can be taken into account and graphical methods such as a variety of plots visualize how close the fitted distribution is to the data. For the details of the statistical methods we refer the reader to, e.g., [7, 10].

As for the specific distributions that have been fitted to the data, our focus was on discrete probability distributions since both the interdemand times (measured in months for our data) and the demand sizes are integer numbers. We found that for all 54 items the geometric distribution fits best for the interdemand times and that the discrete logarithmic distribution fits best to the demand sizes in the majority of cases and at least fits well in all cases. Hence, Croston's interdemand time distribution has turned out to be appropriate and is now validated on the basis of real data. Only his assumption of normally distributed demand sizes has been substituted by validated logarithmic distribution. This supports our belief that the weaknesses are mainly due to inappropriate forecasting techniques.

3.4 Fitting Dependent Data

If the data is assumed to be dependent, our modeling procedure proceeds by fitting the data to an ARMA process. In order to find appropriate ARMA processes that well represent the measured data, we first selected a potentially appropriate class of ARMA models based on partial autocorrelation functions, estimated the corresponding parameters and validated the fit by residual analysis. Roughly speaking, the fit is made iteratively such that processes of different orders are successively fitted and the sum of squares of differences between the fitted process and the time series data is computed. Finally, the order providing the least square sum is chosen.

Informally, the partial correlation between two variables is the correlation that remains if the possible impact of all other random variables has been eliminated and it is much easier to determine the order of an ARMA process via the partial autocorrelation function than via the original autocorrelation function. For our intermittent demand data, inspection of the partial autocorrelation functions suggested that pure AR processes are most appropriate. We have specified an according procedure in S-Plus and specifically used the Akaike information criterion (AIC) where

$$AIC(k) = n\log(\hat{\sigma}^2_{\varepsilon,k}) + 2k$$

is considered and the k for which $AIC(k)$ is minimal yields the estimated order of the AR process. For the details we refer the reader to, e.g., [1, 2, 4, 5, 8].

Even for the two potentially critical items that did not pass the Ljung-Box test for independence, the best fits to ARMA processes resulted in neglecting the MA part and low orders of the AR part, that is, these data were best fitted to purely autoregressive processes of order 2 and 4, respectively, which indicates a very weak dependence. Taking the data as independent and fitting to probability distributions resulted in very accurate fits. Hence, it seems reasonable to assume independence even for these two items. Note that statistical tests give only statements with certain statistical significance, neither proofs nor disproofs of hypotheses.

4 Forecasting and Stock Control

Once we have built stochastic models and validated their appropriateness, it is clear that deterministic point forecasts are not very meaningful but stock control strategies are possible which do not rely on simple point forecasts. We obtain service level guarantees in terms of probability bounds on stock out or item availability, respectively. More specifically, we can compute the probability of a demand size being greater than some given threshold, which is closely related to quantiles and tail probabilities of the fitted demand size distribution. Furthermore, to be useful for inventory management in practice where typically not every month each item stock is checked we also need to consider a larger time horizon as the planning period.

Therefore, we compute the probability of more than a given number of demands within a certain time period, essentially via tail probabilities of sums of random variables. After all, we end up with stock control strategies guaranteeing that, given a desired service level in terms of probabilities and the constraint of minimized inventory costs, for every inventory period sufficiently many units of all items are in stock. Finally, thanks to stochastic similarities, items can be aggregated yielding an integrated inventory control system. Hence, altogether we have a mathematically well-founded model fitting procedure for practicable stock control of slow-moving items that seems to be promising and and overcomes some of the weaknesses of currently practiced methods.

4.1 Stock Control Exemplification

We demonstrate a possible application in practice by a simple example where we assume that the interdemand times and demand sizes are independent and have been properly fitted to probability distributions. We further assume that for an example item a prescribed service levels should be achieved within a time horizon of n months. The service level is considered to be the item availability, that is the probability η that demands for this item can be immediately served by the units in stock. The question arises how many units of this item must be in stock for a given desired availability.

Though the demand size has no theoretical upper bound, the demand size distribution allows statements on the probability of certain demand sizes. We argue that for practicable stock control only demand sizes with some minimum probability can be reasonably taken into account. In other words, we consider a lower probability bound for the demand sizes such that all demand sizes with a smaller probability are neglected. Then, whenever a demand occurs, the units to be hold in stock should equal the η-quantile Q_η of the demand size distribution times the number of months in which the item is actually demanded. Hence, by considering a time horizon of more than one month for the planning period, we are even able to serve extreme demands with unlikely large demand sizes, provided that they do not occur multiple times within the time horizon of the planning period.

The remaining question is how often (in how many months) within the next n months we should be prepared for demands. We do not simply take the expectation of the number of demands within n months, because this does not account for the specific probability distribution. Similarly as for the demand sizes, we consider an upper bound such that only with a small probability demands occur in more months than suggested by the bound. This probability is determined by and can be obtained from the interdemand size distribution.

For numerical illustration, we consider an example item that has according to our fitting procedure geometrically distributed interdemand times with parameter $p = 0.21168$. Hence, the probability that the item is demanded in the next month equals p, which is relatively large. This means we should be prepared for a demand in the next month. Moreover, the number of demands is binomially distributed, that is the probability of exactly j demands within the next n months is given by

$$P_{j,n} := \binom{n}{j} p^j (1-p)^{n-j}.$$

Now, assume that our time horizon of the planning period is six months. Then the probability that the item is demanded in each of these next six months is

$$P_{6,6} = \binom{6}{6} p^6 (1-p)^0 = p^6 = 0.00009,$$

which is extremely small. Such a demand pattern will happen on average every thousand years such that it is reasonable to neglect it. Clearly, other patterns are more likely. For instance, the probability of exactly one demand within the next six months is $P_{1,6} = 0,38667$. Hence, we have to choose a bound ζ such that any probability below this bound will be considered too small to realistically assume the corresponding demand pattern.

Let $\zeta = 0,01$. From the remaining demand patterns with larger probability, choose those with the largest number of months with demand. Multiplying this number with the η-quantile Q_η of the demand size distribution then gives the number m of units required in stock. For our example item the demand sizes are logarithmically distributed with parameter $\vartheta = 0,54295$. With the choice of $\eta = 0.95$, the 95% quantile of the demand sizes computes as

$$Q_{0.95} = min\left\{ x : \frac{-1}{\ln(1-\vartheta)} \sum_{i=1}^{x} \frac{\vartheta^i}{i} \geq 0,95 \right\} = 4.$$

Now, for demonstration purposes, successively compare the probabilities $P_{j,n}$ with $\zeta = 0,01$:

$$P_{1,6} = 0.38667 > \zeta,$$

$$P_{2,6} = 0.25957 > \zeta,$$

$$P_{3,6} = 0.09293 > \zeta,$$

$$P_{4,6} = 0.01872 > \zeta,$$

$$P_{5,6} = 0.00201 < \zeta.$$

Hence, we take four demands within the next six months as sufficiently likely to be prepared for it and the number of units required in stock is therefore

$$m = Q_{0.95} \cdot max\{j : P_{j,n} > \zeta\} = 4 \cdot 4 = 16.$$

The outlined approach is astonishingly simple and provides a useful way of assuring certain service levels. Nevertheless, a couple of issues require further investigation such as a reasonable planning period (time horizon) over which stock control should be considered. Besides, an additional option to react within the planning period when the number of units in stock falls below a critical level would be surely useful. However, with regard to these points, there are usually practical constraints and not everything what is theoretically desirable is practically possible.

5 Conclusion

We have presented a stochastic modeling approach for the demand patterns of slow moving items with regard to intermittent demands forecasting and stock control. The key problem in intermittent demands forecasting and stock control is the demand pattern of slow moving items which renders traditional deterministic exponential smoothing techniques inappropriate. Stochastic models are required to capture the intermittent demand pattern. It turns out that stochastic time series models are well suited. In many cases the intermittent demand patterns even appear to be purely random in the sense that interdemand times and demand sizes are iid random variables corresponding to a white noise process.

A procedure has been suggested for fitting real data to suitable stochastic models based on which forecasting and stock control become well-founded and automated. While this modeling procedure already seems to be mature enough to address items separately without any substantial improvements necessary, further research should deal with aggregated models for multiple items. With regard to stock control for single items, one potential application of the stochastic model has been demonstrated and similar approaches with an aggregated model are highly desirable. Hence, stochastic models for intermittent demands forecasting and stock control appear to be promising and have already proven useful but a lot of future work is still required and already ongoing.

6 References

[1] G. E. P. Box and G. M. Jenkins. *Time Series Analysis: Forecasting and Control*. Holden Day, 1970.

[2] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons, 4th edition, 2008.

[3] J. E. Boylan and A. A. Syntetos. The accuracy of a modified Croston procedure. *International Journal of Production Economics*, 507:511–517, 2007.

[4] P. J. Brockwell and R. A. Davis. *Introduction to Time Series and Forecasting*. Springer, 2nd edition, 2002.
[5] C. Chatfield. *The Analysis of Time Series: An Introduction*. Chapman & Hall, 6th edition, 2004.
[6] J. D. Croston. Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, 23(3):289–303, 1972.
[7] R. V. Hogg, A. T. Craig, and J. W. McKean. *Introduction to Mathematical Statistics*. Prentice Hall, 6th edition, 2004.
[8] G. Kirchgässner and J. Wolters. *Introduction to Modern Time Series Analysis*. Springer, 2007.
[9] G. M. Ljung and G. E. P. Box. On a measure of lack of fit in time series models. *Biometrika*, 2(65):297–303, 1978.
[10] W. C. Navidi. *Statistics for Engineers and Scientists*. McGraw Hill, 2nd edition, 2008.
[11] A. V. Rao. A comment on: Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, 24(4):639–640, 1973.
[12] B Sani and B. G. Kingsman. Selecting the best periodic inventory control and demand forecasting methods for low demand items. *Journal of the Operational Research Society*, 48:700–713, 1997.
[13] L. Shenstone and R. J. Hyndman. Stochastic models underlying Croston's method for intermittent demand forecasting. *Journal of Forecasting*, 24:389–402, 2005.
[14] E. A. Silver, D. F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling*. John Wiley & Sons, 1998.
[15] A. A. Syntetos and J. E. Boylan. On the bias of intermittent demand estimates. *International Journal of Production Economics*, 71:457–466, 2001.
[16] A. A. Syntetos and J. E. Boylan. The accuracy of intermittent demand estimates. *International Journal of Forecasting*, 21:303–314, 2005.
[17] R. Teunter and B. Sani. On the bias of Croston's forecasting method. *European Journal of Operational Research*, Article in Press, 2008.
[18] T. Wild. *Best Practice in Inventory Management*. Butterworth-Heinemann, 2nd edition, 2002.
[19] P. H. Zipkin. *Foundations of Inventory Management*. McGraw-Hill, 2000.