A glass bottle allocation model for Heineken

Master thesis industrial engineering and management

Chris H.J. Beckers

June 2009

Graduation committee:

Dr. Ir. J.M.J. Schutten

J.I.H. Buter

S. Jensma

Management summary

This research aims to improve the allocation of the beer bottle demand of Heineken to its suppliers. To this end, we have developed a goal programming model that is capable of optimizing the allocation on both costs and security of supply. The advantage of a goal programming model for Heineken is that it does not only provide efficient solutions, but also indicates the best possible solutions on all objectives. This way, Heineken improves its insight on the quality of the chosen allocation.

For Heineken, three aspects of the supply are very important: costs, quality, and security. The quality of the bottles and the production process is guarded by regulations and audits. In general, this leads to a high reliability of the suppliers. The production process, however, is complex and occasional disruptions are inevitable. These disruptions can have a serious impact on the delivery of bottles and, therefore, security of supply and costs are considered the most important objectives during the allocation.

Several reasons exist for the high impact of disruptions. The costs of overcapacity at glass factories are very high, due to the high fixed costs. Suppliers therefore tend to minimize free capacity. The production speed is limited by the maximum temperature of the furnace, leaving limited possibilities to increase the capacity. Moulds, used in the production of beer bottles, are machine specific and have a lead-time of approximately three months. Therefore, no quick alternatives can become available at other factories in case of a disruption either. Measures must already have been built-in in the allocation, in order to secure enough supply at times of disruptions.

According to the literature, security of supply is mainly provided by redundancy. Additional suppliers, additional production capacity, and stock can be used to mitigate disruptions.

Allocation problems in the literature are often solved by mathematical programming. Based on the number of objectives for Heineken (minimizing costs and maximizing security of supply), different types of programming models can provide solutions. Goal-programming requires the decision maker to explicitly state the goal values. A comparison of the allocation's cost and security of supply values with these goal values increases the visibility of the quality of the allocation solution. The goal programming model finds a balance between security of supply and costs based on the preference weights of the decision maker.

Table of Contents

List of figures

List of tables

Chapter 1: introduction

Worldwide, approximately 64% of all beer is consumed from bottles (Plato Logic Ltd., 2007). Consequently, the procurement of glass bottles represents a large part of Heineken's total costs. Therefore, selection of the best bottle suppliers to provide these bottles is of major importance.

1.1: Heineken

The Heineken Group is the third largest brewer in the world with a total market share of 9% (Figure 1). Its main brand is the premium beer Heineken, which is the most important premium brand in the world. It sells almost 25 million hectolitres a year (1,5% of the total beer consumption) (Canadean, 2007). Currently, Heineken has 132 breweries located in over 70 countries. Decentralized business units, called OpCos, locally control these breweries.

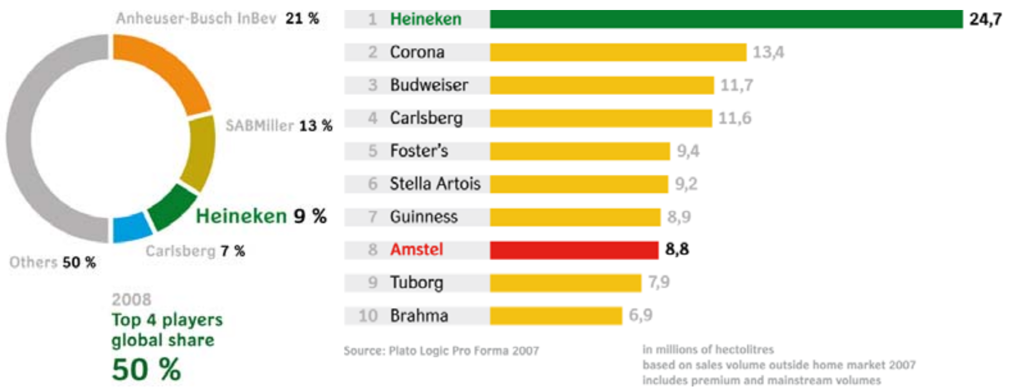

Figure 1: Heineken market share 2007

1.2: portfolio diversification

The increase in the number of different beer types is inherent to a market share growth by take-overs, as the entire portfolio of the acquired company is added to the currently existing portfolio. After recent take-overs, Heineken sells over 170 international, regional, local, and specialty beers all around the globe (Heineken, 2007), about 70% of which in its key markets (Canadean, 2007).

The key markets require a more differentiated portfolio. In the United States, the growing importance of health and wellness led Heineken to introduce the first deviation from the Heineken brand: Heineken light. The acceptance of Hispanic and Asian beer in the US becomes larger as well. For instance in Spain, many marketing initiatives and newly added flavours have led to growth in the discount and the no/low alcohol segment.

1.3: supply market

Glass bottles are produced from commonly available raw materials, which are slowly melted in furnaces at high temperatures. A furnace contains hundreds of tonnes of raw materials and requires constant heating to keep the glass melted. For this reason, it is not economical to shut down a furnace for any period shorter than a month. Glass production is a continuous process. The rate of the production process is bounded by the speed of melting, which in turn is bounded by the temperature allowed by the furnace structure. Capacity flexibility is thus limited by economic and technical reasons.

The construction of new production facilities takes a year and a half of planning. There are high risks involved as well. The plants cost tens of millions of Euros. Moreover, the demand market is limited to a 400-kilometer radius around the plant. The costs of transportation are too high, due to the weight and size of glass bottles, to deliver the bottles further away. The high costs of overcapacity taken into account, suppliers are reluctant to add new production lines. This results in limited sourcing options for glass bottles. Because of the limited inflexible capacity, there is a high interdependence between the sourcing decisions of the different bottles. Capacity usage of one bottle type adds restrictions to the allocation of all the others.

1.4: problem formulation

We have formulated the following research goal:

To improve the sourcing of Heineken's glass bottle demand and to gain more insight into the quality of the sourcing.

In order to reach this goal we answer the following research questions:

1. How does Heineken currently source glass bottles?
2. What does the literature related to Heineken's purchasing process indicate?
3. How can we design an improved sourcing solution in terms of complexity and measure the quality of the improved solution?
4. How can we implement this solution into the working procedures of Group Purchasing?

The outline of this thesis is as follows. Chapter 2 describes the current purchasing procedures, the production process of glass bottles, and the competitiveness of the supply market. Due to the sensitivity of this information, only the main conclusions of this chapter are included in the public report. Chapter 3 gives an overview of the literature regarding this research. Based on the literature review, we conclude what is the best solution type for our problem. Chapter 4 discusses the solution design.

Chapter 2: purchasing process

Heineken procures a large portfolio of bottles of different colors each year. The supplier selection decisions concerning these bottles are interdependent. Two main reasons are that capacity at the factories is limited and the number of factories near the breweries small. Multiple factories are considered to source each bottle type. For every factory, purchasers from Heineken have to decide what types of bottles are sourced there and in what amounts. It is important to be able to quickly find the optimal allocation. Three criteria are important in the allocation process: costs, security of supply, and quality.

Quality is used as a qualification criterion. Heineken has designed policies to ensure the quality of the bottles and the processes of the suppliers. This way, supply disruptions can be largely avoided. However, there is always a risk of disruptions. Security of supply is concerned with the impact of those disruptions. Dealing with the impact is very difficult, since the flexibility of the suppliers is very small. Although sudden small changes in the production schedule are often possible, there is no possibility to significantly increase capacity. Protection against disruptions should therefore be built-in in the allocation.

Chapter 3 will give an overview of the different solutions presented in the literature to solve the supplier selection problem. With respect to the security of supply, it discusses the measurement (what is security of supply?), the ways to increase security, and the ways to incorporate these techniques into the selection.

Chapter 3: literature review

Chapter 2 presented the characteristics of the purchasing process at Heineken. It defined two subjects for which a literature review was required. Section 3.1 describes the literature on supplier selection. Section 3.2 deals with the literature on security of supply. Section 3.3 combines the two previous sections by describing research on supplier selection with incorporated security of supply. Based on the opportunities found in the literature and the purchasing process at Heineken, we propose a solution direction in Section 3.4. Chapter 4 will further develop this solution.

3.1: supplier selection

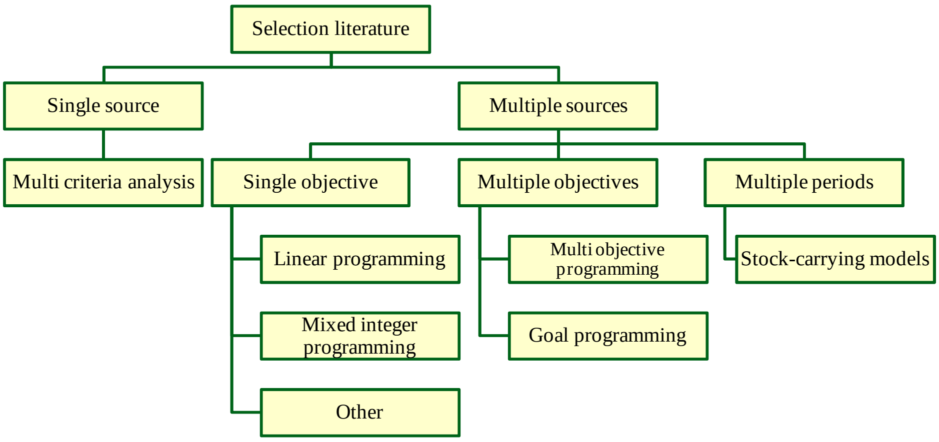

Figure 2: overview of supplier selection literature

During the problem definition phase of supplier selection (as explained in Section 2.2.2), the type of sourcing is identified: single or multiple sourcing. This is the first differentiation in the supplier selection literature (Figure 2).

3.1.1: single sourcing

Single sourcing, where one supplier fills the entire demand for a product, or a package of products, is currently the area in which the largest part of research is conducted (De Boer et al., 2001). The documented benefits of single sourcing are quantity discounts from order consolidation, reduced order lead times, and logistical cost reductions as a result of a scaled down supplier base (Burke et al., 2007). However, it is necessary that the supplier can fully meet the buyer's requests in terms of quantity, quality, or delivery. Single sourcing concerns the identification of the 'best' supplier (Assaoui et al., 2007). Therefore, single sourcing problems are called vendor selection problems. If only one supplier is available, selection is impossible. This situation is called sole sourcing.

Vendor selection problems are usually solved under multiple criteria. Techniques to solve these kinds of problems deal with issues such as weighing criteria and scaling the performances of suppliers. For Heineken, single sourcing is not an option, since the capacity of the suppliers does not allow the sourcing of all bottles at one location.

3.1.2: multiple sourcing

Single sourcing can be dangerous from the perspective of resilience to supplier disruptions (Christopher & Peck, 2004). Therefore, companies might consider multiple sourcing. In case of a disruption at a supplier, other suppliers can continue delivery. Multiple sourcing can thus decrease the dependency on a single supplier. If a supplier's capacity is insufficient to meet peak demand, having additional suppliers can avert the need to build seasonal stock (Jayaraman et al., 1999). Multiple sourcing also motivates suppliers to be price and quality competitive. Especially when switching between suppliers is costly and capacity is limitedly available, threats of switching suppliers do not work as efficiently as multiple sourcing (Cachon & Zhang, 2003). Under multiple sourcing, the problem is twofold: vendor selection and order quantity allocation (Assaoui et al., 2007). Therefore, these problems are generally called allocation problems. The problem at hand in this research is an allocation problem.

Within the area of allocation problems, a difference is made between single and multiple objective problems. Unlike in single sourcing problems, the use of multiple objectives is not that common in multiple sourcing problems (Ignizio & Romero, 2003). Regardless of the number of objectives, allocation problems are most often solved by mathematical programming. Linear programming, mixed-integer programming, multi-criteria programming, and goal programming are all different types of mathematical programming.

Another distinction can be made between single-period or multi-period problems (Assaoui et al., 2007). In case of multi-period problems, decisions have to be made on stock management. In this area, multi-period inventory models have been developed. The basic idea is to find a balance between the fixed costs involved in placing an order, the costs of carrying inventory, and the expected costs of stockouts. For Heineken, the allocation problem can best be solved as a single-period problem.

3.1.3: single objective models

Single objective programming models usually minimize the purchasing costs. Other criteria are often incorporated as well in the form of constraints (Ignizio & Romero, 2003). These constraints, for instance, indicate the minimal quality of the product, the maximum capacity of a supplier, or the minimum number of suppliers. The problem characteristics define the type of programming needed to model the problem. The methods to solve linear programming problems are much more efficient than those for mixed-integer programming problems (Aimms, 2007). As the size of the problem grows, the differences in solving time can rapidly grow larger. Therefore, it is better to use linear programming models where possible.

Other solutions to the single objective problem can be found by methods as dynamic programming, non-linear programming, stochastic programming, or decision theory (Assaoui et al., 2007). These methods are less used than the linear and mixed integer programming models and have the same disadvantage: only one objective can be used, whereas many real-world decision problems involve multiple, conflicting objectives (Ignizio & Romero, 2003). The use of constraints to model these objectives ignores that not all the objectives carry the same weight.

3.1.4: multi-objective models

Multiple objective models are designed to let the decision maker give weights to all objectives. There are two main types of models in this category: multiple objective programming and goal programming models. In multi-objective programming, all objectives are incorporated in the objective function, each with its own weight. In goal programming, not the values of the objectives themselves, but the deviation from the goal value of that objective is used in the objective function. The choice for a method depends on whether it is possible to set goals for all objectives and whether all objectives can be compared to each other directly.

With one goal, there is only one optimal solution value. Multiple objectives, however, give rise to a set of compromise solutions. These solutions are referred to as pareto-optimal. In absence of preference information, none of the trade-offs can be found better than others (Zitzler, 1999). A decision maker has to give preferential weights to the different objectives. This can be done either before or after optimizing. When there is no information on preferences, Chebyshev goal programming can be used (Flavell, 1976). In this method, the maximum deterioration of all objectives is minimized. Next to preferential weights, normalization weights are often necessary too. If both objectives represent different units (Euros and number of bottles, for instance), simple addition of the values would be comparing apples with oranges (Gass, 1986).

Tamiz et al. (1998) give a summary of different types of normalization factors, all with their own pros and cons. The methods differ on stability, with respect to for instance the acceptance of goal values of zero, computation time, and meaning of the final goal function. Zero-one normalization gives most meaning to the goal function, since the value of the objective will be compared to its best and its worst value. The method takes more time than most other methods. The best and worst values have to be found first for all objectives. Therefore, it is best to use this method only for small numbers of objectives with clearly defined minimum and maximum values.

3.1.5: discounts

Additional features can be added to most of the models. Discounts are the most important feature found in the literature. Discounts can become available with scale, but also with special bundles (Lawless, 1991). Especially in environments were multiple products are procured at the same time, special combinations of products can give advantages in terms of price discounts or free products. These deals make the models far more complicated, since ordering more products does not necessarily mean higher costs anymore.

3.2: security of supply

Security of supply is a definition used within Heineken. The literature focuses more on the opposite of security of supply: supply risk. Decreasing supply risk is the same as increasing security of supply. Therefore, we discuss the topic of supply risk instead of security of supply.

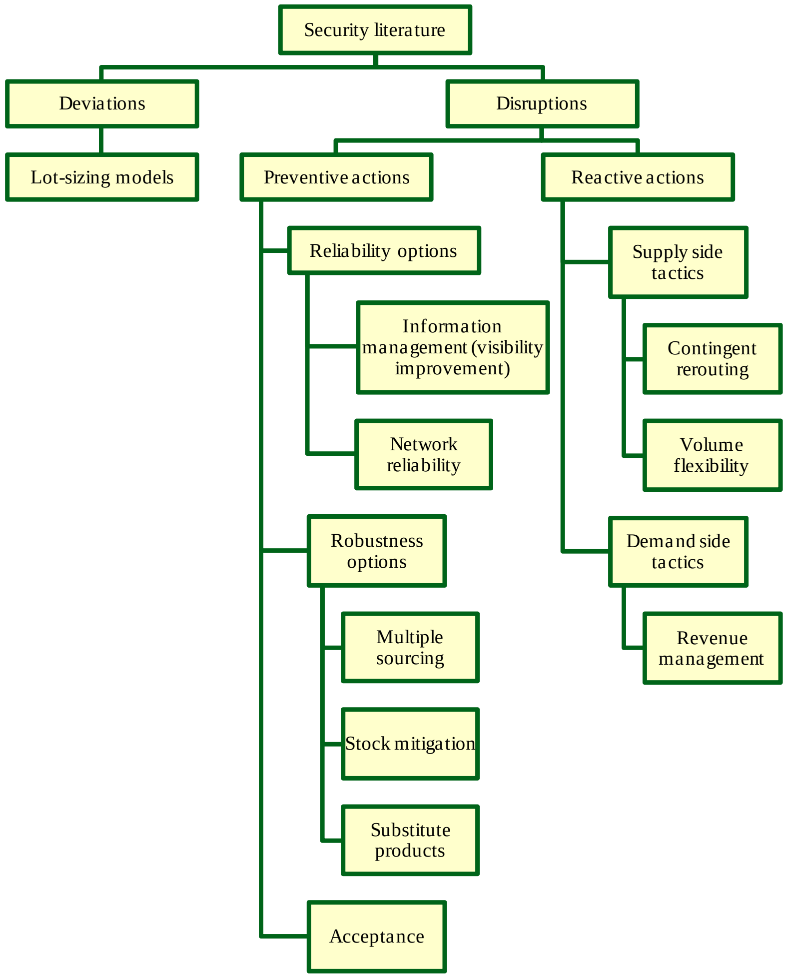

Figure 3: security of supply literature

Supply risk can manifest itself on different levels. Section 3.2.1 elaborates on the differences between these levels. Then, Section 3.2.2 discusses the impact of supply risk. Finally, Section 3.2.3 deals with the strategies to increase the security of supply.

3.2.1: supply risk levels

There are three levels of supply risk: deviations, disruptions, and disasters (Figure 3). Deviations occur on the operational level, when actual delivery is not according to expectations. Two types of deviations can be seen: yield uncertainty, when the quantity of delivered goods is stochastic, and lead-time uncertainty, when the arrival time is subject to changes from the expected time (Snyder & Shen, 2006). The existence of these deviations creates risk of stockouts. Operational ordering systems can decrease the probabilities of a stockout.

Disruptions are problems that occur upstream in the supply chain and lead to a complete stop of supply. We define disruptions to be supplier specific, which means that a disruption occurs at only one supplier at the same time. Examples of disruptions are major problems in the production process of the supplier, supplier insolvency, and transportation problems (Treleven & Bergman Schweikhart, 1988). It also includes the risk that the supplier will terminate the relationship, voluntarily cutting off the supply. Whatever the reason might be, the consequence is that no supply is arriving at all. Although disruptions occur less frequently than deviations do, they are in general more severe (Snyder & Shen, 2006).

Disasters are the most severe type of risk. Disasters can be viewed as large-scale disruptions. They do not affect a single supplier, but entire areas, supply chains, or networks. For instance, in the aftermath of 9/11, border closures caused all international transportation to the United States to be disrupted (Sheffi & Rice, 2005). Due to the size of the problems, their unpredictability, and the mitigating costs, it is not worthwhile to protect against disasters (Gaonkar & Viswanadham, 2003).

In this research, we focus on disruptions. Research in the area of supply disruptions has seriously increased after 9/11 (Paulsson, 2003). One of the main conclusions is that once a disruption occurs, there is not much that can be done regarding the supply chain infrastructure, as strategic decisions cannot be changed quickly (Gaonkar & Viswanadham, 2003). Therefore, it is critical to account for disruptions during the design of supply chain networks (Snyder et al., 2006).

3.2.2: impact of supply risk

Low quality basic material use or bad process control and management are likely to cause disruptions. Therefore, disruptions based on these causes are well predictable. Some disruptions are less likely to occur. Because of the limited number of data points, good estimates of the probability of the occurrence of a disruption and accurate measurements of its impact are difficult to obtain (Tang, 2006). With inaccurate estimates, many firms find it difficult to perform analyses to justify risk reduction programs. Moreover, some disruptions turn out to be unavoidable, at least against reasonable costs (Sheffi et al., 2003). Managers therefore shift the focus from the causes to the effects. Especially when the time between successive disruptions becomes larger, the impact of the disruption becomes more important than its probability.

Some of the recent business trends increase the impact of supply disruptions. (Stecke & Kumar, 2006). Actions that are initiated because of their cost reduction potential can decrease the security of supply at the same time. Globalization, for instance, causes many products to be transported over long distances and with multiple modes of transportation. Not only does this increase the probability of deviations and disruptions, it also increases the response time to problems. Since transportation over long distances is possible, large specialized factories that produce only a few products are created to benefit from economies of scale. Combined with a trend to reduce the overall number of suppliers, this decreases redundancy in production facilities. Next, outsourcing and complex network relationships have decreased the visibility within the supply chain and increased the need for high-quality communication. These factors therefore increase the probability of deviations or disruptions. Another example is the reduction of the throughput time of products in the supply chain. Just-in-time delivery programs, aimed to reduce stock levels, and the reduction of slack time in the process make it impossible to contain disruptions locally.

3.2.3: improvement of security of supply

In general, it holds that improvements to the security of supply cost money. Acceptance of the disruption risk may be the appropriate action in some cases (Tomlin, 2006). Acceptance is not the same as ignoring risks or doing nothing. When a risk is accepted, this risk has been identified and the possibilities to decrease its impact have not been considered worthwhile. Acceptance is thus a deliberate action.

Security improving actions with regard to disruptions can be categorized into two sections: preventive (mitigating) actions and reactive (contingency) actions. The first section can be further divided into measures that decrease the probability of disruptions and measures that decrease its impact. These measures are often conflicting. For instance, the use of multiple suppliers decreases the impact of a disruption, but increases its likelihood.

Multiple sourcing is a type of redundancy. Creating redundancy is the most important robustness action (Sheffi et al., 2003). It concerns the duplication of resources to ensure the availability of a backup solution in case of disruption (Tang, 2006). The use of stock is redundancy of a product type (Schmitt & Snyder, 2007). The use of substitute products is redundancy of a product range. Another way to increase the redundancy in the supply chain is by postponement. Standardizing products or product parts often creates more production locations, supply sources or available stock.

Reliability of Heineken's suppliers is not an issue. Under normal circumstances, quality measures ensure reliable suppliers. For other companies, reliability could be more important than low impact. Reliability can be improved by making the supply network more stable or by improving the flow of information throughout the supply chain. The concept of supply chain reliability is related to network reliability theory. This theory is concerned with maximizing the probability that a graph remains connected after random failures (Gaonkar & Viswanadham, 2003). Difficulties with the use of this theory arise as it is only concerned with on/off systems. In these systems redundant suppliers are expected to be fully able to compensate the loss of other suppliers. The capacity restrictions and lack of volume flexibility at bottle factories prevent this compensation.

Compensation of lost supply requires flexibility and can be used as a reactive action. Suppliers have to be able to (temporarily) increase their processing capacity (Tomlin, 2006). Other options are swift changes in production locations or the use of flexible transportation (Tang, 2006). These actions are contingent rerouting actions, because after the disruption has occurred, the supply chain is quickly rerouted. Firms do not need to rely exclusively on supply-side tactics during a disruption. The ability to use demand-management capabilities to shift demand to alternative products that are less supply constrained can seriously decrease the impact of the disruption (Tomlin, 2006).

3.3: incorporation of security measures in allocation models

Not many allocation models explicitly incorporate security of supply. Some of the models that do incorporate risk are not applicable for Heineken. These models define a trade-off between a higher-cost, reliable supplier and a lower-cost, unreliable supplier (Tang, 2006b). The reliable supplier is assumed to have no disruptions at all, and both suppliers have infinite capacity. However, in Section 3.2 we have established that disruptions cannot be completely ruled out. Capacity of the glass bottle factories is finite as well. This dual sourcing problem is therefore not realistic in most cases.

Since the allocation is done before the disruptions occur, only preventive options are incorporated into the models found in the literature. Snyder et al. (2006) use reliabilities to investigate multiple scenarios of disrupted systems. Based on the probabilities of disruptions in these scenarios and the additional costs of sourcing from a more distant location in case of a disruption, the authors calculate the expected costs of the designed system. For risk-averse decision makers they propose a worst-case model. In both models, infinite capacity is assumed. Another problem is that the number of scenarios grows quickly, making the model unusable for larger problems.

Just as Snyder et al., many authors use fixed values for the reliability of the system or the quality of the supplier. Constraints are then used to create a minimal reliability (in Bundschuh et al., 2003) or a minimal quality solution. The impact of these constraints can be analyzed by comparing the solution to that of a base model, without the constraint. The use of constraints for goals ignores the fact that the decision maker might have different preference levels for the goals.

With respect to the robustness of the solution, Bundschuh et al. present two models as well. In the first model, multiple suppliers are chosen by setting a sourcing limit to the demand allocated to a supplier. In the other model, additional stock can be used to decrease the impact of the disruption. This impact is bounded by a constraint indicating the maximum tolerable loss. The objective is to minimize the costs of both stock and purchasing costs. Bundschuh et al. propose using options to stock, instead of real stock, to decrease the costs. Because of the unique bottles of Heineken, no such options are available in the market.

All of the models discussed above try to minimize the costs. Sometimes, the expected costs of supply disruptions were incorporated in these costs. The robustness model of Gaonkan & Viswanadham (2003) minimizes the expected shortfall from the expected quantity, though. The expected shortfall is combined with the fixed costs of using a supplier. These values are calculated for multiple scenarios of supplier disruptions.

3.4: conclusions

There is no solution in the literature that exactly fits the problem at Heineken. Moreover, very few allocation problems incorporate security of supply measures. We therefore have to design a model ourselves to solve the problem at Heineken. We will do this based on the best practices found in the literature.

The allocation problem is usually solved with mathematical programming solutions. Depending on the number of objectives, single or multi-objective programming is better. At Heineken, there are two important objectives: minimizing costs and maximizing security of supply. With two objectives, both programming types can give satisfactory results. Since only one objective would have to be modelled as a constraint, no preference problems between constraints would arise. However, a goal programming approach gives the most insight into the results by comparisons to the goal values.

Most models calculate the security of supply as the expected costs of disruptions. Gaonkan & Viswanadham (2003) and Bundschuh et al. (2003) use the number of lost products under different risk scenarios to quantify the security of supply. Snyder et al. (2006) suggest minimization of the worst-case scenario.

To improve the security of supply, preventive robustness measures have to be taken. Multiple suppliers can be used for this purpose. The use of a sourcing limit (Bundschuh et al., 2003) suggests that the number of products per factory has influence on the security as well. Stock is a very important measure against supply risks.

Chapter 4: model formulation and results

The previous chapters described the current allocation processes at Heineken and solutions to allocation problems found in the literature. This chapter discusses a solution to the allocation problem at Heineken. An objective of this research is to improve the insight into the quality of the solution. This can be accomplished by the use of a goal programming approach (Figure 4). In this approach, the allocations are characterized by their values on the objectives costs and security of supply. Every solution is a trade-off between optimizing for costs and optimizing for security of supply. The boundaries to the solution range are the best possible cost and the best possible security of supply solutions. In Figure 4, these are indicated by a square and diamond respectively. An allocation is evaluated on the deviation from the goal values relative to the objectives' ranges. The deviation gives immediate clarity about the quality of the solution on an objective.

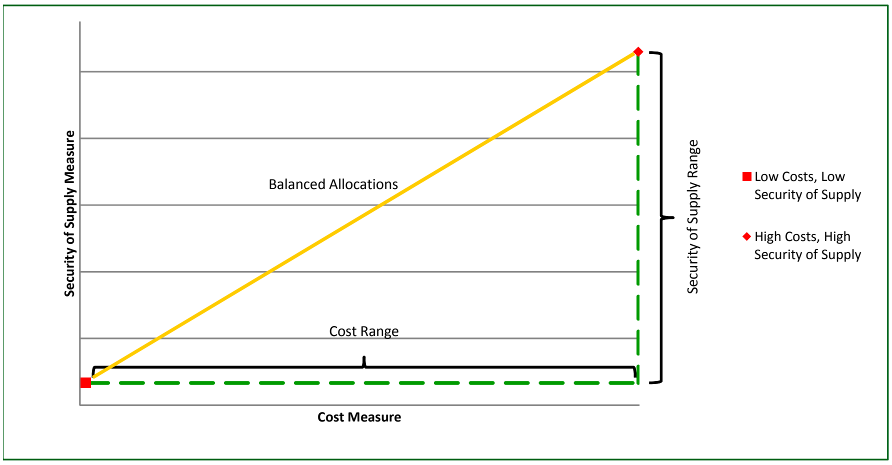

Figure 4: basic elements of the goal programming model

Before we formulate the goal programming model, the security of supply measure and the boundaries of the solution range have to be found. Section 4.1 defines how security of supply can be measured. The boundaries to the solution range can be found by Mixed-Integer Linear Programming models. Section 4.2 formulates these models. Section 4.3 discusses the goal programming model.

4.1: measure of security of supply

The best way for Heineken to measure its security of supply, as concluded after the literature review in Chapter 3, is by calculating the number of bottles the suppliers cannot deliver after a disruption of supply. In this calculation, we try to estimate the disruption mitigation capabilities of an allocation. There are different methods to measure the number of lost products. In this section, we choose an appropriate measurement method for security of supply.

Section 4.1.1 describes the basic formulation used throughout this chapter. It also discusses how redundancy can mitigate a disruption. Finally, it checks whether the formulation is a valid representation of the real problem. Next, Section 4.1.2 discusses the requirements of a good measurement method. Section

4.1.3 evaluates different measurement methods with respect to these requirements. Although in practice the effects of a disruption often extend to multiple SKUs, these sections consider every SKU separately. However, we will show that the interaction between the SKUs is already taken into account in our measurement.

4.1.1: basic disruption formulation

A disruption can have multiple causes that we call problem types. Examples of problem types are furnace leakages, fires in the factory, and machine breakdowns. Every problem type has its duration, which is the expected production time that is lost because of the disruption. Figure 5 visualizes the relations between the different disruption elements used in this chapter. As shown by the bracket, the duration is the number of periods of production that is affected. Every period, an average amount of supply is scheduled to arrive, shown by the two-sided arrows. Multiplying the number of periods (the duration) with the average supply per period leaves the number of products affected by the problem type. This represents the impact of the disruption. The impact and all other elements indicated by one-sided arrows in Figure 5 are counted in number of items.

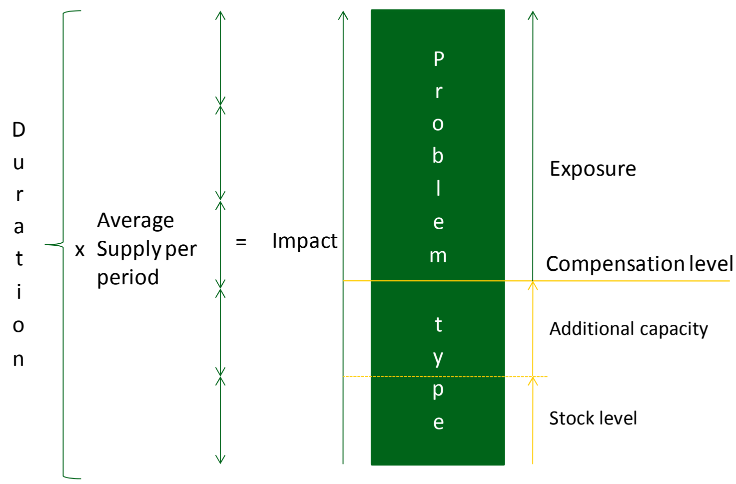

Figure 5: relation between the disruption elements

Heineken might not notice the impact of the disruption. Additional supply sources, available stock, or temporary additional production capacity at other factories can compensate for the loss of production. The additional supply sources are not visible in Figure 5, but influence the impact of a disruption by decreasing the average supply per period of each factory. Exposure is the expected number of products Heineken is actually missing (impact minus the compensation level). Both the impact and the exposure can never be smaller than zero.

We assume the stock level to be constant and independent of the number of factories used to deliver a bottle. Additional production capacity does change with the number of factories used. Suppliers need to have a little flexibility to cope with production losses and set-up times. However, we assume that the factories do not have enough volume flexibility to significantly increase production. The high costs of overcapacity motivate them to minimize free capacity. Therefore, additional capacity for missing SKUs has to be compensated by lowered production for other SKUs. Lowering production of a bottle should not result in a stockout for that bottle (Section 4.2.2 discusses this further). For this reason, it is only accepted if there is enough stock for these SKUs. Since the available additional capacity is dependent on the allocation, the formulation of Figure 5 captures the interdependency of the different SKUs.

Stock is an important provider of redundancy, directly through the stock level and indirectly through the additional capacity. Whether it is best to use the impact or the exposure to calculate the number of lost products is therefore dependent on the availability of stock. At least three factors influence the availability of stock. First, the time between failures: if there is much time between two disruptions, the stock is likely to be rebuilt before a new disruption occurs. Second, the volume flexibility: if the suppliers have flexibility, capacity can be increased to quickly rebuild the stock. Finally, whether the bottles are back-ordered or not: if Heineken's customers do not postpone their purchases until the bottle is available again, less time is needed to recover. There is not much flexibility at Heineken's suppliers. However, Heineken assumes that a large part of demand will not be back-ordered. Moreover, the reliability of the suppliers is high enough to expect a large time between failures. Therefore, we assume that stock is always present if a disruption occurs. We use exposure to indicate the severity of a disruption to Heineken.

4.1.2: measurement requirements

Every SKU at a factory has its own exposure. The security of supply for a SKU can be calculated by combining its individual factory exposures. This is often done by taking the total exposure, the average exposure (Beamon, 1999), or the maximum exposure of a SKU (as suggested by Snyder et al. (2006)). Consider a case with 3 factories delivering a bottle to Heineken. Each factory has a different impact, but all have the same compensation level. Figure 6 shows the security of supply value under the three measurement methods. After this section has described the requirements of a good measurement method, these methods will be evaluated in Section 4.1.3.

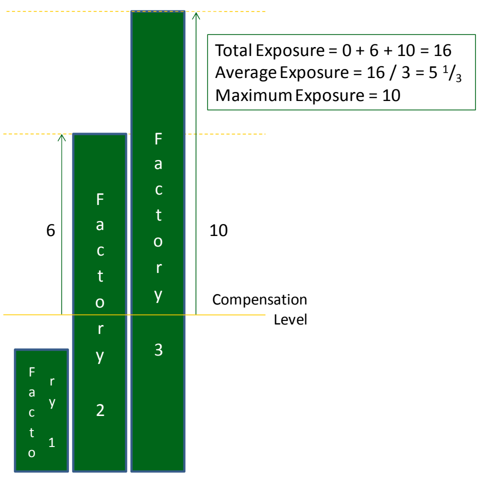

Figure 6: security of supply measurement methods

The measurement method must be aligned with the company's strategy in order to be accepted (Neely et al., 2000). Improvements in the security of supply for Heineken, should be reflected in an improvement of the value of the measurement method. The allocations differ on two points: the number of factories and the number of products allocated to each factory. The distribution of this number of products is expressed by the variance of the exposure. Figure 7 shows how this variance is calculated by the summed squares of the deviation from the average exposure.

The allocation becomes better when there are more factories. The first reason is that the supply is split between multiple factories. As the average supply per factory per period becomes smaller, the impact becomes smaller with a disruption at one of the factories. Moreover, there is a larger probability that one of the suppliers is able to increase production of the missing bottle on short notice.

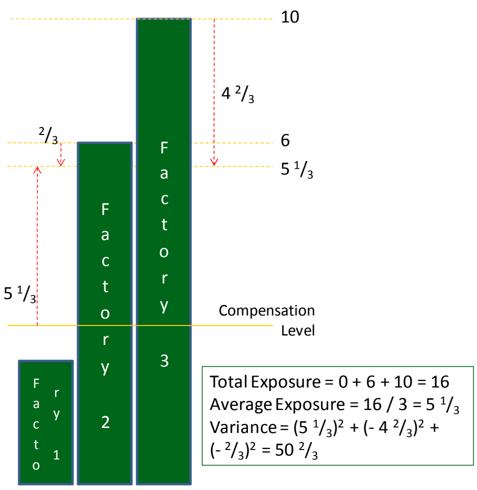

Figure 7: variance of the allocation

The variance of the allocation decreases when the deviation from the average exposure is smaller. If the total exposure is evenly spread between all suppliers of a bottle, the variance is lowest (zero). However, not all factories necessarily have the same amount of bottles allocated in this case. The additional capacity is different for each factory. An equal spread of the risks minimizes the worst-case exposure and is therefore preferred.

The last requirement is the ease of use and understanding. It should be clear what the value of a measurement unit means. The value of the cost measure is clear, but exposure is a more difficult measure. Still, the measure of security of supply and changes to its value should be apprehensible.

There are three important factors (number of factories, variance of exposure, and ease of use) and one of minor importance (individual bottle preferences) on which we will evaluate our methods.

4.1.3: measurement methods

This section discusses the three measurement methods: total exposure, average exposure, and maximum exposure. It also examines combinations of methods. Both methods chosen for this combination meet different requirements. We evaluate which method is the best measure of security of supply for Heineken.

Total exposure
First the method of total exposure will be discussed. This method decreases as demand is allocated to a new factory up to the point where the compensation level equals the impact. The two exposure examples on the left side in Figure 8 show this. Demand is re-allocated from factory 3 to factory 1, decreasing total exposure. In addition, the compensation level becomes larger given a higher number of factories, because there are more opportunities to increase production at other factories. Therefore, the total exposure decreases if the number of factories increases. Changes in the allocated quantity above the compensation level (the right two situations in Figure 8) show that total exposure is indifferent to the distribution of exposure between the factories. The total exposure measurement method does not minimize the variance.

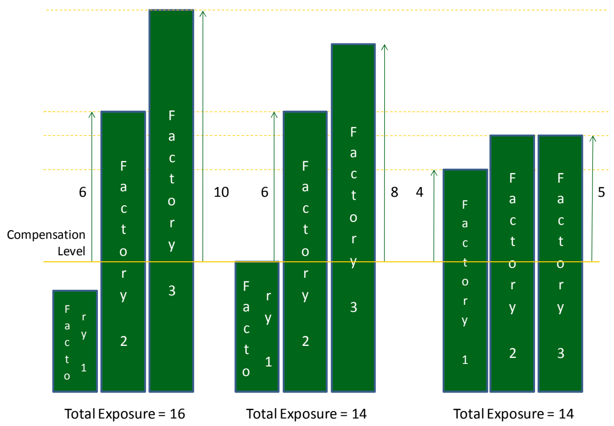

Figure 8: total exposure measurement

Average exposure
Average exposure is based on total exposure and behaves more or less in the same way. It is not affected by a change in the distribution of either, thus failing to minimize the variance. The average exposure responds more intensely to an increase in the number of factories than the total exposure. It is found by dividing the total exposure by the number of factories. An increase in the number of factories therefore gives a strong decrease of the average exposure value. This decrease is strongest for the first factory that is added, and becomes smaller as more factories are used. The average exposure probably overestimates the benefit of an extra factory as using multiple factories already raises the compensation level. Another characteristic however is useful. As the number of factories becomes higher, the benefits of an additional factory becomes lower. This is good, because using a large number of factories makes coordination difficult for the OpCos.

Average exposure improvements become larger with demand, which Figure 9 shows. There are two situations (1 factory on the left side, 2 factories on the right) and two SKUs. For simplicity, we assumed that the compensation level does not increase if two factories are used. SKU 2 has twice the demand of SKU 1. Going from one to two factories decreases the average exposure with approximately 50%. In this case, this causes the net improvement for the large-demand SKU (2) to be twice as large as the improvement for SKU 1.

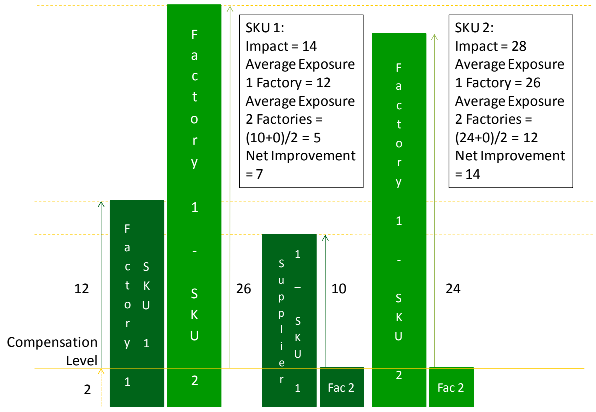

Figure 9: average exposure measurement

Maximum exposure

The maximum exposure is only influenced by allocation changes to the factory with the highest exposure. Reallocating supply from this factory to others improves its value. The maximum exposure therefore hardly improves if factories are added without allocating substantial supply to those factories. Its optimal value is reached if the variance of exposure is zero. Resulting in all factories having the same exposure.

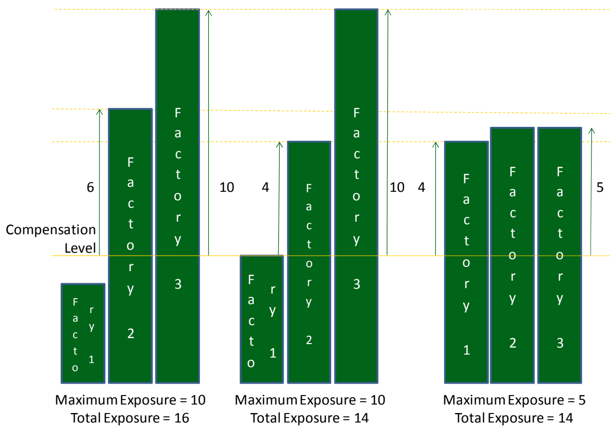

Figure 10: maximum exposure measurement

Allocation changes to other factories than the one with the highest exposure do not influence the maximum exposure value. With more than two used factories, changes are possible that do not affect the factory with the highest exposure. Changes can therefore remain unnoticed to the measurement method. This can cause the total (and average) exposure to become larger than necessary. The left two examples in Figure 10 show this. The model is indifferent between the two situations under maximum exposure measurement, while the second situation is preferred under the other measurement methods. Thanks to the relatively small number of factories (three is in practice the maximum) this is an issue of less importance. Moreover, the optimal allocation under maximum exposure does minimize the variance. The two examples on the right side in Figure 10 show that the optimal maximum exposure also minimizes the total exposure.

Combination of number of factories and variance

Other methods for calculating the security of supply can be used as well. The number of factories per bottle is an option, the variance another. Individually, these methods are not very powerful. They only focus on a single requirement. Combined, these methods might provide a good solution. However, combining measurement methods creates other difficulties. The ease of use of such a method is not high. The number of suppliers and the variance cannot be compared directly. Therefore, a weight is needed to combine the two values. This makes the combined value not well understandable.

Evaluation

Table 1 shows the scores of the four methods on the different requirements. We conclude from the discussion above that the maximum exposure performs best. The only requirement it does not meet, was already considered to be of minor importance. Another benefit of this method is that it reviews the worst-case scenario for a SKU. Therefore, we decide to use this value as a measurement method for the security of supply.

Method Evaluation				
Method	Easy to use	Variance	Nr. suppliers	Demand-based
Total Exposure	Yes	No	Yes	No
Average Exposure	Yes	No	Yes	Yes
Maximum Exposure	Yes	Yes	Average	No
Combination of number & variance	No	Yes	Yes	No

Table 1: measurement method evaluation

Definition:

The maximum exposure of a certain bottle type is the number of bottles that cannot be delivered to a brewery under a realistic worst-case scenario of a disruption at any glass bottle factory.

If the factory has to be rebuilt completely, a disruption can last for multiple months. However, the chance of such a disruption is negligible. Snyder et al. (2006) indicate that it does not make sense to consider all possible scenarios in a worst-case calculation. Therefore, we assume the maximum exposure under a realistic worst-case scenario. We evaluate different problem types and find the worst-case scenario that is worthwhile to protect Heineken against.

The security of supply is found by a summation of all products' maximum exposures. This means that the security of supply also can be defined as a number of bottles that might possibly not be delivered. We want to minimize this number, contrary to the feeling that security of supply should be maximized.

4.2: measurement goals

The quality of an allocation can only be estimated by comparison to references values for both costs and security of supply. These reference values are found by the lowest costs and the highest security of supply scenarios. Section 4.2.1 describes the base model that is used to calculate the cost scenario. Section 4.2.2 extends this model to calculate the highest security of supply allocation. These models can be solved by

Mixed Integer Linear Programming techniques. We have implemented the models using AIMMS modelling software version 3.7 with CPLEX 10.1.

4.2.1: cost base model

The basic allocation problem can be formulated as a network flow problem (Paragon, 2008). This specific situation requires alterations, creating the need for a different structure. The structure consist of multiple origin and destination points with products, in this case beer bottles, transported between them. The question is which product flows to create to fulfil the demand at the destination at minimal costs. Products (indexed as p=1,..,P) flow from supplier factories (f=1,..,F) to breweries (b=1,..,B). The products can have different colors (c=1,..,C), creating subsets (p_c) of products with color *c*, and the suppliers (s=1,..,S) can own multiple factories, creating subsets (f_s) of factories belonging to supplier *s*.

A furnace can only operate one color at a time. Every factory *f* has available capacity ($cap_{f.c}$) to produce products of a specific color *c*. Supplier *s* has a capacity as well ($cap_{s,c}$) for color *c*. This supplier capacity is equal or lower than the sum of the individual factory capacities (Figure 11). The reason is that suppliers are sometimes flexible in their other orders. This means that they can produce these orders at multiple factories, creating flexibility. For example, a supplier has three factories with capacity for 100 tonnes of glass. For another customer, the supplier has to reserve 50 tonnes of capacity, but the products of this customer can be made in any of the three factories. This means that the individual factory capacities are 100 and the supplier capacity is 3*100 – 50 = 250. The capacity is expressed in tonnage, because the output speed of the furnaces is determined by the tonnage of glass that can be melted. We use the weight (w_p) of bottle *p* to link the allocated volume to the capacity.

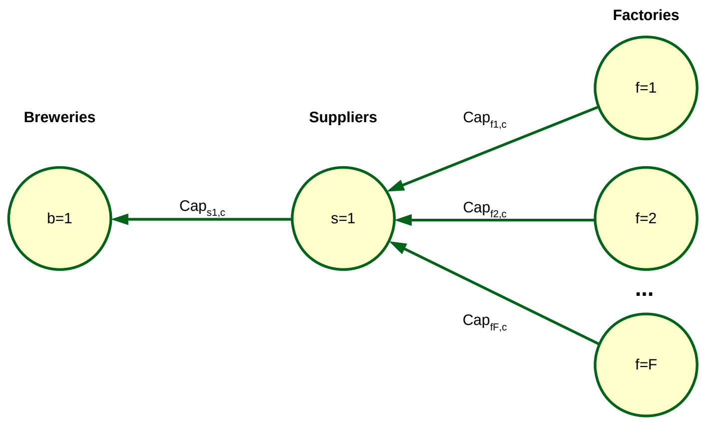

Figure 11: the capacity of the supplier limits the sum of the factory capacities

Brewery *b* of Heineken has demand ($d_{p,b}$) for bottle *p*. The costs ($pr_{p,f,b}$) associated with delivering bottle *p* from factory *f* to brewery *b* include procurement costs, transportation costs, and a payment term discount based on the number of days until payment for that bottle is due. The model decides on the number of bottles ($X_{p,f,b}$) of product *p* that need to be shipped from factory *f* to each brewery *b*. Allocation can only occur if there is a price agreed between the supplier and Heineken. This agreement is indicated by a positive value of the price.

Not all allocation sizes are allowed, because small batches cause high set-up time to production time ratios. Therefore, if there is an allocation ($Y_{p,f,b}$=1; $Y_{p,f,b}$=0 if there is no allocation) of product *p* to factory *f*, destined for brewery *b*, the allocation has to be higher than a minimum value (minv) or a minimal percentage of demand (minp), whichever is less. The objective is to minimize the costs, while guaranteeing enough products to fulfil demand and not exceeding the capacities.

This leads to the following mathematical model formulation:

Indices

P	Products 1,..,P
B	Breweries 1,..,B
F	Factories 1,..,F
S	Suppliers 1,..,S
C	Colors 1,..,C

Parameters

$pr_{p,f,b}$	Price to deliver product p from factory f to brewery b (in Euros)
$d_{p,b}$	Yearly demand for product p at brewery b (in # of products)
w_p	Weight of product p (in kg)
$cap_{f,c}$	Capacity of factory f for color c (in kg)
$cap_{s,c}$	Capacity of supplier s for color c (in kg)
Minv	Minimal value that has to be allocated
Minp	Minimal percentage of demand that has to be allocated
p_c	Products with color c
f_s	Factories of supplier s

Variables

$X_{p,f,b}$	The total number of products p allocated to factory f and destined for brewery b in a year (in # of products)
$Y_{p,f,b}$	Binary assignment variable: 1 if product p is allocated to factory f and destined for brewery b

Constraints

No more allocation is allowed than there is capacity for factory f or supplier s:

$\sum_{p \in p_c, b} X_{p,f,b} * w_p \leq cap_{f,c}$ $\forall f,c$ (1)

$\sum_{p \in p_c, f \in f_s, b} X_{p,f,b} * w_p \leq cap_{s,c}$ $\forall s,c$ (2)

There must be at least enough allocation at all factories to fulfil the demand:

$\sum_f X_{p,f,b} \geq d_{p,b}$ $\forall p,b$ (3)

This constraint causes the binary assignment variable ($Y_{p,f,b}$) to be 1 if there is allocation:

$X_{p,f,b} \leq d_{p,b} * Y_{p,f,b}$ $\forall p,f,b$ (4)

If there is allocation, the allocation size should be at least a percentage of demand or a fixed number of bottles, whichever is less:

$X_{p,f,b} \geq \min(mnv, mnp * d_{p,b}) * Y_{p,f,b}$ $\forall p,f,b$ (5)

No products can be allocated if there is no price agreement between Heineken and the supplier:

$X_{p,f,b} = 0$ $\forall p,f,b \mid pr_{p,f,b}=0$ (6)

No negative amounts of products can be allocated to the suppliers:

$X_{p,f,b} \geq 0$ $\forall p,f,b \mid pr_{p,f,b} \neq 0$ (7)

The assignment variable is binary:

$Y_{p,f,b} \in \mathbb{B}$ $\forall p,f,b$ (8)

Objective function, minimize the total costs of the allocation:

$min \sum_{p,f,b}(X_{p,f,b} * pr_{p,f,b})$ (9)

4.2.2: security of supply base model

An extension of the model of Section 4.2.1 (without objective function (9)) can be used to solve the allocation problem with maximum security of supply. Maximizing security of supply means minimizing the maximum exposure. In order to do this, the model requires two additions. First, the compensation level for each product and factory needs to be calculated. In Section 4.1 we defined this level as the sum of the safety stock level ($st_{p,b}$) of product *p* at brewer *b* and the additional gained capacity. The additional capacity represents the possibility to increase production of a missing bottle at other factories, under the scenario that there is a disruption at a certain factory. Since no supplier can be expected to have spare capacity on the short term, this capacity has to be freed. Production of bottles has to be lowered to compensate for an increase. Therefore, we generalize the concept of gained capacity to production size changes ($PC_{g,p,f,b}$), which can be both positive (in case of gained capacity) and negative. It represents the change in production of bottle *p* at factory *f* destined for brewery *b* after a disruption at factory *g.* The maximum exposure is examined after disruption scenarios at all factories.

It should also become possible to calculate the maximum exposure ($ME_{p,b}$) of product *p* for brewery *b*, based on the allocation over the problem duration (pd) and the compensation level. Many restrictions to this model have to be considered. Some of these restrictions are required for the model to make sense. Others are based on our assumptions and are meant to keep the model simple and fast, and the results realistic in practice:

1) No production size changes can occur at the factory that has a disruption. Although it might be possible that the factory is still partly functional after a disruption, this is very dependent on the problem type at hand. Moreover, this would only be of help if at least one line or furnace of the same color would still be operational, which is not only dependent on the problem type, but on the factory layout as well. Because this becomes very complex and many different scenarios might exist, we assume that the entire factory is unavailable after a disruption. This is the worst–case scenario.
2) No changes can occur for products that are not already produced at a factory. It takes too much time to make moulds for a new product. Therefore, we assume that only factories that already produce the bottle can increase production. The other way around, the production size of a bottle obviously cannot be decreased if that bottle is not produced at a factory (this is further specified at point 4).
3) Production of a bottle cannot be decreased by more than the amount of bottles that are in stock for that bottle type. In other words, it is not allowed to cause shortages by decreasing production. In order to understand this assumption, first, suppose that it would be allowed to cause a shortage for other bottle types. This would give two options. Either the shortage could be solved by a production increase at yet another factory or the shortage is not compensated. The first scenario is reasonable in theory, but would be very complex in practice in terms of coordination and cost control. The second scenario would result in a trade-off between the loss of two bottle types. Without information about preferences of both bottles, this decision becomes useless.

Therefore, it is best to keep the solution simple and not allow the decrease to be larger than the stock level. This limitation is shown in Figure 12.

4) Production of a bottle cannot be decreased by more than the size of the production over the problem duration. This is also shown in Figure 12. The supplier reserves no more capacity during this period.
5) The tonnage of freed capacity at each factory is equal to the tonnage of increased production. The output of the furnace is the bottleneck for capacity and therefore no additional tonnage can be produced of a specific color. Figure 12 shows that the increase of production is limited by the freed capacity.
6) Only for missing bottles, needed because of the disruption, it is allowed to increase production. Moreover, it is not allowed to decrease production of missing products. In Figure 12, changes for missing bottles cannot be negative, whereas changes for other bottles cannot be positive. We make these assumptions to keep the model simple and the coordination requirements relatively low.

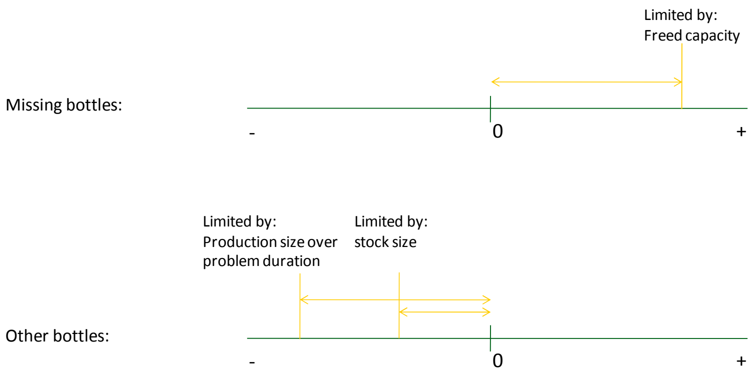

Figure 12: production changes and its limitations

Based on these restrictions the model calculates the compensation levels. Then the exposure per factory and product can be calculated as the difference between the impact (allocated volume times the problem duration) and the compensation level. The maximum exposure is the largest of the product's exposures per factory. The model in this section minimizes the sum of all products' maximum exposures.

We added the restriction that only missing products can have a positive production change. Consequently, the only factories that need to be considered are those that produce at least one product that the factory with a disruption produces as well. In order to keep the problem small and efficient, we therefore introduce a subset (f_g) of factories that have at least one product in common with the factory *g* that has a disruption. The model is only used on this subset.

It is impossible to increase the production of a bottle beyond the available capacity. In order to identify the capacity that a factory has for a bottle of color *c* we add the parameter c_p to indicate what color product *p* has.

This leads to the following mathematical model formulation. We only identify the changes from the previous model:

Parameters

c_p	Color c belonging to product p
Pd	Problem duration (in years)
$St_{p,b}$	Stock for product p at brewery b (in # of products)
f_g	Factories 1,..,F_G with at least one product in common with factory g

Variables

$PC_{g,p,f,b}$	The change in production size over the problem duration pd of product p for brewery b at factory f after a disruption at factory g (in # of products)
$ME_{p,b}$	Maximum exposure of product p for brewery b (in # of products)

Constraints

Constraints (1)-(8) and:

The following three constraints have to be seen together. They make sure that the value of change can be positive if and only if the bottle is produced at f (10) and missing at g (11). In addition, the change is nonnegative if the bottle is missing at g (12). The size of the change can never be larger than the capacity over the problem duration allows:

$$PC_{g,p,f,b} \leq Y_{p,f,b} * \frac{\sum_{c|c \in c_p} cap_{f,c} * pd}{w_p} \qquad \forall\, g,p,f,b \mid f \in f_g \qquad (10)$$

$$PC_{g,p,f,b} \leq Y_{p,g,b} * \frac{\sum_{c|c \in c_p} cap_{f,c} * pd}{w_p} \qquad \forall\, g,p,f,b \mid f \in f_g \qquad (11)$$

$$PC_{g,p,f,b} \geq (Y_{p,g,b} - 1) * \frac{\sum_{c|c \in c_p} cap_{f,c} * pd}{w_p} \qquad \forall\, g,p,f,b \mid f \in f_g \qquad (12)$$

A larger decrease than the stock level is not allowed:

$$\sum_{f|f \in f_g} PC_{g,p,f,b} + St_{p,b} \geq 0 \qquad \forall\, g,p,b \qquad (13)$$

A larger decrease than production over the duration of the problem is not allowed either:

$$PC_{g,p,f,b} + X_{p,f,b} * pd \geq 0 \qquad \forall\, g,p,f,b \mid f \in f_g \qquad (14)$$

The positive changes in production output of a factory equal the negative changes:

$$\sum_{p,b|p \in p_c} PC_{g,p,f,b} * w_p = 0 \qquad \forall\, g,f,c \mid f \in f_g \qquad (15)$$

The maximum exposure is larger than the exposure per factory for product p destined for brewer b. This exposure is calculated by the difference between impact and compensation level. Moreover, the maximum exposure is nonnegative. Note that here *f* indicates the factory that has a disruption:

$$ME_{p,b} \geq X_{p,f,b} * pd - \left(\sum_g PC_{f,p,g,b} + st_{p,b}\right) \qquad \forall\, p,f,b \mid pr_{p,f,b} > 0 \qquad (16)$$

$$ME_{p,b} \geq 0 \qquad \forall\, p,b \qquad (17)$$

Production size change is only allowed in factories that have overlap with the disrupted factory:

$PC_{g,p,f,b} = 0 \qquad \forall\, g,p,f,b \mid f \notin f_g$ (18)

Objective function, minimize the maximum exposure:

$min \sum_{p,b} MS_{p,b}$ (19)

4.2.3: goals and worst-case values

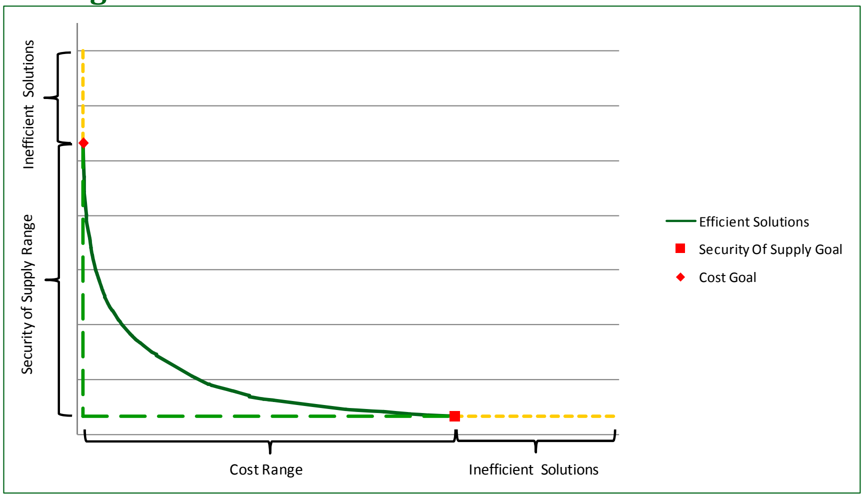

Figure 13: efficient versus non-efficient solutions

No better solutions can be found for the two criteria, than the goal values. That does not mean that there is only one solution with this optimal value. Often multiple allocations can be found with the same cost or security of supply value, but different values for the other criterion. A solution is efficient if it cannot improve on any of the criteria without deteriorating another. We would never accept an inefficient solution, because it can be improved without deteriorating a criterion. The model is indifferent between efficient and inefficient solutions if it only optimizes for one criterion. Figure 13 shows a possible line of efficient solutions between the security of supply goal and the cost goal. Although the goals have fixed points here, the solutions to the base models of the previous sections can be anywhere on the range of inefficient solutions. Note that this representation of the model solution is different from the one in the beginning of this chapter. The reason is that the measure of security of supply is optimal when it is minimized.

Accepting any solution from the base model would not only give an inefficient solution, it would also increase the cost or security of supply range. The solutions of the balanced goal programming model are evaluated based on these ranges. This makes it important to keep the ranges as small as possible. Therefore, we are interested in the marked efficient goal points and define the values of the corresponding criteria as worst-case values. These points are not really the worst-case solutions, rather the worst-case solutions we would accept.

In order to find the worst-case values the base models have to be adjusted a little. The worst-case costs model minimizes the costs, given the optimal value of the security of supply. It uses (1)-(18) and the following constraint to keep the maximum exposure at its optimum (sg):

$$\sum_{p,b} ME_{p,b} \leq sg \qquad (20)$$

The worst-case security of supply model uses (1)-(8) and (10)-(19) with the addition of a constraint to keep the costs at its optimum (cg):

$$\sum_{p,f,b} \left(X_{p,f,b} * pr_{p,f,b} \right) \leq cg \qquad (21)$$

4.3: glass bottle allocation model

Now that the maximum and minimum values for both objectives are found, the zero-one goal programming model can be formulated. The model consists of the same variables and constraints as the previous models ((1)-(8) and (10)-(18)) and minimizes the weighted deviation from the goal value of the two objectives scaled over their ranges. The importance of closeness to the goal value of a criterion is given by a weight. The preference weight for costs (PWC) is defined as α and has a value between zero and one. The preference weight for security of supply is found by 1- α. Depending on the preference for optimizing costs or security of supply, the model balances the improvement of one criterion against the deterioration of the other criterion. These improvements and deteriorations are evaluated as a percentage of their range. This guarantees that the two criteria can be compared even though they have different units and sizes. The goal programming model uses these additional parameters:

Parameters

Cg	Cost goal value
Cwc	Cost worst case value
Sg	Security of supply goal value
Swc	Security of supply worst case value
Α	Preferential weight of costs

The objective function is as follows (where *cg* and *sg* are the cost and security goals respectively, *cw* and *sw* the worst-case values):

Objective function

$$min\left(\alpha * \frac{\sum_{p,f,b}\left(X_{p,f,b} * pr_{p,f,b}\right) - cg}{cwc - cg} + (1-\alpha) * \frac{\sum_{p,b} ME_{p,b} - sg}{swc - sg}\right) \qquad (22)$$

The model will only provide efficient solutions for PWC values between zero and one. The base models find the efficient solutions at these boundaries.

Bibliography

AIMMS, 2007. 'Optimization Modeling.' *Published on* www.aimms.com.

Aissaoui, N., Haoauri, M. & Hassini, E. 2006. 'Supplier Selection and Order Lot Sizing Modeling: A Review.' *Computers & Operations Research, in press.*

Beamon, B.M. 1999. 'Measuring supply chain performance.' *International Journal of Operations & Production Management.* Vol. 19, No. 3, 275-292.

Boer, L. de, Labro, E. & Morlacchi, P. 2001. 'A review of methods supporting supplier selection.' *European Journal of Purchasing & Supply Management.* Vol. 7, No. 2, 75-89.

Burke, G.J., Carrillo, J.E. & Vakharia, A.J. 2007. 'Single versus multiple supplier sourcing strategies.' *European Journal of Operations Research*. Vol. 182, No. 1, 95-112.

Cachon, P.C. & Zhang, F. 2003. 'Procuring Fast Delivery, Part I: Multi-sourcing and Scorecard Allocation of Demand Via Past Performance.' *Working Paper, Wharton School, University of Pennsylvania, Philadelphia, PA.*

Canadean, 2007. *Global Brewer Handbook. Annual Report.*

Christopher, M. & Peck, H. 2004. Building the Resilient Supply Chain. *The International Journal of Logistics Management*, Vol. 15, No. 2, 1-13.

CIWMB*, California Integrated Waste Management Board Website* (www.ciwmb.ca.gov) (accessed November 2008).

Flavell, R.B. 1976. 'A new goal programming formulation.' *Omega*, Vol. 4, 731-732.

Gaonkar, R. & Viswanadham, N. 2003. 'A Conceptual and Analytical Framework for the Management of Risk in Supply Chains.' *Submitted to IEEE Transactions on Automation & Systems Engineering.*

Gass, S.I. 1986. 'A Process for Determining Priorities and Weights for Large-Scale Linear Goal Programmes.' *Journal of the Operational Research Society*. Vol. 37, No. 8, 779-785.

GPI. *Glass Packaging Institute website* (www.gpi.org) (accessed March 2009).

Heineken N.V., 2007. *Annual Report.*

Ignizio, J. P. & Romero, C. 2003. 'Goal programming.' *Encyclopedia of Information Systems.* Vol. 2, 489–500.

Jayaraman, V. Srivastava, R. & Benton, W.C. 1999. 'Supplier selection and order quantity allocation: A comprehensive model.' *The Journal of Supply Chain Management.* Vol. 35, No. 2, 50–58.

Kardamanidis, D. (account manager Ardagh Glass), interview. November 2008.

Kraljic, P. 1983. 'Purchasing Must Become Supply Chain Management'. *Harvard Business Review.* September-October.

Neely, A., Mills, J., Platts, K., Richards, H., Gregory, M., Bourne, M. & Kennerley, M. 2000. 'Performance measurement system design: developing and testing a process-based approach'. *International Journal of Operations and Production Management*. Vol. 20, No. 10, 1119-1145.

OI. *Owens-Illinois Website* (http://www.o-i.com) (accessed November 2008)

Paulsson, U. 2003. 'Managing Risks in Supply Chains – An Article Review.' *NOFOMA Conference, Oulu, Finland.*

Plato Logic, 2008. *Western Europe Beer Report.*

Porter, M.E. 1979. ‘How competitive forces shape strategy’. *Harvard Business Review.* March/April.

Schmitt, A. & Snyder, L.V. 2006. ‘Infinite-horizon models for inventory control under yield uncertainty and disruptions.’ *Working paper, Lehigh University, Bethlehem, PA*.

Sheffi, Y., Rice Jr., J.B., Fleck, J.M. & Caniato, F. 2003. ‘Supply Chain Response to Global Terrorism: A Situation Scan.’ *EurOMA POMS Joint International Conference*.

Snyder, L.V., Scaparra, M.P., Daskin, M.S. & Church, R.L. 2006. ‘Planning for Disruption s in Supply Chain Networks.’ *Tutorials in Operations Research INFORMS.*

Snyder, L.V. & Shen, Z.M. 2006. *Supply and Demand Uncertainty in Multi-Echelon Supply Chains*. Lehigh University Working Paper.

Stecke, K.E. & Kumar, S. 2006. ‘Sources of Supply Chain Disruptions, Factors that Breed Vulnerability, and Mitigating Strategies.’ *Working paper, University of Texas, Dallas, Texas.*

Tamiz, M., Jones D. & Romero, C. 1998. ‘Goal programming for decision making: an overview of the current state-of-art.’ *European Journal of Operations Research*. Vol. 111, 569-581.

Tang, C.S. 2006. ‘Robust strategies for mitigating supply chain disruptions.’ *International Journal of Logistics: Research and Applications*. Vol. 9, No. 1, 33-45.

Tang, C.S. 2006b. ‘Perspectives in Supply Chain Risk Management: a Review.’ *International Journal of Production Economics.* Vol. 103, 451-488.

Telgen, J., Buter, J. & Schotanus, F. 2006. ‘Purchasing Management.’ *Twente University Lecture Reader.*

Tomlin, B. 2006. ‘On the Value of Mitigation and Contingency Strategies for Managing Supply Chain Disruption Risks.’ *Management Science*, Vol. 52, No. 5, 639–657.

Treleven, M. & Bergman Schweikhart, S. 1988. ‘A Risk/Benefit Analysis of Sourcing Strategies: Single vs. Multiple Sourcing.’ *Journal of Operations Management.* Vol. 7, No. 4, 93-114.

Weele, A.J. van. 2006. ‘Purchasing and Supply Chain Management: Analysis, Planning and Practice.’ *Thomson Learning*.

Zitzler, E. 1999. ‘Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications.’ Doctoral Thesis at the Swiss Federal Institute of Technology Zurich.