Syngenta Uses a Cover Optimizer to Determine Production Volumes for Its European Seed Supply Chain

24 November 2014

Peter Comhaire, Syngenta Supply AG, Schwarzwaldallee 215, 4002 Basel, Switzerland, peter.comhaire@syngenta.com

Felix Papier, ESSEC Business School, Av. Bernard Hirsch, 95021 Cergy, France, Phone +33-1-34-43-36-58, papier@essec.edu

Abstract

The European seed business of Syngenta relies on its supply chain to supply seed products to the different European markets, which led to more than 1.2 billion USD revenues in 2013. The seed supply chain is, however, exposed to a high level of uncertainty – from the demand side as well as from the supply side. Determining optimal production volumes in a highly volatile environment and more than one year before the sales period is not only a complex business decision but also one which strongly affects the company's profitability through lost sales and unsold supply. In order to better handle the production volume planning, Syngenta has developed a planning tool which determines optimal production volumes by taking the different levels of uncertainty into account. We report on this tool, the impact it has achieved, its integration into the planning process at Syngenta, and its technical design. In 2013, its first year of application, the production optimization tool has already avoided approximately 1.5 million USD in supply discards and has led Syngenta to revise the way how it handles uncertainty in its supply chain planning.

Key words: agricultural supply chain; production; inventory; optimization; decision support

Introduction

Syngenta Crop Protection AG is one of the world's leading companies in the agricultural seed and crop protection industry. It has been formed in 2000 as a joint venture between Novartis Agribusiness and Zeneca Agrochemicals. Since then, the company has achieved a remarkable market position among its competitors such as Bayer Crop Sciences, Monsanto, and Dupont Pioneer. In 2013, Syngenta had total revenues of 14.7 billion USD, thereof 3.2 billion USD (22%) in its seed division. Syngenta produces seed for various types of field crops such as barley, cereals, corn, sunflower, soybean, and vegetable and flower seeds, and sells the seed through wholesalers and agricultural co-operatives to farmers.

The European seed business is a key business division in Syngenta's portfolio with total sales of 1.2 billion USD in 2013, up by 12% from the previous year. While the European seed business is showing profitable growth, the management of its supply chain is nevertheless a challenging task. The long lead times of the supply chain process (more than 12 months), the spread of the Europe-wide production network with more than 40 seed processing sites and thousands of supply farms, and the difficult regulatory conditions in each of the ~30 European sales countries constitute the complexity of the planning task. All of the above three factors contribute to the immense uncertainty that the European seed supply chain faces, from the demand side (the forecast error for seed products ranges from 15-85%) as well as from the supply side (the standard deviation of the supply yield ranges from 15-60%).

The European seed supply chain at Syngenta is managed by the EMEA[1] Supply Operations team, which is based in the corporate headquarters in Basel, Switzerland. The Supply Operations team

[1] EMEA = Europe, Middle East, Africa

is responsible for the coordination of the operational activities along the whole supply chain, from the field production to the final processing, certification, and distribution of the seed. One of the key challenges in this supply chain is the determination of the seed production volumes for the next sales season. The annual seed production plan (SPP) determines how many fields will be planted of each crop type and variety in the different European countries. Since most crops have a single planting time per year, the seed production volumes have to be determined approximately twelve to fifteen months before their sales period. The seed production has a strong impact on the financial position of Syngenta because it does not only directly determine a large part of the operational costs (e.g., field, harvesting, and processing costs), but it also directly affects the revenue side through the availability of the seed products in the market. With product margins as high as 70%, lost sales lead to high opportunity costs and allows Syngenta's competitors to grab additional market share. But despite the high margins, safety buffers are built with caution into the seed production plan. Seed products that are not sold in a sales cycle do not only lead to costs for storage to the next sales season, many units will also not pass the required quality germination test and become unsellable or even require additional costs for disposal. The trade-off between lost sales on the one side and leftover supply on the other side is the focus of this paper.

In the past, Syngenta has set capacity levels to satisfy the need for supply reliability without explicit consideration of the financial implications of lost sales versus costs for unsold supply. With the recent increase of competitive pressure and a decline of profit margins in the European seed industry, Syngenta has taken a closer look at the profitability of its business portfolio. As a consequence, Syngenta has decided to make use of advanced analytical tools to improve the profitability of the seed supply chain in Europe. In a cross-functional project between marketing, sales, production, supply operations, and finance, the company has developed the Cover

Optimizer (CO) tool and has integrated it into the seed planning process. The CO tool supports the planning team in determining optimal production volumes in the supply chain and in performing scenario analyses for alternative supply chain configurations. In this article, we report on the CO tool and the approach that Syngenta has taken. In addition to the development of the planning tool, a key part of the CO project has been devoted to changing the mindset of the planners involved in the seed planning process, from a linear, deterministic approach of forecast, supply, and budget planning to an approach that explicitly takes uncertainty from multiple sources as well as financial profitability into account.

Determining optimal production volumes in the seed industry is related to determining optimal safety stock levels in inventory management. The initial idea of the CO tool is therefore based on the well-known newsvendor model for safety stock planning in a single-period setting (see Khouja 1999 or Ahumada and Villalobos 2001 for reviews). The CO tool, however, has been designed to deal with the more complex situation of the seed supply chain, where seed can be stored to some extent between multiple sales cycles, where uncertainty arises not only from the demand side but from also from the supply side in production and storage, and where production and sales take place in multiple European countries.

Jones et al. (2003) also reported on a project in the context of seed planning for Syngenta in the United States. However, their focus is on a second production option through introducing a second planting time in South America to better match supply with demand, while our article focuses on matching supply and demand in a complex, multi-country environment in Europe. In contrast to a homogeneous market such as the one of the United States, the European market requires that the specifics of every country are respected, for example with regard to differences in production costs, demand and supply uncertainty levels or sales prices.

The article is structured as follows: In the next section, we discuss Syngenta's seed supply chain and the different sources of demand and supply uncertainty (Section *Syngenta's European Seed Supply Chain*). Then, we give an overview on the implementation of the CO tool at Syngenta and on the integration of the tool into the supply chain process (Section *Implementation of the Cover Optimizer at Syngenta*). Afterwards, we describe the technical design of the CO tool (Section *Description of the Cover Optimizer*). Finally, we report on the impact and benefits of the CO project (Section *Impact and Benefits*). A mathematical description of the CO tool is contained in the appendix.

Syngenta's European Seed Supply Chain

Syngenta sells seed products for a variety of different crops. For each crop, the company offers several hundreds of different variants, so-called *hybrids*, which differ in their genetic composition. Different hybrids have to be separately planned, grown, and processed, even though they may be from the same crop. The seed products are sold by local cooperatives and local wholesalers to farmers. Seed sales at Syngenta are spread over more than 30 countries in Europe. The seeds are used by farmers to grow the respective crops, mainly to feed animals or to sell to the commodity market. Farmers decide on the crop and on the hybrid based on the expected yield, the expected price development of crops, the weather conditions, the soil conditions, and other factors.

In order to meet the need of the farmers, Syngenta has set up a supply chain process for its seed products, which consists of a set of key steps:

1. <u>Creation of initial forecasts.</u> The commercial side (marketing and sales) determines the first forecasts for the different sales countries for the planning year, *year N+1,* and for the

year after (*year N+2*). The forecasts of the current year (*year N*) are also updated to reflect the latest information about demand. These forecasts are made on hybrid-level and are based on the market environment, the relative competitive position of the production portfolio, and other factors.

2. Creation of the seed production plan (SPP). The long-term planning team within Supply Operations coordinates the creation of the SPP, which determines the number of fields to plant for each hybrid and in each production country. Seed is grown by farmers who have a production contract with Syngenta. Farmers have to fulfill a series of requirements in order to be eligible for seed production. When determining the SPP, the strategic crop development plan is taken into account. The *strategic crop development plan* is defined on a long-term basis by the crop teams and considers factors such as product margins, strategic market developments, market shares, and competitor information.
3. Field production. In line with the SPP, the farmers that are contracted by Syngenta plant, grow, and harvest the seed. They receive technical support from Syngenta's field production team. Farmers are mainly compensated for harvested fields only.
4. Seed processing. The seed processing takes place in two steps. First, the seed is cleaned and filtered for size to ensure that it meets the quality specifications. The second phase concerns the finishing of the seed products and consists of treatment of the seed with seed care products and of the final packaging. This is done both in Syngenta's own processing plants in Europe and in external processing plants. After the processing, the seed is stored in warehouses or directly distributed to the market.
5. Allocation and distribution. Based on revised short-term forecasts for the market demand and updated numbers for the available supply, the allocation team decides on the

allocation of seed to the different markets. After the allocation the logistics team organizes the distribution of the finished products to the markets.

6. Sales. The seed is sold during the sales seasons of the different sales countries. Sales seasons typically only last for a couple of weeks in each country.
7. Collection of unsold supply. In the last step of the supply chain cycle, the unsold supply is collected from the different warehouses and prepared for storage to the next sales season, one year later. Much of the stored supply will not pass the germination test in the next season and cannot be sold. The disposal of this supply is referred to as *discard* and entails significant write-offs in Syngenta's balance sheet. If seed has not been treated with chemicals, it may be sold as ordinary crop in the commodity market at a price which is significantly lower than the sales price of seed. If seed has been treated with chemicals, it has to be incinerated at an additional expense. Storing seed for the next season also incurs significant inventory costs, such as costs for rental of storage space, for re-testing the quality, and for holding the capital.

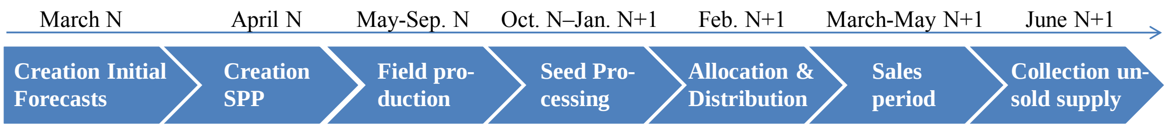

Figure 1: *Overview of the seed supply chain process with the typical time line for corn seed.*

The seed supply chain process is summarized in Figure 1. The figure also shows the time line of the seed supply chain for the example of corn seed (the time lines vary between crops). As can be seen, the lead time of the planning process is quite long. The seed is planted and grown from May to September of the previous year (*year N*), processed from October to January, and finally sold from March to May of the following year *(year N+1)*. The first forecasts and production decisions have consequently to be made 12-15 months before the sales period.

The seed supply chain is exposed to uncertainty from different sources, from the demand as well as from the supply side. The market demand for seed of a certain crop is already uncertain and may depend on different, unpredictable events. For example, after the first political crisis in the Ukraine in early 2014, the financial system in the country made loan conditions for local farmers much stricter than they were before, which nearly led to a complete collapse of the seed market in this country. On a hybrid-level, the demand is even more uncertain, which is a consequence of the complexity of the product portfolio. Syngenta as well as its main competitors offer hundreds of hybrids for each crop, which renders forecasting the demand for a particular hybrid difficult. The average error of the initial forecasts on country- and hybrid-level ranges from 15-85%, with an average value of 30-40%.

The uncertainty on the supply side is also immense and has various sources, such as the variability of the carry-in supply, the risk of losing whole fields, and the variability in the seed processing stage. For the carry-in supply, i.e., seed that has been stored in the previous season, the average discard rate is approximately 10%, which means that one out of ten bags of the carried-over supply has to be discarded. In addition, the carry-in supply is not fully known at the time when the SPP for the next season is decided, because the sales season is still ongoing. For the field production, there is a binary risk of losing a complete field. This can happen when bad weather conditions prevent the farmer from planting the crop, when heat and drought stress prevent pollination, or when the genetic purity of the seed is lost due to contamination with weed or diseases. Finally, the seed processing phase also has a strong variability, in particular because the product quality is tested at the end of this stage. The standard deviation of the seed processing stage ranges from 15-60%, depending on the type of hybrid.

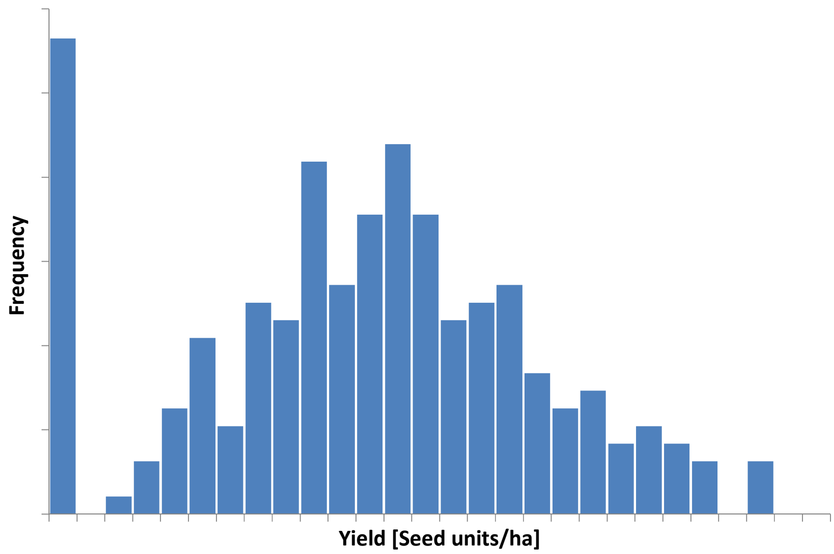

Figure 2: *Histogram of supply variability (realized yield versus plan).*

Figure 2 shows a histogram of realized yield of fields of a sample crop, taken over five years. The name of the crop and the values of the horizontal axis have been removed for confidentiality reasons. The figure demonstrates the variability of the finished seed per field versus the plan. It also indicates that the risk of losing complete fields (the bar at zero yield) is significant in this crop example.

As a consequence of the complexity of the seed supply chain, it happens that even if the total demand exceeded the total supply of a certain hybrid, supply is left at the end of the sales season. This may occur if the chemical treatment of the remaining seed does not correspond to the market demand or if farmers change their crop plans in the last minute such that the supply cannot be redeployed to other regions. These process inefficiencies can attain values between 2% and 30% of the overall supply.

Implementation of the Cover Optimizer at Syngenta

The CO project was started in 2012 with the participation of the different functions that are involved in the seed supply chain process, such as marketing, sales, production, Supply Operations, and finance. One outcome of the project was the CO tool. The CO tool has been installed on the desktop computers of the long-term planning team in Basel, Switzerland, in Toulouse, France, and in Mezotur, Hungary. The tool was used in 2013 for the first time for the planning of two major crops with more than 300 hybrids.

The CO project did not only focus on the development of the tool. Much effort was also spent on changing the mindset of the participants of the supply chain planning process. A major objective was to create awareness for the impact of uncertainty in the business planning. In the past, the budgeting process did not systematically take the sources and level of uncertainty into account. The demand side forecasted, for example, a demand of 100,000 units for a certain hybrid. The supply side consequently planned supply of 100,000 units with 30-35% additional inventory as buffer, and the finance team budgeted sales of 100,000 units. However, because sales are the minimum of market demand and available supply and because both factors are subject to uncertainty, the expected sales were often lower than 100,000 units and budget plans on hybrid-level were missed. We performed workshops with the participants of the seed planning process to discuss that “100 units forecasted demand and 120 units of planned supply make only 90 units of sales”. This change in mindset, where participants explicitly consider the levels of uncertainty and their impact, significantly helped to better align the three plans (forecasting, supply planning, and sales budgeting) and rendered the business planning more robust.

The CO tool is used for the creation of the annual SPP. The long-term supply chain planning team uses the tool during the multiple iterations of the planning process, when demand forecasts are updated and when the first production quantity decisions have to be made. The CO tool proposes an optimal production volume, which consists of a base quantity (equal to the demand forecast) and an additional safety stock, referred to as *risk cover*. The additional cover compensates for lost supply in the case that production yield turns out to be lower than anticipated and it generates additional revenues and profit in case that demand turns out to be higher than forecasted. Figure 3 shows an example of optimized risk covers for a sample of hybrids of a major crop. The graph demonstrates that the optimal risk cover varies substantially from hybrid to hybrid, depending on the individual forecast accuracy and the production variability of the product as well as its footprint in the European production network. Some hybrids have an optimal risk cover of 100% or above because the hybrids have a small forecast, such that the size of a single field exceeds more than twice the forecast.

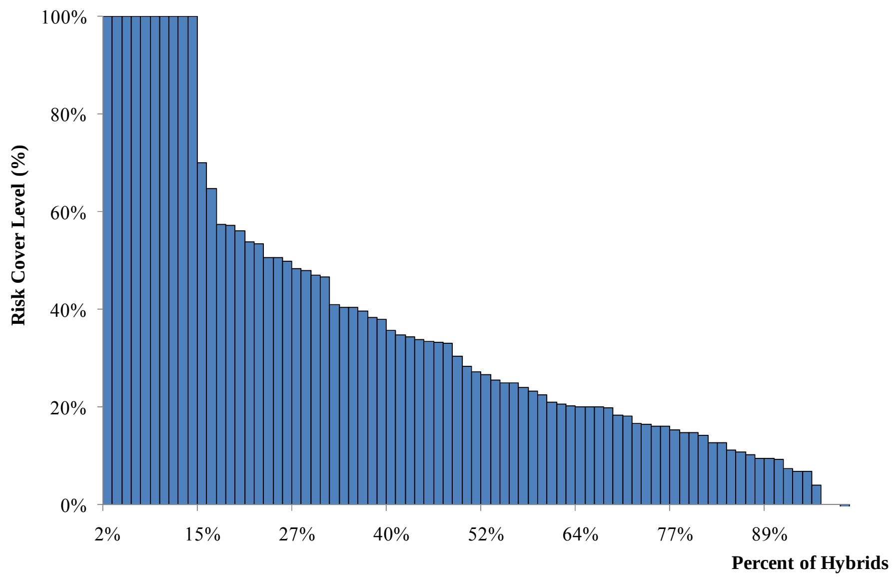

Figure 3: *Distribution of optimal risk cover levels for a major crop.*

The process of creating the SPP is done in several steps. In a first step, participants from the demand side (marketing, crop teams, and portfolio managers) and the supply side (long-term planning) determine the demand forecast figures and align these figures with strategic and financial considerations. In a second step, the long-term planning and the field production teams review the input parameters for the production planning (e.g., target yields, available field capacity per country). After these steps, the long-term planning team updates the parameters of the CO tool with estimates from historic data or expert judgment (see also the discussion at the end of this section) and calculates the optimal risk covers and planned production volumes. This information is then discussed and manually refined in a third step and a draft SPP is created that aligns optimal supply decisions with the financial targets and constraints and with the overall portfolio strategy. The planners also verify and stress-test the outcome of the CO tool. In case that the results are surprisingly low or high, the long-term planning team verifies the input parameters for these products and revises the calculation if necessary. The different drafts converge to a final SPP, which is signed off by management. The SPP is the single biggest financial commitment of Syngenta's seed business and forms the basis for all supply-chain-related decision making throughout the year. The SPP is updated monthly by Supply Operations to integrate new information, such as revised forecasts or information about supply yield.

The CO tool significantly changed the way production volumes are determined, not only through the calculation of the optimal cover levels but also through the facilitation of the discussion among the stakeholders of the seed supply chain. For one example, a hybrid with a strong growth potential, planners expected a high optimal risk cover to protect the growth. However, the CO tool proposed a surprisingly low risk cover. After some analysis, the planning team discovered that the discard rate of this hybrid was particularly high, which makes unsold supply excessively expensive. Based on this insight, the team proposed a medium production quantity that considers the future growth potential but which also respects the particular cost situation of the hybrid. Another example, a hybrid that has strong strategic importance to Syngenta, was expected to have a relatively high optimal risk cover. However, the CO tool proposed an even higher cover than was expected, because, in addition to the strategic importance, the cost situation of the hybrid allowed for high risk cover levels to capture even unlikely demand upsides.

The CO tool also supports the long-term planning team in managing budget constraints on the supply chain. In several cases, Syngenta decided to limit the amount of working capital in the supply chain of certain crops, even though such a restriction does not lead to the maximum expected operational profit. The CO tool is used to simulate the impact of working capital constraints on the supply chain. Finally, the CO tool is used to perform simulations of the impact design improvements of the supply chain processes and production practices, such as the location and spread of planting areas, the improvement of production techniques for irrigation or seed storage, and the option to sell unsold seed on the commodity market rather than to carry it over to the next year. The CO tool estimates the expected sales and unsold supply volumes for each hybrid under the different scenarios.

Description of the Cover Optimizer

In this section, we describe the features of the CO tool as it is implemented at Syngenta. The tool has been developed in Visual Basic for Applications and runs in Microsoft Excel on the desktop computers of the long-term planning team. Figure 4 shows the user interface of the tool (data has been removed for confidentiality reasons). Every line contains the planning parameters of one hybrid-country combination. The tool optimizes hybrids on a Europe-wide scope and proceeds from hybrid to hybrid by first searching the parameter table for all hybrid-country pairs that correspond to the product that is optimized.

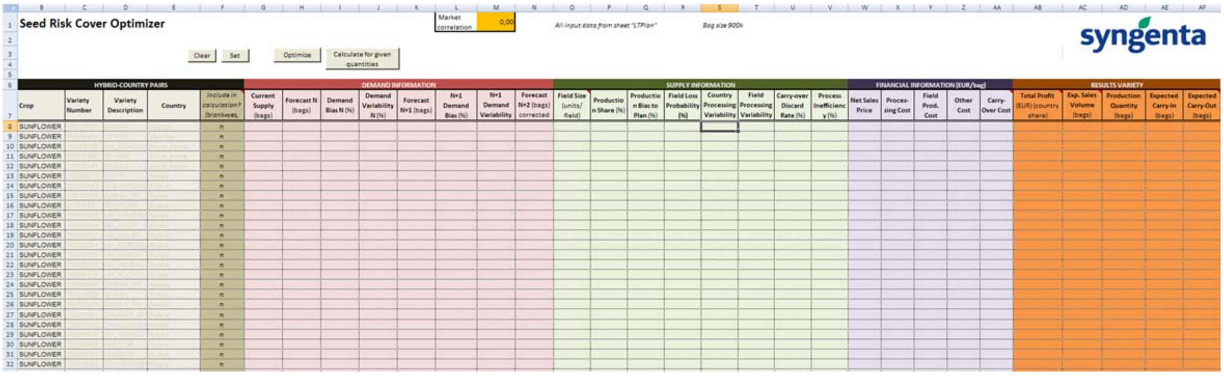

Figure 4: *Screenshot of the user interface of the CO tool.*

The CO tool determines for any given hybrid the optimal number of fields to plant in each European country. The project first decided to include an optimization of the regional spread of the production across different countries, but it was finally decided that the spread among production countries for each hybrid should be determined outside the CO tool and should be considered as an input parameter. This decision has been made because the geographic spread of production takes more factors into account than just financial profitability and is considered a decision that is difficult to optimize quantitatively. For example, it could be necessary for political or strategic reasons to produce part of the overall supply in some of the main sales countries to increase acceptance of the brand within the country. In some cases, production should be spread across different geographic regions for better genetics. Therefore, the CO tool receives a *target production share* for each country. For example, for a certain hybrid, the target share could require that production is split 40%, 40%, and 20% between France, Spain, and Italy.

Syngenta can only contract whole fields with farmers, which means that the CO tool has to plan for integer values for the number of fields to plant. In addition, the size of the fields differs from country to country. The CO tool determines an integer partition of the fields to plant across the European countries, such that the expected profit is maximized and the target production share is respected as closely as possible. To evaluate the expected profit for a given integer partition, the tool builds three different stochastic distributions: For the market demand, for the existing supply that is carried-over from the previous year *(carry-in)*, and for new supply that is produced in the current year (*production*). The three distributions are considered to be independent because the production is done in the season before the sales and is separate from the sales process and the storage process.

For the calculation of the distribution of the market demand, the CO tool considers the current forecast, the forecasting bias, and the standard deviation of the forecast error for each hybrid-country pair. The forecasting bias is the average level of over- or under-forecasting. The forecasting bias is tracked over time and is used to re-adjust the forecasts before they are used for planning. Correlation between market demands in the different countries is modeled through an additional parameter, although historic data showed that the market demand is not significantly correlated. Based on data about forecast errors in previous years, the project team concluded that a normal distribution represents well the distribution of the forecast error and this assumption was consequently used.

A key question concerned the modeling of the allocation step in the supply chain process. After discussion with the allocation team, the project team found that an equal-share allocation based on the short-term forecasts represents best the complex, manual allocation process. Under an equal-share allocation, all countries receive a fixed percentage of their forecasts, independent of their prices and other factors. The allocation is based on the short-term forecasts, which are already available when the allocation is decided, and which are significantly more precise than the original forecasts on which the production volumes have to be decided. As a consequence of this modeling decision, the marginal net sales price that is earned for an additional unit of supply can be approximated by the forecast-weighted average net sales price of the different countries. Furthermore, the allocation based on short-term forecasts is more precise than an allocation based on the initial forecasts. This observation is modeled by assuming that the total available supply can to a large extent be used to serve the total market demand in Europe. This is not entirely true in reality because supply that is allocated based on short-term forecasts but which cannot be sold completely in this country may not be used for other countries in the same season. This allocation effect can be modeled in the CO tool through the parameter *process inefficiency*, which reflects the average loss in sales even if the Europe-wide supply would have been sufficient to cover the total market demand.

For the distribution of the total production volume, the CO tool considers the uncertainty in field production through an individual field loss probability. The individual field loss probability of each field leads to a binomial distribution of the total number of harvested fields in a country. For the processing yield, the CO tool uses a normal distribution with an additional production bias (the difference between planned and average realized yield). The data sample shown in Figure 2 confirms this choice: neither of the Chi-Squared, the Kolmogorov-Smirnov, and the Anderson-Darling test rejects the normal distribution on a 5% level of significance (see Thode 2002 for more information about the statistical tests). The variance of this distribution depends on the number of fields and on an additional variability on country-level. The reason for using two parameters for modeling the production variability is that a part of this variance stems from the individual conditions of the growing field, whose impact decreases with an increasing number of fields (pooling effect). Another part of this variance stems from country- or region-wide factors, such as extreme heat or cold, that affect the yield, independent of the number of fields. This variability is modeled through the country-level variability parameter. Figure 6 illustrates the density function of the yield distribution for different numbers of fields. In the example shown, the field loss probability is 10%, the field-level variability 36%, and the country-level variability 0%.

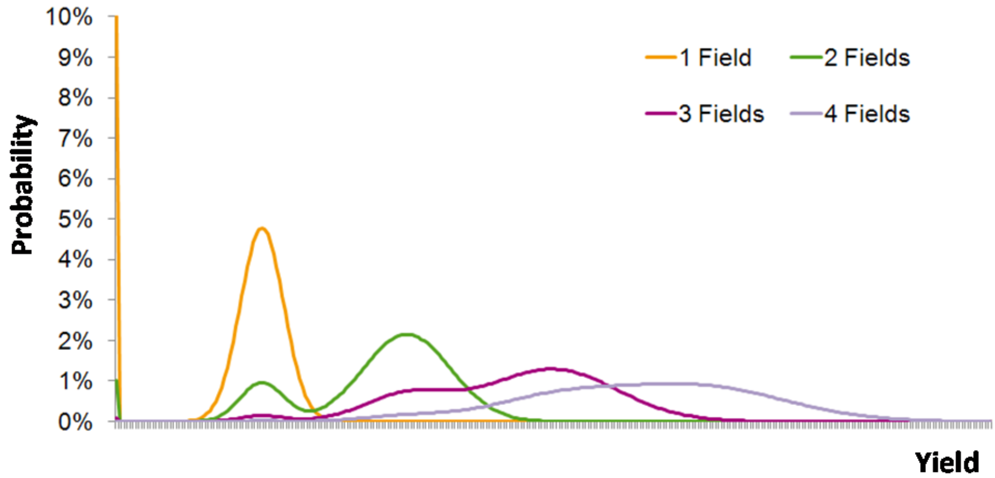

Figure 5: *Yield distributions for different number of planted fields.*

The computation of the distribution of supply that is carried-over from year N to the planning year N+1 is complicated by the fact that the final sales in year N are not known when the production quantity for year N+1 has to be decided (for corn, in March of the previous year). Therefore, the CO tool first calculates the distribution of unsold supply by considering the difference between existing supply (a fixed value, because the available supply of year N is already known when the production volumes are planned) and the uncertain market demand of year N, which is specified through the short-term forecast, the forecast bias, and the forecast error of year N. The difference of the existing supply and the market demand in year N determines the total, Europe-wide *carry-out* supply. The carry-out supply is diminished by the discards, which are given by a hybrid-specific discard rate, which finally leads to the *carry-in* supply of year N+1.

Once the three distributions (demand, production, and carry-in supply) are built, the CO tool computes the distribution of the total sales as the minimum of market demand and total supply, given by the sum of the carry-in supply and the realized production volume. Unsold supply is given as the difference between the total supply and the sales volume. The final output for each hybrid consists of the optimal number of fields to be planted in each country, the expected sales volumes and revenues per country, the expected field and processing costs per country, and the expected carry-in and carry-out volumes. The calculation flow is summarized in Figure 6. Figure 7 shows the distributions calculated by the CO tool for an example of a sunflower hybrid. The distribution of the total sales volumes is shown as a thick, purple line. It can be seen that the mode of the sales distribution is less than the modes of the distributions of the market demand and of the total supply, which corresponds to the observation that we made earlier (see Section "*Implementation of the Cover Optimizer at Syngenta*").

For the field costs and the processing costs, every hybrid-country pair has individual cost rates to allow for cost differences in farming and processing operations between countries and hybrids. Carry-out supply is evaluated at the production cost that is avoided from having carry-in supply in the next year. The current policy of storing seed at Syngenta specifies that at most the forecast for the next year can be carried-over. Consequently, the carry-over supply in the CO tool is limited to the forecast of year N+2 and the remaining carry-out supply is discarded or sold to the commodity market, if possible.

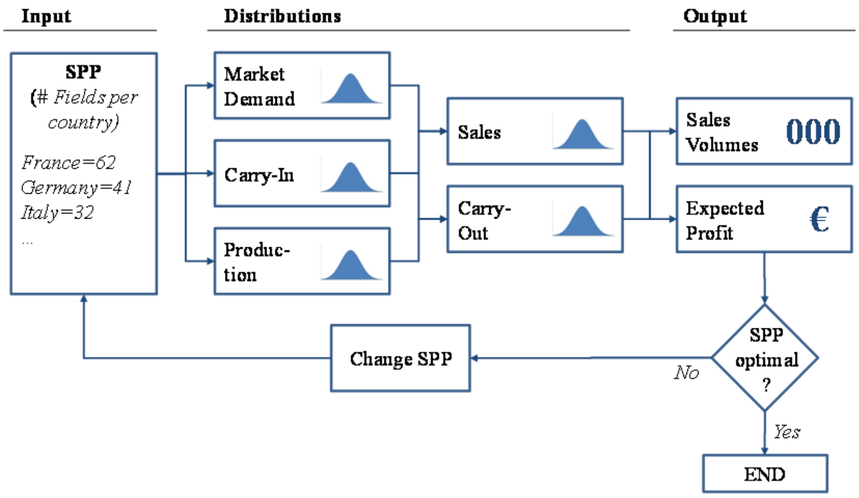

Figure 6: *Flow scheme of the CO calculation logic.*

The CO tool can be used in two ways: It either calculates the above performance indicators for a given field configuration (to evaluate business scenarios) or it determines the optimal field configuration that leads to maximum expected profitability by respecting the target production share as closely as possible. The run time for the optimization is approximately 2-3 hours per crop. The project team is currently working on an extension to allow the specification of an overall financial budget for risk covers for a crop and the determination of the optimal allocation of this budget to the different hybrids. A mathematical formulation of the evaluation and optimization logic of the CO tool is contained in the appendix.

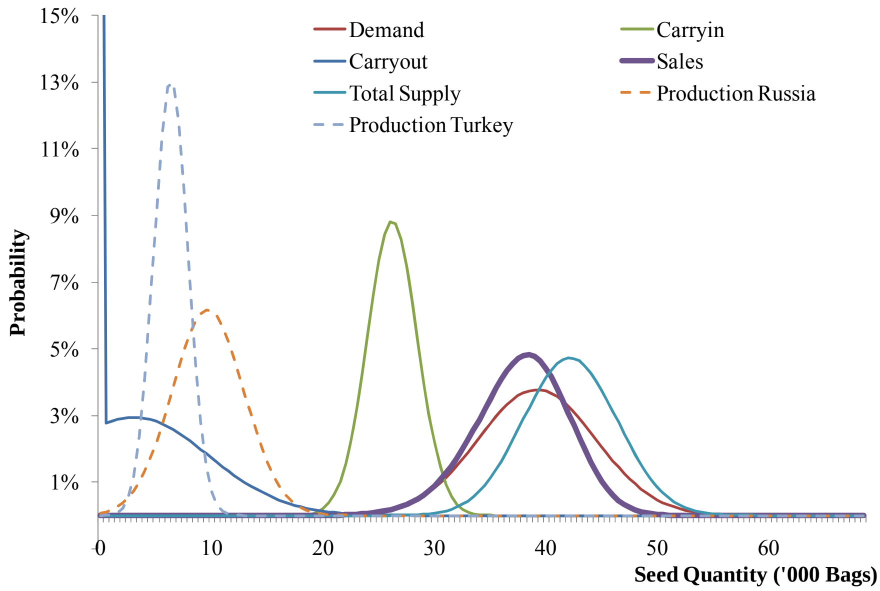

Figure 7: *Demand and supply distributions for a sunflower hybrid, as created by the tool.*

One of the key challenges of the CO project was the parameterization of the model. Some of the parameters are directly taken from planning systems, such as current forecasts for years N, N+1, and N+2 or available supply in year N. Most of the parameters, however, have to be estimated either from historic data or through expert judgment. Forecast biases and errors, field loss probabilities, and production variability could be derived from historic data of previous years. In cases where not enough data was available for an individual hybrid, values were estimated on the level of maturity groups or crops. For some parameters, expert judgment had to be used. These parameters were estimated by a panel of experts from field production, production and processing technology engineers, crop specialists, R&D, and other functions. Sales prices and cost rates were derived from the financial systems with the support of the financial team. The finance team uses the controlling systems to derive the average realized costs per country and hybrid. Net sales prices were forecasted by the finance team, based on extrapolations of current prices, expected commodity prices, and other factors. The CO project has defined an annual data update process for the CO tool, which is led by the long-term planning team. During this process, all parameters of the model are updated based on new historic data.

Impact and Benefits

Determining an effective seed production plan is a difficult but strategically important task for Syngenta. Long lead times, high levels of uncertainty from different sources, and the diversity of a Europe-wide production network render the planning task highly complex. The CO project that we described in this article supports Syngenta in determining production quantities that optimally balance write-offs from unsold supply with lost revenues from supply shortages.

The annual impact of the CO tool on the bottom-line of Syngenta depends strongly on the year because in some years market demand is low while in other years market demand is high. Therefore, in some years, the avoidance of supply discards is more prevalent, while in other years, the reduction of lost sales is more important. In its first year of application, 2013, the CO tool led to a reduction in discards in the order of 1.5 million US dollars compared to the initial, manually determined production quantities. In this year, the carry-in supply from the previous year was particularly large, which has been well anticipated by the CO tool. In addition to the financial impact, the CO project led to an improved seed planning process, in which the level of uncertainty is explicitly considered and mutually accepted by the production, supply, and commercial teams.

A key lever for the benefits of the CO tool is the differentiation in the risk cover levels for individual hybrids. While before, a standard risk cover level was applied to most hybrids, with some manual adjustments, the CO tool calculates individual risk cover levels for every hybrid.

For example, two related and strategically important hybrids traditionally had a common risk cover level of 35%. The CO tool confirmed the overall risk cover level but suggested a cover level of 15% for one hybrid and a cover level of 50% for the other hybrid. The subsequent discussion in the seed planning process also led to a better understanding of the factors that influence the optimal risk cover levels. Furthermore, the CO tool identified some smaller hybrids for which a zero production quantity is optimal. This is an interesting insight and is mainly a consequence of the large field sizes of these products compared to their annual forecasts. Based on the results of the CO tool, the seed planning team introduced a new production policy and encouraged the portfolio managers to remove some of the smaller hybrids from the product portfolio, or to, at least, evaluate the financial impact if the hybrid was kept in the product portfolio.

In addition, the CO tool led to an improved design of the supply chain processes. The CO tool is used to simulate the impact of design changes on the expected sales quantities and profit as well as on the business risk. One example of such design changes concerned the target production share. Before the CO project, the direct cost impact of a revised target production share could already be determined from the differences in field and processing costs, but the effect of this change on the variability of the supply distribution could not be evaluated, which left the assessment incomplete.

The cover optimizer supports Syngenta in better assessing the uncertainty in demand and in supply and in making more conscious and data-driven decisions in the creation of the SPP; a plan that determines the majority of the annual expenditures of Syngenta's seed business unit. In an increasingly competitive environment, this is particularly important for Syngenta, which seeks to preserve its historically high profit margins.

Verification Letter

Alain Millet, Head of Demand Planning Seeds, writes: “This letters verifies the application of the ‘Cover Optimizer’ (CO) tool in the Syngenta seed supply chain in Europe. Seed production planning is related to different factors dealing with a lot of uncertainties and volatilities. We welcomed the development and the use of the cover optimizer (CO) tool for planning the supply chain production volumes of our large field crops seed products (barley, corn, and sunflower). The total annual revenues of these products in EAME amount to $ 770 million per year. The CO has been developed as a joint project under the supervision of Peter Comhaire and with the involvement of our supply chain, production, quality, marketing and commercial teams as well as with the support of Felix Papier from the ESSEC Business School. The CO has been intensively tested by different experts and crops before its implementation in the seed production planning processes. Since then, it helps us to take appropriate decisions for reducing the impact of lost sales and over-production of seed. In addition to the monetary benefits from a reduction of discard costs, the tool and the surrounding project has led us to a more focused and data-driven discussion of the strategic priorities in the management of our portfolio and our seed supply chain. This has re-enforced our approach of long-term supply chain planning, an area of our business process which has high strategic priority. We are currently extending the tool to other crops to further increase the benefits. Combined with the expertise and the experience of our teams, I am confident that the CO will continue to be an integral part of our long-term supply chain planning process in the next years.”

References

- Ahumada, O.; J.R. Villalobos (2001): Application of planning models in the agri-food supply chain: A review. *European Journal of Operational Research*, Vol. 195, 1-20.
- Jones, P.; G. Kegler; T.J. Lowe; R.D. Traub (2003): Managing the Seed-Corn Supply Chain at Syngenta. *Interfaces*, Vol. 33, 80-90.
- Khouja, M. (1999): The single-period (news-vendor) problem: literature review and suggestions for future research. *Omega*, Vol. 27, 537-553.
- Thode, H. (2002): *Testing For Normality. Statistics: a Series of Textbooks and Monographs (Book 164).* CRC Press, Boca Raton, Florida.

Appendix - Mathematical Model Formulation

In this section, we give a mathematical formulation of the optimization problem of the CO tool. We write $\boldsymbol{a} \cdot \boldsymbol{b} = a_1 b_1 + \cdots + a_J b_J$ for the inner (dot) product and $\boldsymbol{a} \circ \boldsymbol{b} = (a_1 b_1, \ldots, a_J b_J)$ for the Hadamard product of two vectors **a** and **b.** Furthermore, we write $\mathbf{a}^2$ for $\mathbf{a}\cdot\mathbf{a}$ and $\|\boldsymbol{a}\| = \sqrt{\boldsymbol{a} \cdot \boldsymbol{a}}$ for the norm of a vector **a**. **1** is a vector of ones and $[x]^+$ is the positive part of a scalar x. All vectors are row vectors. The basic unit for production and sales is, unless specified otherwise, a standard bag of seed. An overview over the notation is given in the table of symbols at the end of the section.

The objective of the CO tool is to maximize the expected profit of year N+1, $\mathbf{E}_{\text{D,CI,PR}(\mathbf{u})}\Pi(\mathbf{u})$, by deriving the optimal value for the decision variables $\mathbf{u}=(u_1,\ldots,u_J)$, where u_j represents the number of fields to be planted in country j. J represents the total number of countries in which the hybrid is sold or produced. The maximization problem can be written as:

$$\Pi^* = max_{\boldsymbol{u}\geq\boldsymbol{0}}\mathbf{E}_{\text{D,CI,PR}(\mathbf{u})}\Pi(\mathbf{u}), \qquad (1)$$

with respect to

$$\|\boldsymbol{u} \circ \boldsymbol{FS} - (\boldsymbol{u} \cdot \boldsymbol{FS})\boldsymbol{t}\| = min_{\boldsymbol{u}'|\boldsymbol{u}'\cdot\boldsymbol{1}=\boldsymbol{u}\cdot\boldsymbol{1},\boldsymbol{u}'\geq\boldsymbol{0}} \|\boldsymbol{u}' \circ \boldsymbol{FS} - (\boldsymbol{u}' \cdot \boldsymbol{FS})\boldsymbol{t}\|. \qquad (2)$$

Equation (2) ensures that the number of fields are chosen in such a way that the deviation from the target production shares, $\mathbf{t}=(t_1,\ldots,t_J)$, is minimized, where t_j represents the target share of a country's production volume on the total production volume (target production shares sum up to 1). Vector $\mathbf{FS}=(FS_1,\ldots,FS_J)$ denotes the vector of field sizes FS_j of each country j. D, CI, and PR(**u**) are random variables to describe the demand (D), the carry-in (CI), and the production volume (PR) of year N+1.

The expected profit $\Pi(\mathbf{u})$ in Equation (1) is defined as

$$\Pi(\mathbf{u}) = \tilde{p}SV(\boldsymbol{u}) + \tilde{r}\min(CO(\boldsymbol{u}), \mathbf{1}\cdot \boldsymbol{FC}^{N+2}) - \left(\boldsymbol{c}^{Field} + \boldsymbol{c}^{Prod} \circ (\mathbf{1} + \boldsymbol{PB}) \circ \boldsymbol{FS}\right) \cdot \boldsymbol{HF}(\boldsymbol{u}), \quad (3)$$

where $\tilde{p}$ denotes the average net sales price, SV($\mathbf{u}$) the sales volume, $\tilde{r}$ the net value of carry-out supply, CO($\mathbf{u}$) the carry-out volume of year N+1, and $\mathbf{HF}(\mathbf{u})$ a vector of random variables to describe the number of harvested fields per country. The profit function takes into account that Syngenta pays only for harvested fields and not for planted fields. Vectors $\mathbf{FC}^{N+2}$, $\mathbf{c}^{\mathrm{Field}}$, $\mathbf{c}^{\mathrm{Prod}}$, and $\mathbf{PB}$ denote the forecasts (year N+2), the field and production costs, and the production biases per country. Note that all forecast error, production variability, and bias values are given as percentage values of the respective demand forecasts or average production yields.

Total market demand in year N+1 is given by D as the sum of the market demand in the different countries. The distribution of the random variable D is given as the convolution of the demand distributions in the different markets. From the assumption that demand follows a normal distribution in each country, it follows that the total European market demand follows a normal distribution with mean value $(\mathbf{1} + \boldsymbol{FCB}^{N+1}) \cdot \boldsymbol{FC}^{N+1}$ and variance

$$(1-\gamma)\left((\mathbf{1} + \boldsymbol{FCB}^{N+1}) \circ \boldsymbol{FC}^{N+1} \circ \boldsymbol{FCE}^{N+1}\right)^2 + \gamma\mathbf{1}\left(\left((\mathbf{1} + \boldsymbol{FCB}^{N+1}) \circ \boldsymbol{FC}^{N+1} \circ \boldsymbol{FCE}^{N+1}\right)^T \left((\mathbf{1} + \boldsymbol{FCB}^{N+1}) \circ \boldsymbol{FC}^{N+1} \circ \boldsymbol{FCE}^{N+1}\right)\right)\mathbf{1}^T, \quad (4)$$

where $\mathbf{FC}^{\mathrm{N+1}}$, $\mathbf{FCB}^{\mathrm{N+1}}$, and $\mathbf{FCE}^{\mathrm{N+1}}$ are vectors of the forecasting value, bias, and error per country in year N+1 and γ is the correlation factor between market demands. From the fact that allocation is done based on an equal-share principle and based on short-term forecasts that are significantly more precise than the initial forecasts, the average net sales price can be approximated by the forecast-weighted average of the net sales prices,

$$\tilde{p} = \frac{\boldsymbol{p} \cdot \left((\mathbf{1} + \boldsymbol{FCB}^{N+1}) \circ \boldsymbol{FC}^{N+1} \right)}{\mathbf{1} \cdot \left((\mathbf{1} + \boldsymbol{FCB}^{N+1}) \circ \boldsymbol{FC}^{N+1} \right)}, \tag{5}$$

where **p** is the vector of net sales prices per country.

The carry-in supply of year N+1 is given by the random variable CI, which can be written as $CI = (1 - DR)\left([US^N]^+ + PI(\mathbf{1} \cdot \boldsymbol{SU})\right)$, where DR is the discard rate, PI is the factor for process inefficiency, **SU** is the vector of supply that is available in each country at the time the tool is operated, and US^N is the random variable of unsold supply of year N (excluding the part resulting from process inefficiencies). The unsold supply in year N, US^N, is distributed with a normal distribution with mean value

$$\left((1 - PI)\boldsymbol{SU} - (\mathbf{1} + \boldsymbol{FCB}^N) \circ \boldsymbol{FC}^N\right) \cdot \mathbf{1}, \tag{6}$$

and variance

$$\left((\mathbf{1} + \boldsymbol{FCB}^N) \circ \boldsymbol{FC}^N \circ \boldsymbol{FCE}^N\right)^2. \tag{7}$$

The additional supply from production is given by the sum of the random variables for the production volume of each country. The number of successfully harvested fields in country j, $HF_j(u_j)$, follows a binomial distribution, i.e., $HF_j(u_j) \sim Binom(u_j, 1 - FL_j)$, where FL_j is the field loss probability of country j. The distribution of the production yield per country j, PY_j, follows a normal distribution with mean value $(1 + PB_j)FS_jHF_j(u_j)$ and variance

$$\left((1 + PB_j)FS_jHF_j(u_j)\right)^2 \left(\frac{PV_j^2}{HF_j(u_j)} + CPV_j^2\right), \tag{8}$$

and the total production volume is then given by $PR(\boldsymbol{u}) = \boldsymbol{PY} \cdot \mathbf{1}$, where PV_j^2 and CPV_j^2 are the respective standard deviations of the field-level and country-level yield variability.

From the random variables defined above, the missing quantities can be derived. The sales in year N+1 are defined as $SV(\boldsymbol{u}) = \min\left(D, (1 - PI)\big(CI + PR(\boldsymbol{u})\big)\right)$. The carry-out supply of year N+1 is given as the difference between total available supply (carry-in and production) and the sales volume, so that we can write $CO(\boldsymbol{u}) = CI + PR(\boldsymbol{u}) - SV(\boldsymbol{u})$.

The tool uses an iterative search procedure that returns a solution in acceptable time (approx. 2-3 hours for 300 hybrids on a standard PC) to the maximization problem. It can be mathematically shown that the solution returned is the optimal one that maximizes expected profit without violating the constraint of minimizing the deviation from the target production share.

It is straightforward to extend the model to a model that also derives the optimal allocation of the planting fields among the different countries, i.e., a model without the target production share **t**. In this case, the constraint of Equation (2) is dropped and the solution method has to be adapted, taking into account the second-order properties of the objective function.

The Table of Symbols

- $\mathrm{c}_{\mathrm{j}}^{\mathrm{Field}}, \mathbf{c}^{\mathrm{Field}}$ Field production cost in country j, vector of field costs [USD/field]
- $\mathrm{c}_{\mathrm{j}}^{\mathrm{Prod}}, \mathbf{c}^{\mathrm{Prod}}$ Processing cost in country j, vector of processing costs [USD/unit]
- CI Total carry-in supply in year N+1 [units]
- $\mathrm{CO}(\boldsymbol{u})$ Total carry-out supply in year N+1 [units]
- $\mathrm{CPV_j}$ Yield variability on country-level in country j [%]
- D Total European market demand in year N+1 [units]
- DR Discard rate [%]
- $FC_j^N, \boldsymbol{FC}^N$ Forecast for country j in year N, vector of forecasts [units]
- $FCB_j^N, \boldsymbol{FCB}^N$ Forecasting bias for country j in year N, vector of forecasting biases [%]
- $FCE_j^N, \boldsymbol{FCE}^N$ Forecasting error for country j in year N, vector of forecasting errors [%]
- $\mathrm{FL_j}$ Field loss probability in country j [%]
- $\mathrm{FS_j}, \mathbf{FS}$ Field size of country j, vector of field sizes [units]
- γ Correlation factor of market demand
- $\mathrm{HF_j}(u_j), \mathbf{HF}(\mathbf{u})$ Number of harvested fields for country j, vector of harvested fields
- J Number of countries for selling or producing (the particular hybrid)
- $\Pi(\mathbf{u}), \Pi^*$ Profit of year N+1 for given **u**, optimal expected profit of year N+1 [USD]
- $\mathrm{p_j}, \mathbf{p}$ Net sales price in country j, vector of net sales prices [USD/unit]
- $\tilde{\mathrm{p}}$ Demand-weighted average net sales price [USD/unit]
- PI Parameter for process inefficiency [%]
- $\mathrm{PB_j}, \mathbf{PB}$ Production bias in country j, vector of production biases [%]
- $\mathrm{PR}(\mathbf{u})$ Total production yield for given vector **u** [units]
- $\mathrm{PV_j}$ Yield variability on field-level in country j [%]
- $PY_j, \boldsymbol{PY}$ Realized production yield per country, vector of yields [units]
- $\tilde{\mathrm{r}}$ Net value of carry-out supply [USD/unit]
- $SU_j, \boldsymbol{SU}$ Available supply in year N in country j, vector of supply [units]
- $SV(\boldsymbol{u})$ Europe-wide sales volume in year N+1 [units]
- $\mathrm{t_j}, \mathbf{t}$ Target production share for country j, vector of target shares [%]
- $\mathrm{u_j}, \mathbf{u}$ Number of planned fields for country j, vector of field quantities [fields]
- $\mathrm{US^N}$ Unsold supply in year N (excl. supply from process inefficiencies) [units]