

2 Inventory Management

The objective of inventory management is to fulfill the customer's needs by stocking the right quantity at the right time such that the resulting costs are minimal. It is a crucial function in most companies, and in general it cannot be separated from other functions. For example, the optimal inventory policy will certainly depend on promotion campaigns conducted by the marketing department. Nevertheless, this chapter deals with the technical part of inventory management and methods in order to find optimal inventory policies. Additionally, when focussing on retail companies, only inventories of finished goods will be considered, which are referred to as stock keeping units (SKUs).

This chapter is divided into three parts. Section 2.1 deals with the measurement of the supply chain performance. Thus, it describes what is meant by fulfilling the customer's needs. After this, Section 2.2 gives an overview of the relevant costs, their relationships, and how they can be estimated based on a company's balance sheet. The last and main part of this chapter introduces different inventory policies and the decision variables which can be set in order to fulfill the customer's needs at minimal costs. It focuses on stochastic inventory models and methods to find an optimal inventory policy.

2.1 Supply Chain Performance Measurement

This section considers the measurement of supply chain performance, i.e. what separates a good inventory management system from a bad one. Three different measures are described, namely the α -, the β -service, and

the inventory turnover. All measures are widely used in theory and practice and feature two different views of what is meant by supply chain performance. The α - and the β -service measure the performance from the customer's perspective whereas the inventory turnover is a technical measure to describe the supply chain performance in sense of capital efficiency. This section presents these three performance indicators and gives an example to state the properties of those views.




Figure 2.1: Inventory levels of one SKU over five weeks

To introduce the terms and methods of inventory management, Figure 2.1 shows the continuously sampled inventory level of an exemplary SKU over five weeks. The inventory level starts at 10 pieces, and the underlying inventory policy is static. This means that at the beginning of each week a delivery of 20 pieces arrives, and the inventory level increases. The continuous tracked demand has a relatively high variance. Combined with the static inventory policy, this leads to negative inventory levels at the end of weeks 2 and 4. This backlog, i.e. the demand which waits to be satisfied, is marked with a gray shade. The backlog adds up to 10 pieces at the end of week 2 and 1 piece at the end of week 4. By way of representing every single inventory change, the figure gives a very detailed view of what is really happening with the inventory level. Nevertheless, in most circumstances the data is not that detailed, but aggregated for different periods. Thus,

Table 2.1 lists the weekly aggregated data of Figure 2.1. The demand is aggregated as the sum of sold pieces within a week. The second row lists the inventory level at the end of each week. Therefore, the inventory level at week 0 is the starting inventory level of the first week. The inventory level in each period is equal to the inventory level of the previous week plus deliveries of the current week minus the demand during the current week. For example, the inventory level at the end of week 3 is 5, because it starts with -10 , the demand is 5 and a quantity 20 is delivered at the beginning of the week ($(-10) - 5 + 20 = 5$). In this setup it is assumed that customer demands which cannot be fulfilled are backordered, which means the customer will wait for the next delivery if the SKU is understocked. This assumption will not hold in every case, but if one considers the inventory of an online store, this might be realistic. The assumption that unsatisfied customers will wait for the next delivery reduces the complexity and is therefore well suited for an introductory example.

	Week 0	Week 1	Week 2	Week 3	Week 4	Week 5
Demand		20	40	5	26	10
Inv. level	10	10	- 10	5	- 1	9
Deliveries		20	20	20	20	20

Table 2.1: Exemplary demand series with corresponding inventory

α -Service Level

The α -service focuses on customer satisfaction. It measures the probability of serving the demand in a period directly from stock. In other words the α -service measures the probability of completely satisfying a customer. The α -service is given by:

$$\alpha = P(y_t \leq I_t) \quad , \quad (2.1)$$

where y_t is the demand of a customer or the demand within a period and I_t is the inventory level at time t . For the given example the α -service is 60%, because in 2 out of 5 periods the whole demand could not be delivered directly from stock, no matter if the backlog is 10 or just 1 piece as in period 4.

Fixing an α -service target can be useful if the customer is unsatisfied no matter if the shortage is one or a hundred pieces, as is the case in the automotive sector, in which the production line stops at any shortage (Nahmias, 2009).

β -Service Level

In contrast to the customer focus of the α -service, the β -service level focuses on delivered quantities. It measures the share of demand which could be delivered directly from stock with no delay. Because of that the β -service is usually used as a target for a warehouse. It is given by:

$$\beta = 1 - \frac{E(y_t | y_t > I_t)}{E(y_t)} \quad , \quad (2.2)$$

where $E(y_t | y_t > I_t)$ is the conditional expectation of the demand which exceeds the current inventory level. For the given example β equals 89.1% because 11 of 101 pieces could not be delivered directly from stock.

It can be seen that by choosing one service level over the other totally different values can occur. Thus, there is no overall right decision. It depends upon circumstances. If it does matter how many pieces of an order could not be delivered, one should use the β -service, but if it is an all or nothing evaluation of whether the inventory management is successful, the α -service is the appropriate performance measure.

Inventory Turnover

The inventory turnover η indicates how often the average inventory level is sold within a given period. Therefore, it focuses on the fixed capital of the inventory. Additionally, as a dimensionless measure it allows to compare the inventory level of two different SKUs. It is defined as the ratio of the sales of a SKU and its average inventory level:

$$\eta = \frac{\sum_{t=1}^T y_t}{\frac{1}{T} \sum_{t=1}^T I_t} \quad . \quad (2.3)$$

Compared with the service measures, one can see that the inventory turnover does not regard the customer at all. It is a purely inventory focused measure and is not used to find an optimal inventory policy. In most setups inventory management needs to fulfill a certain service constraint while minimizing the total inventory costs. Thus, a high inventory turnover results from a good inventory policy. The average inventory level in each period is the mean value of the inventory level at the beginning of the period (inventory level at the end of the previous period plus deliveries) and the inventory level at the end of the period. Therefore, the average inventory levels of the five periods are 20, 10, 7.5, 13, and 14. For the given example, η equals 7.83 because in total 101 pieces are sold and the overall average inventory level is 12.9.

2.2 Relevant Costs

In order to minimize the total inventory costs, it is crucial to have information about what the relevant costs are. This fact might be trivial, but Themido et al. (2000) state that there is a lack of appropriate cost information in many companies. In addition, Askarany, Yazdifar, and Askary (2010) argue that there is still a high potential to increase the accuracy of the internal cost estimates of companies. It is not the focus of this work to describe the methods to estimate costs, but every optimal inventory control system is based on those estimates. Therefore, in addition to the description of the relevant costs, there will also be a reference to estimation methods.

The relevant costs of an inventory control system are split into three different parts. First, the order costs which arise with an order. These are costs like the handling, transportation or labor costs, as well as the variable costs of an order, which may vary due to quantity discounts (Axsäter, 2006, p. 44ff). The order costs per unit decrease as the order quantity increases. So there is a positive impact of high order quantities on order costs. This part of the inventory costs can be estimated with data provided by internal accounting. Themido et al. (2000) and Everaert et al. (2008) discuss the use of activity-based costing models to determine the handling, transportation and order costs. They argue that the use of activity-based costing models can increase the accuracy of a company's logistic costs estimation.

The second part of the inventory costs are the holding costs, which occur when keeping SKUs on stock. The holding costs again consist of different parts. There is the interest which needs to be paid for the fixed capital and the opportunity costs of using this fixed capital in a different way. There are the costs of running the warehouse, and even if they are non-cash expenses, the risk of spoilage of perishable goods and thievery needs to be regarded because they increase at higher inventory levels (Axsäter, 2006, p. 44). Overall the holding costs will rise with the order quantity, thus there is a

negative impact of a high order quantity on holding costs. The estimation of the warehouse costs, spoilage and thievery can be done based on data provided by internal accounting whereas the opportunity costs of capital can be estimated using the capital asset pricing model (CAPM).

The CAPM was independently developed by Sharpe (1964) and Lintner (1965) to value capital assets based on their risk. The concept of this approach is to assume a linear relationship between the rate of an asset's return which an investor expects and its risk. Thus, for a given risk, this model can be used to estimate the expected rate of return, i.e. the opportunity costs of capital (Singhal and Raturi, 1990). Sharpe (1964) and Lintner (1965) propose a risk measure which is based on the relationship between the risk of the asset and the risk of the market portfolio. It can be calculated as follows:¹

$$\omega_i = \frac{\text{cov}(u_M; u_i)}{\text{var}(u_M)} \quad , \quad (2.4)$$

where $\text{cov}(u_M; u_i)$ denotes the covariance between the rate of return of the asset and the rate of return of the market. $\text{var}(u_M)$ denotes the variance of the return of the market portfolio. Based on this risk ω_i measure, the expected rate of return can be calculated using the following equation:

$$u_i = u_f + \omega_i \cdot (u_M - u_f) \quad , \quad (2.5)$$

where u_f is the return of a risk free asset. Jones and Tuzel (2013) showed a relationship between the risk measure of the CAPM and the inventory levels of a company in an empirical study. Therefore, there is empirical evidence that the estimated ω_i influences inventory decisions. The CAPM is used to estimate the opportunity costs of capital in different industry sectors in Europe in order to give an indication about the inventory holding costs. This might not be a very accurate measure, but it will provide a

¹The proposed risk measure is usually denoted as β , but for reasons of consistency ω will be used.

suitable estimate. Table 2.2 lists the results of an empirical analysis using the CAPM. It is based on the stock returns of 9833 European companies in 39 industry sectors between January 2000 and May 2014 (Bloomberg, 2014). The different ω -values were calculated by comparing the returns of the different sector portfolio, i.e. an equally weighted portfolio containing all the companies of a sector, and the market portfolio containing all 9833 shares. The risk-free rate was assumed to be 0.045, which equals the average return of a German federal bond during this period. In addition to the expected interest rate, the table also lists the 95% confidence interval of the expected interest rate, the average gross margin, and the number of sampled companies of the sector.

The third part of the inventory costs arise with a stock out. They are called shortage or lost sales costs. To define these costs, it needs to be distinguished whether, in case of a shortage, the customer is willing to wait until the product is available again. If so, there are additional costs of express deliveries or compensation discounts. If the customer is not willing to wait, the sale is lost and therefore, the company will not earn the gross margin. In addition, the customer is unsatisfied. This raises the risk that the customer may switch to a competitor. The costs of an unsatisfied customer are called loss of goodwill costs and whereas the gross margin is relatively easy to estimate (see Table 2.2), the estimation of the loss of goodwill costs is almost impossible (Crone, 2010, p. 39). The risk of a stock out falls with higher order quantities. Thus, higher order quantities have a positive impact on the lost sales costs.

After describing the three parts of the inventory costs, it can be seen that there is a trade-off between high order quantities, which reduce the order and lost sales costs, and the low order quantities, which cut the holding costs. Therefore, one goal of inventory management is to find an order quantity, which balances these different costs.

ICB sector name	u_i	CI _{0.95}	gross margin	n
Aerospace & Defense	0.115	[0.103;0.127]	0.268	40
Alternative Energy	0.137	[0.119;0.154]	0.217	66
Automobiles & Parts	0.178	[0.161;0.195]	0.207	74
Beverages	0.071	[0.061;0.081]	0.471	92
Chemicals	0.134	[0.124;0.144]	0.284	120
Construction & Materials	0.221	[0.160;0.283]	0.233	245
Electricity	0.076	[0.069;0.084]	0.382	105
Electronic & Electrical Equipment	0.113	[0.106;0.120]	0.366	219
Equity Investment Instruments	0.110	[0.102;0.117]	0.770	344
Financial Services	0.064	[0.055;0.072]	0.468	652
Fixed Line Telecommunications	0.110	[0.065;0.156]	0.419	47
Food & Drug Retailers	0.106	[0.098;0.115]	0.223	54
Food Producers	0.071	[0.066;0.076]	0.278	209
Forestry & Paper	0.123	[0.111;0.136]	0.292	42
Gas, Water & Multiutilities	0.083	[0.072;0.093]	0.376	40
General Industrials	0.109	[0.100;0.118]	0.246	84
General Retailers	0.201	[0.174;0.228]	0.376	207
Health Care Equipment & Services	0.104	[0.096;0.113]	0.519	212
Household Goods & Home Construction	0.108	[0.100;0.116]	0.281	136
Industrial Engineering	0.125	[0.117;0.133]	0.299	264
Industrial Metals & Mining	0.216	[0.190;0.242]	0.132	78
Industrial Transportation	0.089	[0.082;0.097]	0.257	241
Leisure Goods	0.094	[0.079;0.108]	0.405	78
Media	0.087	[0.070;0.105]	0.447	317
Mining	0.095	[0.076;0.114]	0.238	202
Mobile Telecommunications	0.233	[0.187;0.278]	0.429	33
Nonequity Investment Instruments	0.110	[0.104;0.115]	0.483	1625
Nonlife Insurance	0.121	[0.112;0.130]	0.587	73
Oil & Gas Producers	0.156	[0.135;0.177]	0.334	159
Oil Equipment, Services & Distribution	0.128	[0.117;0.140]	0.359	81
Personal Goods	0.106	[0.095;0.116]	0.453	149
Pharmaceuticals & Biotechnology	0.129	[0.120;0.138]	0.633	240
Real Estate Investment & Services	0.071	[0.066;0.076]	0.577	403
Real Estate Investment Trusts	0.089	[0.081;0.096]	0.817	97
Software & Computer Services	0.107	[0.098;0.115]	0.559	512
Support Services	0.102	[0.095;0.110]	0.312	375
Technology Hardware & Equipment	0.136	[0.122;0.150]	0.402	165
Tobacco	0.128	[0.110;0.146]	0.379	4
Travel & Leisure	0.072	[0.059;0.086]	0.282	333

Table 2.2: Expected interest rate and gross margin of European industry sectors.

2.3 Inventory Policies

Inventory management answers the question of when to order and how much should be ordered. The inventory policy is the strategic framework which determines the way those questions are answered, i.e. whether an order will be placed on a regular basis or whether the order quantity is constant. The different inventory policies are defined by the variables the inventory management decides on. The literature provides a vast number of inventory policies to find optimal solutions for various different circumstances (see Bakker, Riezebos, and Teunter (2012) and Goyal and Giri (2001) for an overview). Therefore, this section only describes a subset of inventory policies which are frequently used in practice.

The literature distinguishes mainly between four different inventory policies. They can be differentiated from each other by considering the flexibility of the time when an order is placed and the amount which is ordered. There could either be a fixed interval r between two consecutive orders, or an order could be placed whenever the inventory level falls below a certain level s . The same scheme holds for the order quantity, which can be either fixed or variable. If the quantity is fixed for all orders, it is denoted as Q . If instead the order size depends on the current inventory level and should fill the stock up to a given level, this order up to level is denoted as S . The third variable represents the review interval, i.e. whether if the inventory level is known at every point in time or just at given intervals T . If $T > 0$ the inventory policy is referred to as periodic review. However, if $T = 0$ the policies are called continuous review because the inventory level is known at each point in time. The different inventory policies are defined by the combination of those decision variables. Thus, there are four different cases:

- (r, Q, T) - fixed time between orders, fixed order quantity
- (r, S, T) - fixed time between orders, variable order quantity
- (s, Q, T) - variable time between orders, fixed order quantity

- (s,S,T) - variable time between orders, variable order quantity

These four inventory policies also include several special cases. For example, the news vendor model for highly perishable goods can be formulated as a base-stock policy (Bouakiz and Sobel, 1992). The intuition behind those base-stock policies is to order whenever the inventory level is reduced such that the inventory level is S at the beginning of each period. They are in fact (s,S,T) policies, but due to their special properties they are referred to as $(S-1,S,1)$ policies.




Figure 2.2: Different inventory policies

To state the behavior the four different inventory policies, Figure 2.2 shows the inventory levels and order sizes over 30 weeks of simulated inventories using each of these different policies. The rows separate the decision about when an order is placed whereas the columns separate a fixed order size from a flexible one. All four graphs are based on the same demand series

and parameters. The order size Q and the order up-to-level S are set to 30 the fixed interval between two orders r is 10, and the reorder point s is 5. Once an order is placed, it takes 2 weeks until the delivery arrives. Thus, the lead time L is 2 and the review period T is 1 week for all graphs. The upper-left graph plots the inventory and order sizes of the (r, Q, T) policy based on the given parameters. The first order is placed in period 4 and arrives two weeks later in period 6. One can see that all order sizes are equal, and the time between two consecutive orders is fixed at 10 weeks. The dotted line denotes the inventory position, which equals the inventory level plus the future deliveries. Therefore, the inventory position and inventory level are equal if no delivery is in progress. As in Figure 2.1 the gray shaded inventory marks a backlog. For the first policy this is the case in period 5, 15, 23, 24 and 25.

The second graph shows the results of a (r, S, T) policy and has the same structure as the first graph. Additionally, the horizontal dashed line marks the order-up-to level S . As in the first graph, the orders are equidistant, but in this graph they differ in size. The order size is selected in the way that the order will raise the inventory position to S , in other words the order size is equal to $S - I_t$. Thus, the order-up-to level is an upper bound of the inventory position. The inventory level will not reach the order-up-to level in most cases, even if the inventory position does, due to the demand during the lead time L .

The third graph shows the results of a (s, Q, T) policy. In this case an order of Q pieces is placed whenever the inventory position drops below s (dashed line). The inventory position is used for the order rule because it gives clear instructions, unlike the inventory level. Consider an order policy based on the inventory level. An order would be placed not only in week 3, but also in week 4, because the inventory level is still below s since the delivery has not arrived yet. In contrast, a rule based on the inventory position will not lead to further orders. Only one order is placed in week 4, because the

inventory position regards future deliveries. Additionally, it can be seen that the reorder point s is a lower bound of the inventory position.

The last graph shows the results of a simulated (s, S, T) policy. The inventory position fluctuates between the reorder level s and the order-up-to level S . Similar to the third graph, the first order is placed in week 3, whereas the second and third orders are placed in period 13 and 20. Backlogs appear in period 14 and 21. This inventory policy has the highest adaptability, and Sani and Kingsman (1997) show that in case of intermittent demand series, the (s, S, T) policy suits best. Therefore, a (s, S, T) policy is described and used in the following.

2.3.1 Stochastic Inventory Models

The challenges which arise from an uncertain demand are based on the combination of two circumstances, the uncertain demand, and the positive lead time. For example, if there is no lead time ($L = 0$) one would place an order whenever the inventory level falls to zero ($s = 0$) no matter how volatile the demand is because it would be delivered immediately. If the demand is certain and the lead time is positive, one would place the order such that the remaining stock lasts until the next order arrives. As this generally also holds for the case of a positive lead time and an uncertain demand, a new problem arises. It could be the case that the reorder point s , i.e. the amount that should last until the new order arrives, is too low. Thus, the demand cannot be satisfied from stock, and will be backlogged. This leads to shortage costs and unsatisfied customers. Therefore, an inventory policy based on stochastic demands needs to regard those uncertainties and the resulting costs. This can be done by finding a set of inventory variables, e.g. the reorder point s and the order-up-to level S , which minimizes the total inventory costs as the sum of order, holding, and shortage costs. But, as mentioned above, it is hard to estimate the loss of goodwill costs which provide the main part of the shortage costs (Crone, 2010, p. 39). To avoid this issue, the selection of s and S can be based on minimizing the sum of the holding and order costs while satisfying a certain service constraint, e.g. find s, S such that an α -service level of 95% is achieved at minimal cost. This method avoids the need to determine the shortage costs because they are implicitly assumed by selecting a service target. Those service constrained inventory policies will be used in the following.

2.3.2 Determination of the Order-Up-To-Level S

As described above by using a (s, S, T) policy the inventory position will vary between the two bounds s and S . There are some approaches to determine S based on the demand series directly (e.g., see Teunter, Syntetos, and

Babai (2010)), but in most cases S results from selecting a reorder point s and the gap between s and S , denoted as D . This position gap D is selected such that the sum of the order and inventory costs are minimized. Thus, the remainder of this section describes the economic order quantity model to determine the gap D .

The widest used inventory model is the economic order quantity model developed in F. W. Harris (1990).² This model can be denoted as a $(r, Q, 0)$ policy because the order interval and quantity are constant, and the inventory is known at every point in time. It determines the cost optimal position gap D^* based on the assumption that the lead time is zero and that the underlying demand series is a continuous-time continuous-space time series with a constant demand rate. The demand needs to be fulfilled immediately, i.e. shortages are not allowed. Based on these assumptions, the optimal solution can be derived in a straightforward manner. As shortages are prohibited and the lead time is zero, a cost optimal order is placed everytime the inventory level reaches 0. Therefore, the total costs $C(D)$ are determined by the order and holding costs:

$$C(D) = \underbrace{\frac{\mu_y}{D} \cdot K}_{\text{order costs}} + \underbrace{\frac{D}{2} \cdot h}_{\text{holding costs}}, \quad (2.6)$$

where D is the inventory position gap, μ_y is the demand in a given period, K are the fixed order costs and h are the holding costs per unit and per period. F. W. Harris (1990) defines the order costs as product of the number of orders ($\frac{\mu_y}{D}$) and the fixed order costs K . The holding costs are defined as the product of the holding costs per unit and per period h and the average inventory level in the period ($\frac{D}{2}$). Thus, in order to calculate the optimal inventory position gap using the *EOQ* model, D^* results from the minimum of the total costs function $C(D)$, i.e. the root of the first

²The original publication by Harris dates back to 1913 while this citation refers to the digital reprint in 1990.

derivative.

$$\underset{D}{\text{minimize}} \quad C(D) = \frac{\mu_y}{D} \cdot K + \frac{D}{2} \cdot h \quad \Rightarrow \quad D^* = \sqrt{\frac{2 \cdot \mu_y \cdot K}{h}} \quad . \quad (2.7)$$

The underlying assumptions of the *EOQ* are very strict and will not hold in many practical setups. Therefore, the literature provides extensions to the *EOQ* in order to reduce the strictness of the assumptions. Axsäter (2006, p. 55 ff) considers the *EOQ* in case of finite production rates and discounts. Grubbström and Erdem (1999) derive an adapted version of the *EOQ* with backlogging and Weiss (1982) and Goh (1994) regard nonlinear holding costs. There is also a wide range of literature dealing with imperfect SKU quality (see Khan et al. (2011) for a review). A historical review and several other extensions of the *EOQ* are given in Choi (2013).

Haneveld and Teunter (1998) show that in case of slow moving demand, an inventory policy based on maximizing the discounted cash flows instead of minimizing the inventory costs leads to better results. This addresses the main drawback of the *EOQ* model. The strict assumption of a certain and fixed demand rate. While this leads to the straightforward derivation of the optimal order quantity D^* , it does not reflect the majority of practical settings in which the future demand is uncertain and variable.

Nevertheless, the original *EOQ* model is still frequently used due to its easy implementation and robustness against violations of the assumptions (Drake and Marley, 2014). Therefore, and for reasons of comparability, the *EOQ* will be used to determine the inventory gap D even if there may be potential issues.

2.3.3 Determination of the Reorder Point s

The determination of the reorder point s differs depending on whether an α - or β -service constraint is specified. As described in Section 2.1, the α -service is the probability of satisfying the entire demand in a period, or in

other words not going out of stock. In general the literature provides the following rule to determine the reorder point for a given α -service constraint (Schneider, 1978):

$$\int_{-\infty}^s f_{ltd}(x) dx = \alpha \quad . \quad (2.8)$$

s is implicitly defined as the α -quantile of the lead time demand. This formula only holds if the lead time demand is approximated by a continuous distribution, which is always a simplifying assumption if one considers the demand series. Thus, it needs to be transferred into the discrete case:

$$\sum_{i=0}^s p_{ltd}(i) \geq \alpha \quad . \quad (2.9)$$

This rule has the same interpretation as (2.8) and it is theoretically much closer to an observed demand series, but it has the drawback that it does not have an unique solution for s . The reorder point is the smallest s which satisfies condition (2.9). This is equal to the definition of the Value-at-Risk, a frequently used risk measure for financial assets in portfolio theory. Therefore, the determination of s for a given α -service level can also be interpreted by means of portfolio theory. Figure 2.3 shows the quantile function of the lead time demand. The dashed lines indicate an α -service level of 95% and the resulting s .

If the lead time distribution is approximated by a continuous distribution, Schneider (1978) derives an implicit calculation rule to determine s for a given α -service constraint. In most cases the assumption of a Gaussian lead time demand distribution is used to calculate s . The main reason for this is the central limit theorem and the convolution invariance of the normal distribution. This means that the sum of two independent normally distributed random variables, is normally distributed again. The assumption of a normal distribution might be no problem for fast moving SKUs, but in




Figure 2.3: Reorder point determination subject to an α -service constraint

the case of intermittent demand series this assumption could lead to poor results. Intermittent demand series have a low mean value and a relatively high variance. Therefore, a normal distribution based on those values will have a considerable positive density in the negative range. To avoid these problems, the literature provides inventory policies which are based on the gamma distribution (e.g., see Dunsmuir and Snyder (1989) and Moors and Strijbosch (2002)). This distribution is only defined for positive values and is therefore a suitable supplement to the normal distribution. In addition, the gamma distribution is also completely defined by the first two central moments and is convolution invariant. Thus, in the following the reorder points are approximated using the normal distribution and additionally the gamma distribution.

If a normal distribution is assumed, the reorder point s can be found as the root of the following approximation:

$$f(q) = -\frac{(1-\alpha)\left(D + \frac{\mu_{2ltd}}{2 \cdot \mu_{ltd}}\right)}{\sigma_{ltd}} - \frac{1}{2}q \operatorname{erfc}\left(\frac{q}{\sqrt{2}}\right) + \frac{e^{-\frac{q^2}{2}}}{\sqrt{2\pi}} \stackrel{!}{=} 0 \quad , \quad (2.10)$$

$$s = \mu_{ltd} + q \cdot \sigma_{ltd} \quad , \quad (2.11)$$

where μ_{ltd} is the expectation, μ_{2ltd} is the second moment and σ_{ltd} is the standard deviation of the lead time demand. $\operatorname{erfc}(x)$ is the complementary error function (Abramowitz and Stegun, 1972, p. 297). If a gamma distribution is assumed, s can be found as the root of the following approximation

$$f(s) = \frac{\mu_{ltd} \cdot \Gamma(p+1, b \cdot s)}{\Gamma(p+1)} - \frac{s \cdot \Gamma(p, b \cdot s)}{\Gamma(p)} - (1-\alpha) \left(D + \frac{\mu_{2ltd}}{2 \cdot \mu_{ltd}}\right) \stackrel{!}{=} 0 \quad , \quad (2.12)$$

where $\Gamma(x, y)$ is the incomplete gamma distribution (Abramowitz and Stegun, 1972, p. 260), $p = \frac{\mu_y}{\sigma_y^2}$ and $b = \left(\frac{\mu_y}{\sigma_y}\right)^2 \cdot (L+1)$.

In case of a β -service constraint, s generally can be found by using the following rule if a continuous lead time distribution is assumed. s is implicitly defined as the limit for which the upper partial moment of the lead time distribution, i.e. the expected lost sales, is equal to the lost sales restriction $((1-\beta) \cdot D)$.

$$\int_s^\infty (x-s) f_{ltd}(x) dx = (1-\beta) \cdot D \quad . \quad (2.13)$$

As in (2.8), Equation (2.13) is transferred into (2.14) for a discrete space lead time distribution. Again this transformation results in a non-unique solution for s . Because the expected lost sales reduce with a higher reorder point, the reorder point is selected as the lowest s which satisfies the

inequation (2.14):

$$\sum_{i=s}^{\infty} (i-s) p_{ltd}(i) \leq (1-\beta) \cdot D \quad . \quad (2.14)$$

As mentioned above, the α -service can be interpreted by means of portfolio theory. This is also the case for the β -service even if it is not a perfect match. The expected shortfall is a risk measure from portfolio theory. In contrast to (2.14) it tracks the lower partial moment but the general intuition remains the same. Therefore, both the α - and β -service may also be interpreted as risk measures. Figure 2.4 shows the expected lost sales (upper partial moment) of the lead time demand distribution in relation to the order size. The dashed lines indicate the resulting reorder size s if the β -service is equal to 95% ($1-\beta = 0.05$).




Figure 2.4: Reorder point determination subject to a β -service constraint

Schneider (1978) also derives an implicit solution of s for a continuous lead time distribution assumption if a β -service constraint is given. If a normal distribution is assumed, the reorder point s can be calculated using the following approximation. The root of (2.15) can be interpreted as the safety

factor (Schneider, 1978), while s results from the reverse z-transformation of q .

The notation is equal to (2.10).

$$f(q) = \frac{(\beta - 1)(2D\mu_{ltd} + \mu_{2ltd})}{\sigma_{ltd}^2} + \frac{1}{2}(q^2 + 1) \cdot \operatorname{erfc}\left(\frac{q}{\sqrt{2}}\right) + \frac{e^{-\frac{q^2}{2}}}{\sqrt{2\pi}} q \stackrel{!}{=} 0 \quad , \quad (2.15)$$

$$s = \mu_{ltd} + q \cdot \sigma_{ltd} \quad . \quad (2.16)$$

In case of a gamma distributed lead time demand Schneider (1978) argue that the provided approximation only holds if the value of D is high. This might be a problem due to the low mean of intermittent demand series. Thus, it is referred to the procedure of Dunsmuir and Snyder (1989). Their approach is also based on the upper partial moment of the gamma distribution, but they regard the properties of intermittent demand series. The implicit rule provided in Dunsmuir and Snyder (1989) in order to calculate the reorder point s has the same structure as equation (2.13). It is also based on the integral over the gamma density function, but this integral can be solved using the incomplete gamma function. This leads to the following implicit calculation rule for the reorder point s .

$$\frac{e^{-v \cdot s} \cdot (v \cdot s)^{v-1} \cdot s - \Gamma(v, v \cdot s) \cdot (s-1)}{\Gamma(v)} = \frac{(1-\beta) \cdot S}{\pi_{ltd}} \quad , \quad (2.17)$$

$$v = \frac{(\mu_{ltd}^+)^2}{(\sigma_{ltd}^+)^2} \quad , \quad (2.18)$$

where the parameters μ_{ltd}^+ and σ_{ltd}^+ are the expectation and variance of the positive demands during lead time. They can be calculated using the following rule (Dunsmuir and Snyder, 1989, p. 17):

$$\mu_{ltd}^+ = L \cdot \frac{\mu_Y}{\pi_{ltd}} \quad , \quad (2.19)$$

$$(\sigma_{ltd}^+)^2 = \frac{L \cdot \sigma_y^2 + (1 - \pi_{ltd}) \cdot \pi_{ltd} \cdot (\mu_{ltd}^+)^2}{\pi_{ltd}} \quad , \quad (2.20)$$

$$\pi_{ltd} = 1 - (1 - \pi_Y)^L \quad , \quad (2.21)$$

where π_Y is the share of positive demands.



<http://www.springer.com/978-3-658-14061-8>

Managing Intermittent Demand

Engelmeyer, T.

2016, XV, 157 p. 65 illus., Softcover

ISBN: 978-3-658-14061-8