Interfaces

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

USG Uses Stochastic Optimization to Lower Distribution Costs

Amy David, David Farr, Ross Januszyk, Urmila Diwekar

Interfaces
Vol. 45, No. 3, May–June 2015, pp. 216–227
ISSN 0092-2102 (print) | ISSN 1526-551X (online)
http://dx.doi.org/10.1287/inte.2014.0786

USG Uses Stochastic Optimization to Lower Distribution Costs

Amy David
Krannert School of Management, Purdue University, West Lafayette, Indiana 47904,
david11@purdue.edu

David Farr, Ross Januszyk
USG, Chicago, Illinois 60661
{dfarr@usg.com, rjanuszyk@usg.com}

Urmila Diwekar
Vishwamitra Research Institute, Crystal Lake, Illinois 60012,
urmila@vri-custom.org

We present a case study of a large-scale stochastic optimization problem for USG, a building supply manufacturer with plants and customers throughout North America. USG seeks to minimize total delivered cost (including production and freight costs) of products in its Durock® product line, subject to capacity constraints and uncertainties in both demand and production costs. We first demonstrate that demand uncertainty, rather than production-cost uncertainty, is the main cause of month-to-month variations in total cost. We then use the chance constraint method to optimize the network, and propagate uncertainty through the cost models, applying a penalty cost for unfulfilled constraints. We show that we can reduce theoretical costs by approximately 4.8 percent by optimizing the network for the 50th percentile of demand, as compared to the base case that uses demand and cost data for a single month. We implemented the new network plan via sourcing rules in both USG's order fulfillment system and Oracle's advanced supply-chain planning module. Several practical delivery concerns limit the benefits realized to an amount less than the theoretical cost reductions, but savings are still considered to be substantial.

Keywords: supply chain planning; optimization; programming: probabilistic.
History: This paper was refereed. Published online in *Articles in Advance* April 3, 2015.

USG is a leading building supplies manufacturer in North America. For its Durock® product line on which this project focuses, it produces several dozen items in three locations in the United States. These products may pass through any one of a network of warehouses before being delivered to customers throughout the United States and Canada (see Figure 1).

Currently, USG makes its production and distribution decisions by using a large-scale linear program (LP). The goal of the LP is to minimize the total delivered cost, which includes production, freight, and handling for all items in the planning network, while meeting the managerial constraints of capacity at each plant and demand at each customer location. The input parameters (production, freight, and handling costs, and demand) reflect a single point in time.

The problem we are addressing is large; the current model used by USG contains approximately 1,500 customer locations and more than 90,000 decision variables; these include dozens of items, 54 warehouse locations, and the selection of rail or truck modes of freight.

Currently, USG solves the LP using commercial software (the INFOR tactical planner) to create an input file of production costs, demands, freight costs, and loading costs, which it then optimizes using IBM's CPLEX. The tactical planner interprets the resulting output and converts it into sourcing rules that link the optimal plan to USG's order fulfillment system (OFS), ensuring that customer orders are placed at the correct plant or warehouse and that warehouse replenishment orders are placed at the correct originating plant. The sourcing rules are also

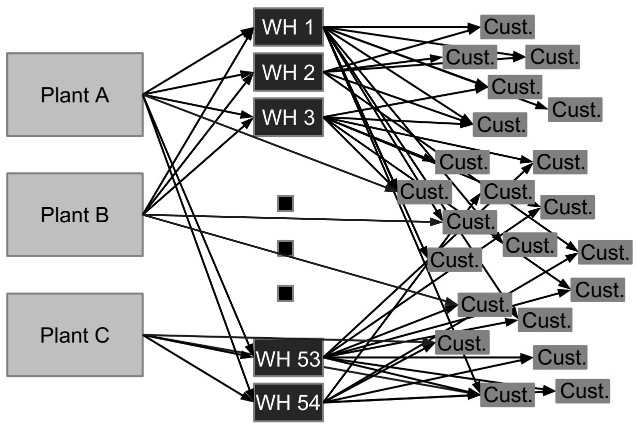

Figure 1: The Durock line includes several dozen items produced in three locations and shipped to more than 1,500 customer locations through a network of 54 warehouses.

used as an input to the advanced supply chain planning (ASCP) module of the Oracle enterprise resource planning (ERP) system, which then uses them to create the plant production schedule.

Therefore, the current planning process does not account for the variability in the historical data that are entered into the optimization software. Most notably, the customer demand at each location changes from month to month; at a location and (or) item level, this can result in almost 6,000 input parameters for which we do not have a known value (i.e., uncertainties). Further, the cost for each item at each production plant may vary, thus adding more than 100 uncertainties to the problem. In USG's current process, one month of data are entered into the plan as though they are deterministic parameters, a method that is unlikely to produce the optimal stochastic solution; therefore, we are seeking a method for optimizing the distribution network that considers these uncertainties.

Problem Definition

We wish to minimize the total expected value of production and distribution of 37 commodities produced at three plants and distributed to customers throughout the United States and Canada. Each commodity may be either shipped directly from a plant to a customer or shipped through any of 54 intermediate warehouses. Outbound shipments to customers use trucks, but shipments from plants to warehouses can use truck or rail mode.

A product shipped directly from a plant to a customer incurs production, truck-loading, and outbound freight costs. A product that passes through an intermediate warehouse incurs costs for production, rail or truck loading, rail or truck freight from the plant to the warehouse, unloading the inbound rail or truck at the warehouse, loading the outbound truck to the customer, and outbound truck freight to the customer. Although the freight and loading and (or) unloading costs incurred are typically greater for a product passing through an intermediate warehouse than for one shipped directly from a plant, the absence of freight lanes (truck or rail routes) from each plant to each customer location sometimes necessitates using this option.

Our objective is therefore to minimize the total cost of production, loading and unloading, and freight by determining which plant or warehouse should ship to each customer, and which plant should ship to each warehouse used to meet customer demand. The problem is subject to three sets of constraints:

1. Demand: demand must be met at each customer location.
2. Capacity: capacity cannot be exceeded at any production location.
3. Warehouse balance: the amount of each product leaving each warehouse must be less than or equal to the amount entering that warehouse.

Approach

To optimize USG's production and distribution network while accounting for demand and cost uncertainty, we first characterize the uncertainty of the historical data. To do so, we use one year of monthly data for customer demand and three years of monthly data for production costs. As we discuss in the *Statistical Fit of Historical Data* section, we fit these data to probability distributions. In the *Problem Formulation* section, we solve the base case to simulate USG's current planning process. We then propagate the demand and cost uncertainties through the base model to determine their relative impacts. Based on our analysis of the historical uncertainty, we determine that the chance constraint method is appropriate for our problem, and use the cumulative probability distributions to convert the deterministic demand constraints into probabilistic constraints. Finally, we sample the cost and demand uncertainties and propagate them through the models obtained using the chance constraint method for several constraint fulfillment levels, applying a penalty for constraint violations to determine the total cost. We present our findings in the *Results* section.

Review of Literature

The problem of optimizing distribution through a supply chain has been long studied. In their seminal work, Geoffrion and Graves (1974) optimize the location of distribution facilities between plants and customers using Bender's decomposition. Thomas and Griffin (1996), Vidal and Goetschalckx (1997), and Klose and Drexl (2005) present reviews of other works that focus on deterministic production and distribution planning.

The benefits of ASCP approaches are best evidenced in implementation cases such as those undertaken by IBM; representative examples include Lin et al. (2000), which describes the development of IBM's asset management tool for strategic and tactical supply chain planning, and Degbotse et al. (2013) in which a large-scale mixed-integer program (Denton et al. 2006) is blended with heuristics to optimize material flow through IBM's semiconductor supply chain. At Heracles General Cement Company, an example in the building supplies industry, which is the focus of our work, the development of a custom supply chain optimization tool lowered sea transportation costs by €6.9 million (Dikos and Spyropoulou 2013).

The uncertain nature of many elements of the supply chain, however, often means that deterministic approaches lead to suboptimal results. Stochastic programming, rather than linear programming, may better capture the realities of a business in which future conditions are uncertain (Sen and Higle 1999). A number of authors have addressed uncertainty in planning a single tier of the supply chain, such as production planning and scheduling or transportation decisions, as Sahinidis (2004) and Mula et al. (2006) review.

More recently, several authors have used stochastic optimization methods to optimize production and distribution decisions in a supply chain under demand uncertainty, similar to the problem in our study. Gupta and Maranas (2003) propose a model to handle demand uncertainty by considering that "the manufacturing decisions are modeled as 'here-and-now' decisions ... the logistics decisions are postponed in a 'wait-and-see' mode" (p. 1). Similarly, Kempf (2004) suggests breaking the problem into two stages: a linear programming formulation that determines the material available and a model predictive-control formulation for the actual fulfillment of demand. In both works, the uncertainty is addressed in the second stage, as in a two-stage problem with recourse; Birge and Louveaux (2011) provide information on this type of problem. Guillén et al. (2005) use a multiobjective two-stage stochastic model to maximize both profit and demand fulfillment, and Santoso et al. (2005) apply a sample average approximation scheme to optimize the supply chain configuration to minimize total production and distribution costs.

Other literature on supply chain uncertainty is motivated by a desire to formulate optimal responses to natural disasters (i.e., the root cause of uncertainty is a single uncertain event that can be modeled as one of many possible scenarios). In an early paper on the topic, Haghani and Oh (1996) present several heuristic methods for solving a multicommodity, multimodal network-flow problem under time constraints. Barbarosoğlu and Arda (2004) develop a two-stage stochastic programming model for distribution of critical goods under uncertainty in both demand and transportation network capacity. Mete and Zabinsky (2010) assume a discrete number of disaster scenarios to solve a storage location and inventory problem with a vehicle routing subproblem. Chang et al. (2007) similarly model both the location and distribution decisions, but assume that the supply chain has several echelons, each containing several entities that must work in cooperation to respond to a disaster.

Another common approach is robust optimization for addressing uncertainty in supply chain planning. Mudchanatongsuk et al. (2007) use robust optimization to solve a network design problem with uncertainty in both demand and transportation costs, and demonstrate its advantages over worst-case scenario planning. Atamtürk and Zhang (2007) propose a robust optimization approach for solving two-stage network-flow problems with uncertain demand. They show that their method is less computationally complex than scenario-based stochastic optimization and allows for parameterization of the conservatism of the solution. Chen and Lee (2004) investigate a three-tier supply chain with uncertainty in both demand and price preferences among the buyers and sellers, and use robustness measures to minimize the degree to which the objective values vary with demand uncertainties.

Our work is most closely related to Tsiakis et al. (2001) who model a multiechelon supply chain network under demand uncertainty, optimizing the number, location, and capacity of warehouses and distribution centers, the transportation links among locations, and the production rates of materials. Although we optimize the production and flow of materials through warehouses in already-fixed locations, the choice among many warehouses in the USG problem can be thought of as analogous to the decision of choosing warehouse locations from a discrete number of options. Similarly, the choice of freight lanes for moving products in our case is similar to determining which transportation links to establish.

Statistical Fit of Historical Data

The data available to us consist of one year of monthly data for the historical demand and 34 months of monthly data for the historical production cost. In this section, we develop probability distributions that fit the data so that we can later simulate a large number of scenarios by drawing samples from those distributions.

Historical Demand

Looking at historical demand, we aggregate data into 162 geographical areas (hereafter referred to as areas) and 37 items. The areas represent states, provinces, or in some cases, parts of states; for example, we divide California into two areas because historical shipments to customers in northern California come from a different warehouse than those to customers in southern California. This gives us 5,994 area-item combinations; however, for the period studied (2012), only 689 of these combinations have sales in at least one month. We therefore heuristically eliminate the 5,305 combinations that have no historical demand, and work to model the demand of the 689 combinations that do.

For combinations with sales in 10 or more months, we model the demand as a truncated normal distribution in three steps:

1. We calculate the mean and standard deviation (μ and σ) of the historical data.
2. To find the bounds within which 99.9 percent of samples will fall, we multiply the standard deviation by 3.08 (taken from a standard normal z table), and subtract this from the mean to determine the lower bound. We add the same value to the mean to determine the upper bound.
3. If the lower bound is negative ($L < 0$), we replace the lower bound with zero, and the upper bound with twice the mean (2μ). This ensures that the demand samples are both nonnegative and centered around the mean.

Using this method, we model 41 combinations with a lower bound greater than zero and 148 combinations as a normal distribution truncated at a lower bound of zero.

For the 500 area-item combinations with sales in nine or fewer months, we model the demand as a binned uniform distribution. For this type of distribution, we specify one or more bins from which samples should be drawn, each defined by a lower and upper bound, and the percentage of samples that should be drawn from each bin. For each combination, the smallest bin has bounds $[0, \epsilon]$, where ϵ is a mathematically insignificant value (i.e., 1.0E–9) meant to ensure that some of the samples taken for this combination are effectively zero.

Of these 500 combinations, 201 have three or fewer distinct values for demand in the 12 months of data studied. We therefore create a separate bin for each value, and then assign a percentage of samples equal to the percentage of historical data points matching that value (see Table 1). For example, if an area-item combination has historical demand of 15 in two months, 30 in one month, and zero in the remaining months, we create three bins: $[0, \epsilon]$, $[15, 15 + \epsilon]$, and $[30, 30 + epsilon]$, and assign them probabilities of 0.750, 0.167, and 0.083, respectively.

For the remaining 299 combinations, each of which has four or more distinct values for demand, we use a zero bin as described previously and a nonzero

Bins (excluding zero bin)	Number of area-item demand combinations
1	400
2	49
3	51

Table 1: Area-item combinations with three or fewer distinct values for demand in the dataset are modeled as a binned uniform distribution with a bin for each distinct value.

bin with lower and upper bounds corresponding to the smallest and largest nonzero demand represented in the 12 months of data studied. As an example, suppose an area-item combination has sales of five in one month, 10 in one month, 15 in one month, 20 in one month, and zero in the remaining eight months. We would model this combination in two bins: $[0, \epsilon]$ and $[5, 20]$, with respective probabilities of 0.667 and 0.333.

Historical Production Cost

Historical production cost data are available for 36 items at three plants, and one additional item at two plants, for a total of 110 item-plant combinations.

As summarized in Table 2, we model 108 of these combinations as normal distributions using 34 months of monthly historical data from 2010 to 2012 to calculate the means and standard deviations. Of these combinations, 76 pass a Pearson's chi-squared goodness-of-fit test for a normal distribution at a p-value of 0.05 using all data points available. Of the remaining 32 combinations, eight pass the test when a single outlier is dropped (and the mean and standard deviations recalculated accordingly), and 15 pass when two outliers are dropped. In the latter case,

Distribution	No. of data points used	No. of item-plant cost combinations
Normal	34	91
Normal	33	8
Normal	32	15
Uniform	34	1
Binned uniform	34	1

Table 2: Item-plant combinations are modeled as normal, uniform, or binned uniform distributions, depending on the shape of the available historical data.

two months of extremely high production costs follow; we assume they are an uncommon production issue that will not recur. The final nine combinations do not pass at the five percent level, but pass for larger p-values, and do not more closely resemble any other common distribution; thus, the normal approximation is still used.

One combination is modeled as a uniform distribution, and one as a binned uniform distribution, as we describe previously. Both pass the goodness-of-fit test for these distributions at the five percent level.

Problem Formulation

The Base Case

We first optimize the network using a single month of data (October 2012, a month in which the total demand was close to the 12-month average) to simulate the current planning process and therefore create a baseline plan against which we can measure other plans. Our objective is to minimize total cost, which consists of the following costs: production, freight from plants to warehouses, freight from plants to geographical areas, freight from warehouses to geographical areas, and handling. For the solution to be feasible, it must meet three sets of constraints:

1. Demand: the amount shipped from all plants to a geographical area plus the amount shipped from all warehouses to a geographical area must be greater than or equal to demand in that area.
2. Capacity: the amount produced at each plant must be less than or equal to the capacity at that plant.
3. Warehouse balance: the amount of product shipped from all plants to a warehouse must be greater than or equal to the amount of product shipped from that warehouse to all geographical areas.

The appendix shows the mathematical formulation.

The Average Demand Case

We next optimize the network by using average costs and demands to quantify the cost of the plan when the deterministic optimization techniques are applied to the average of one year of data. In this case, our problem remains the same as in the base case; however, we replace each production cost with the 12-month average production cost and each demand with the 12-month average demand.

The Impact of Uncertainty

To determine how cost and demand uncertainty affect total cost, we create 10,000 scenarios by sampling the demand and cost distributions (their development is shown in the previous section), and propagate them through the model optimized for average cost and demand, ignoring constraint violations. We generate them using software provided by the Vishwamitra Research Institute to implement a Latin hypercube sampling technique, and complete cost propagation in MATLAB using standard functions.

As an example, suppose the model, optimized for average cost and demand, suggests that customer area j should receive item l from plant i. The sampling step produces 10,000 costs for producing item l at plant i, and 10,000 demands for item i in customer area j. Therefore, for each sample scenario, we can multiply the production and freight costs times the demand and determine the cost of providing item i to customer j. The total network cost for each sample scenario is the summation of each combination of customer area-item.

We first propagate both demand and production cost uncertainty, repeat the process for demand only, and then repeat it for production cost only. To protect the confidentiality of USG's cost data, we represent the average total cost in the base case optimization as \$BC, and show means in terms of a percent increase or decrease from this value and standard deviations as a percentage of this value. For the case in which costs and demand constraints are based on average demands, we find an average total cost of \$BC + 0.3 percent. As Table 3 shows, neither demand nor production-cost uncertainty affects the mean total cost, and the demand uncertainty accounts for most of the variance in total cost: the standard deviation propagating only demand uncertainty through the cost model is more than 97.5 percent of the standard deviation when both demand and production-cost uncertainty are used. We can therefore justify using an optimization method that focuses primarily on finding a best solution with respect to demand uncertainty.

Uncertainties	Mean	Standard deviation
Demand and production cost	\$BC + 0.3%	\$208,900
Demand	\$BC + 0.3%	\$203,700
Production cost	\$BC + 0.3%	\$41,500

Table 3: Standard deviations of the total cost under the propagation of demand and cost uncertainties make it evident that demand uncertainty has the largest effect on total cost uncertainty.

The Chance Constraint Method

To account for uncertainty in the demands, we use the chance constraint method, as first proposed by Charnes et al. (1958). In this method, we seek a solution to our optimization problem that will allow each demand constraint to be met a specific percentage of the time, once implemented. For example, if we choose a 75 percent chance of meeting demand for a single item in Montana, we would expect that in any given year, the planned amount of product distributed to Montana would meet or exceed the demand in nine of 12 months, and fall short in the other three months. Customers in Montana would still receive the product required in the other three months (provided that total network demand is less than total network capacity, an assumption we address next), but from a nonoptimal location, and thus at a higher cost.

We therefore divide our calculation of total cost into two phases:

Phase 1. We optimize the plan assuming a certain probability (or chance) p that all demand constraints will be met.

Phase 2. We sample the demand distributions using a Latin hypercube sampling technique (McKay et al. 1979) and propagate the uncertainty through the model developed in Phase 1, applying a cost penalty each time a constraint is violated.

The problem in our study lends itself to this method because of the zero demands common in our historical data set. Of the 689 area-item combinations with sales in 2012, 201 have no sales in at least nine of 12 months, and 367 have no sales in at least six of 12 months. Therefore, a chance constraint formulation allows us to consider what happens when we focus the optimization on the most likely demands, and decrease the influence of less likely occurrences (i.e., sales in area-item combinations that occur with low frequency).

The chance constraint method also enables the conversion of the stochastic program into a standard LP.

Given that USG already has the software to solve a large-scale LP, using a method such as this one, which will allow the company to do future planning of the network using existing capabilities, is advantageous.

To optimize the plan under chance constraints (Phase 1), we begin with the base case formulation. For each demand constraint, we take the cumulative probability distribution developed earlier and locate the pth percentile. We then replace the right side of the demand constraint with this value, thus ensuring that demand will be fulfilled with a probability of p. Because we have monthly data for demand, we begin with $p = 6/12 = 0.5$, and then later repeat the chance constraint optimization for other fractions of 12: 0.58, 0.67, 0.75, 0.83, 0.92, and 1. We formalize this mathematically in the appendix.

For example, to model the 83rd percentile of demand, we used the third-largest demand value of the 12 available. Figure 2 shows an example. By using the 83rd percentile of demand, we assume that if USG implements this optimized plan, 83 percent of the demand constraints would be met each month.

In the second phase, we sample 10,000 data points from the demand and cost distributions using a Latin hypercube sampling scheme, and propagate each sample through the optimized model. Because we assume demands in many of the area-item combinations to be zero in Phase 1, if they take a nonzero value when sampled, we assume that the product shipped from the location found in the optimization with average demand constraints, provided capacity was available at that location, and we apply the appropriate cost. If capacity was not available at that location, we use the average cost from the other two production locations, and add the average freight per unit cost of the entire optimized plan as the penalty for shipping from a nonprimary location. The penalty value is based on calculations provided by USG for shipments from nonoptimized locations.

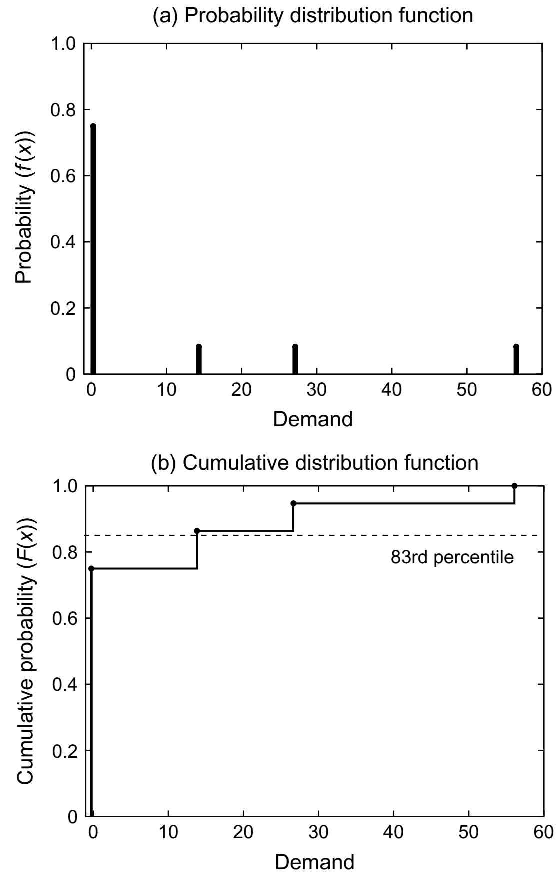

Figure 2: The probability distribution function and cumulative distribution function of a sample area-item combination show that the 83rd percentile of demand for this combination has a value of 14.

Results

We model all problems described in the previous section as LPs in GAMS, and solve them using the IBM CPLEX engine. Computational time for each is under five minutes. Again, we represent the average total cost in the base case optimization as \$BC, and all other values in terms of a percentage increase or decrease from this value.

Table 4 gives the descriptive statistics for monthly cost using a chance constraint-optimized model applied to 10,000 samples of uncertain parameters, and with constraint violation penalties applied. Because many of the area-item combinations have sales in fewer than 12 months, we also provide the number of nonzero demand combinations. Note that the 92nd and 100th percentiles could not be optimized in this manner, because for those percentiles, the total network demand exceeds total network capacity.

Optimizing at the 50th percentile and applying the penalty for violated constraints gives us a 4.86 percent improvement over the average, the lowest cost of all the demand schemas used. This indicates that USG should use the median demands when optimizing the

Percentile	Mean	Std. dev.	Avg. no. of violations
Base case	\$BC	6.05% of mean	262
Average	\$BC − 1.37%	5.70% of mean	256
42nd	\$BC − 4.58%	7.07% of mean	307
50th	\$BC − 4.86%	6.84% of mean	295
58th	\$BC − 4.35%	6.28% of mean	287
67th	\$BC − 1.37%	5.75% of mean	264
75th	\$BC + 9.08%	4.30% of mean	251
83rd	\$BC + 29.07%	3.15% of mean	230

Table 4: Descriptive statistics for the expected monthly total cost with constraint violation penalty show that the 50th percentile of demand gives the greatest cost savings.

network, rather than average demands or demands from a single month, as it does now.

When we propagate uncertainty through the model optimized at the 50th percentile of demand, more constraints are violated and then penalized by a higher cost. Therefore, as expected, the standard deviation is larger when optimizing at the 50th percentile than at any of the higher percentiles. Because the constraints violated differ from sample to sample across the 10,000 used, the results at the 50th percentile show more variance than those at higher percentiles. Therefore, we note the trade-off between the expected lowest total cost and the risk of a high total cost in any individual month. For comparison, at the 58th percentile, we see a 4.4 percent improvement over the base case; however, the standard deviation is almost eight percent smaller. Although optimization at the 50th percentile gives us the lowest total cost, it forces us to accept a greater risk of high network costs in any given month.

Looking at the capacity constraints, only one is active in the initial optimization at the 50th percentile, indicating that all capacity will be used at the lowest-cost plant (LCP). In many cases, distributing product from the LCP through an intermediate plant to a customer is more cost-effective than servicing that customer directly from a more expensive plant. Qualitatively, this is the main difference between the solution from the base case and from that of the 50th percentile optimization: the latter ensures that the capacity at the LCP is utilized better. However, when uncertainty is propagated through the base case model, the full capacity at the LCP is used 87 percent of the time; through the 50th percentile model this increases to 96 percent.

Value of the Stochastic Solution

The value of the stochastic solution is the difference between the solution obtained by using average values for each of the uncertain inputs and the solution obtained by considering uncertainty in our choice of modeling methods (Birge 1982). In this case, it is the difference between the base case solution and the solution obtained using the chance constraint method at the 50th percentile. Therefore, we calculate the value of the stochastic solution to be 4.86 percent.

This indicates that the stochastic solution is valuable to USG if it can be executed and implemented for less than 4.86 percent of the total monthly cost. As previously mentioned, the chance constraint method uses USG's existing software capabilities; thus, the execution and implementation costs are very small.

Expected Value of Perfect Information

The expected value of perfect information (EVPI) is the difference between the stochastic solution and the total cost if all costs and demands are known prior to optimizing the network (Howard 1966). Optimizing the network individually for each month with available demand data gives us an average cost approximately 3.61 percent less than the expected cost when optimizing it at the 50th percentile. This indicates that a perfect forecast of both production costs and demands would save the company 3.61 percent over the stochastic solution.

We can interpret this to mean that USG could save up to 3.61 percent by reducing the uncertainty in production costs and demands, either through an actual reduction in the month-to-month variability of production costs and demands or through improved forecasting methods that would allow network optimization based on narrower sampling distributions. For example, production cost variability may be reduced through stricter process control and demand variability may be reduced by using safety stock inventories. However, we must weigh the cost of investing in these improvements against the EVPI, which we see offers less than a 3.61 percent improvement over the stochastic solution.

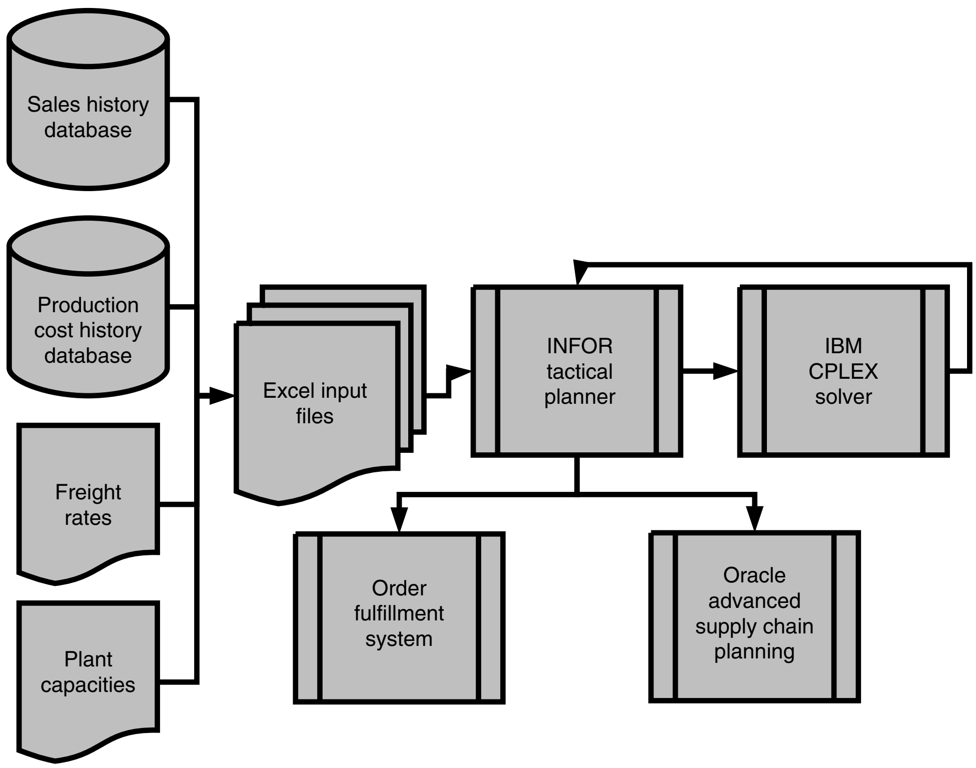

Figure 3: Information flows through various systems in the semiannual sourcing rule creation process.

Implementation

USG implements its network plan through a set of sourcing rules in its OFS. When a customer places an order, the OFS places the order at the plant or warehouse location dictated by the rule. Therefore, it converts the output of the CPLEX solver into a table of sourcing rules by creating a sourcing rule any time the decision variable has a positive value, and no sourcing rule if the decision variable is equal to zero (i.e., if the decision variable representing the volume of Durock products shipped from plant i to the customer area j is positive, we create a sourcing rule in OFS so that Durock orders placed by customers in area j would be shipped from plant i).

In parallel, the sourcing rules are uploaded into the Oracle ASCP software (Oracle 2014) that uses this information, and inventory, forecast, and capacity information to create the master schedule for purchasing and production. The existing sourcing rule creation process is shown in Figure 3.

Going forward, a member of USG's network optimization planning team will generate the sourcing rules semiannually as follows:

Step 1. Find the 50th percentile demands.

(a) Compile historical data on sales for each customer area-item using monthly buckets within an Excel spreadsheet.

(b) Sort the monthly demands from high to low.

(c) Choose the seventh-largest monthly demand, representing the 50th percentile of demand.

Step 2. Find the average production costs:

(a) Compile historical data on production costs by plant and item using monthly buckets within an Excel spreadsheet.

(b) Calculate the average (mean) demands.

Step 3. Enter the following problem parameters into the INFOR tactical planner software via the Excel spreadsheet:

- Demands found in Step 1.
- Average production costs found in Step 2.
- Current freight rates.
- Current plant production capacities.

Step 4. Run the INFOR tactical planner to create a CPLEX input file.

Step 5. Upload the CPLEX input file to the CPLEX engine and initiate the LP solver, producing a CPLEX output file.

Step 6. Upload the CPLEX output file to the INFOR tactical planner to view output in a graphical user interface; download the file in a comma-separated values (CSV) format.

Step 7. Make changes to any sourcing rules in the OFS program via a CSV upload.

Step 8. Make changes to any sourcing rules in the Oracle ASCP via a CSV upload.

Note that Steps 4–8 remain unchanged from the original process, and the only difference in Step 3 is the use of 50th percentile demands and average costs, rather than data from the previous month. Therefore, the only additional work is in Steps 1 and 2, both of which involve a larger data query and some brief spreadsheet manipulation that the original process did not include. We estimate the additional workload to be negligible—less than one full day annually by a member of the network optimization planning team.

As we discussed in the *Results* section, the theoretical results suggest potential average savings of 4.9 percent in each month. To achieve these savings, USG began work to implement the new sourcing rules in January 2014. In implementing the sourcing rules, however, USG found several limitations that prevent it from obtaining the entire cost savings. The largest limitation is the practice of shipping Durock orders in less-than-full truckload quantities. Only about half of Durock orders ship in full-truckload quantities, with the remaining orders shipping in mixed truckloads with items from other product lines. Typically, the volume of the other product line exceeds that of the Durock products; therefore, the other product line's sourcing rules supersede those of the Durock product line.

Other customer service limitations also prevent orders from shipping from the optimal plant; for example, customers who pick up their products at a specific warehouse select the warehouse. Taking all these limiting factors into consideration, the average monthly savings are somewhat less than the theoretical amount, with some variation in the savings amount from month to month. Although the actual savings are only about one-third of the theoretical savings, they still represent a substantial annual cost reduction to USG.

Conclusions

We sought to incorporate the uncertainty of production costs and demand into the optimization of a distribution network to find the lowest-possible cost of servicing 162 geographical areas from three production sites, either directly or through one of 54 intermediate warehouses. To do so, we first fit the historical data for each of 800 uncertainties to either a normal or binned uniform distribution. We then used the chance constraint method to convert probabilistic demand constraints to deterministic demand constraints based on their cumulative distribution functions. We propagated uncertainty through the chance constrained model and applied a penalty for each constraint violation. Optimizing the network for the 50th percentile of demand with constraint violation penalties provided the lowest network cost—a 4.8 percent theoretical improvement over the base case, in which we used one month of demand. As implemented, the savings are smaller than the theoretical amount; however, they still represent a significant monthly cost reduction compared to the base case.

Acknowledgments

The authors thank two anonymous reviewers for their feedback in improving our manuscript. This research was made possible through a 2013 grant from USG.

Appendix. Mathematical Formulations

Mathematical formulation of base case problem: where:

x_{ijl}: amount of item l shipped from plant i to customer j via truck;
x^t_{ikl}: amount of item l shipped from plant i to warehouse k via truck;
x^r_{ikl}: amount of item l shipped from plant i to warehouse k via rail;
x^w_{kjl}: amount of item l shipped from warehouse k to customer j via truck;
p_{il}: cost of producing item l at plant i;
l^t_i: cost of loading a truck at plant i;
lr_i cost of loading a rail at plant i;
n_t: units per truck;
n_r: units per rail;
f_{ij}: truck freight per unit from plant i to customer j;
f^t_{ik}: truck freight per unit from plant i to warehouse k;

f_{ik}^r: rail freight per unit from plant i to warehouse k
f_{kj}^w: truck freight per unit from warehouse k to customer j.

Mathematical formulation of chance constraints:

$$P\left(d_{jl} \geq \sum_i x_{ijl} + \sum_k x_{kjl}\right) \geq p \quad \forall j, l,$$

$$\sum_i x_{ijl} + \sum_k x_{kjl} \geq F_{d_{jl}}^{-1}(p) \quad \forall j, l,$$

where $F_{d_j}^{-1}$ is the inverse cumulative probability distribution of demand at location j.

Mathematical formulation of chance constraint problem:

$$\begin{aligned}\text{Minimize } E\Bigg[&\sum_i\sum_j\sum_l\left(p_{il}+f_{ij}+\frac{l_i}{n_t}\right)x_{ijl}\\&+\sum_i\sum_k\sum_l\left(p_{il}+ft_{ik}+\frac{l_i}{n_t}\right)xt_{ikl}\\&+\sum_i\sum_k\sum_l\left(p_{il}+fr_{ikl}+\frac{lr_i}{n_r}\right)xr_{ikl}\\&+\sum_j\sum_k\sum_l\left(fw_{kj}+\frac{l_i}{n_t}\right)x_{kjl}\Bigg]\end{aligned}$$

$$\begin{aligned}\text{subject to: } &\sum_i x_{ijl} + \sum_k x_{kjl} \geq F_{d_{jl}}^{-1}(p) \quad \forall j, l,\\&\sum_j x_{ijl} + \sum_k (xt_{ikl} + xr_{ikl}) \leq c_i \quad \forall i, l,\\&\sum_i (xt_{ikl} + xr_{ikl}) \leq \sum_j xw_{kjl} \quad \forall k, l,\end{aligned}$$

where p is the probability that each demand constraint is met.

References

Atamtürk A, Zhang M (2007) Two-stage robust network flow and design under demand uncertainty. *Oper. Res.* 55(4):662–673.

Barbarosoğlu G, Arda Y (2004) A two-stage stochastic programming framework for transportation planning in disaster response. *J. Oper. Res. Soc.* 55(1):43–53.

Birge JR (1982) The value of the stochastic solution in stochastic linear programs with fixed recourse. *Math. Programming* 24(1):314–325.

Birge JR, Louveaux F (2011) *Introduction to Stochastic Programming* (Springer, New York).

Chang M-S, Tseng Y-L, Chen J-W (2007) A scenario planning approach for the flood emergency logistics preparation problem under uncertainty. *Transportation Res. Part E* 43(6):737–754.

Charnes A, Cooper WW, Symonds GH (1958) Cost horizons and certainty equivalents: An approach to stochastic programming of heating oil. *Management Sci.* 4(3):235–263.

Chen C-L, Lee W-C (2004) Multi-objective optimization of multi-echelon supply chain networks with uncertain product demands and prices. *Comput. Chemical Engrg.* 28(6):1131–1144.

Degbotse A, Denton BT, Fordyce K, Milne RJ, Orzell R, Wang C-T (2013) IBM blends heuristics and optimization to plan its semiconductor supply chain. *Interfaces* 43(2):130–141.

Denton BT, Forrest J, Milne RJ (2006) IBM solves a mixed-integer program to optimize its semiconductor supply chain. *Interfaces* 36(5):386–399.

Dikos G, Spyropoulou S (2013) Supply chain optimization and planning in Heracles General Cement Company. *Interfaces* 43(4): 297–312.

Geoffrion AM, Graves GW (1974) Multicommodity distribution system design by Benders decomposition. *Management Sci.* 20(5):822–844.

Guillén G, Mele FD, Bagajewicz MJ, Espuna A, Puigjaner L (2005) Multiobjective supply chain design under uncertainty. *Chemical Engrg. Sci.* 60(6):1535–1553.

Gupta A, Maranas CD (2003) Managing demand uncertainty in supply chain planning. *Comput. Chemical Engrg.* 27(8):1219–1227.

Haghani A, Oh S-C (1996) Formulation and solution of a multi-commodity, multi-modal network flow model for disaster relief operations. *Transportation Res. Part A* 30(3):231–250.

Howard RA (1966) Information value theory. *Systems Sci. Cybernetics IEEE Trans.* 2(1):22–26.

Kempf KG (2004) Control-oriented approaches to supply chain management in semiconductor manufacturing. *Proc. 2004 Amer. Control Conf.* (Institute of Electrical and Electronics Engineers, Washington, DC), 4563–4576.

Klose A, Drexl A (2005) Facility location models for distribution system design. *Eur. J. Oper. Res.* 162(1):4–29.

Lin G, Ettl M, Buckley S, Bagchi S, Yao DD, Naccarato BL, Allan R, Kim K, Koenig L (2000) Extended-enterprise supply-chain management at IBM Personal Systems Group and other divisions. *Interfaces* 30(1):7–25.

McKay MD, Beckman RJ, Conover WJ (1979) Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 21(2):239–245.

Mete HO, Zabinsky ZB (2010) Stochastic optimization of medical supply location and distribution in disaster management. *Internat. J. Production Econom.* 126(1):76–84.

Mudchanatongsuk S, Ordóñez F, Liu J (2007) Robust solutions for network design under transportation cost and demand uncertainty. *J. Oper. Res. Soc.* 59(5):652–662.

Mula J, Poler R, Garcia-Sabater JP, Lario FC (2006) Models for production planning under uncertainty: A review. *Internat. J. Production Econom.* 103(1):271–285.

Oracle (2014) Advanced supply chain planning. Accessed October 1, 2014, http://www.oracle.com/us/products/applications/ebusiness/scm/051323.html.

Sahinidis NV (2004) Optimization under uncertainty: State-of-the-art and opportunities. *Comput. Chemical Engrg.* 28(6):971–983.

Santoso T, Ahmed S, Goetschalckx M, Shapiro A (2005)A stochastic programming approach for supply chain network design under uncertainty. *Eur. J. Oper. Res.* 167(1):96–115.

Sen S, Higle JL (1999) An introductory tutorial on stochastic linear programming models. *Interfaces* 29(2):33–61.

Thomas DJ, Griffin PM (1996) Coordinated supply chain management. *Eur. J. Oper. Res.* 94(1):1–15.

Tsiakis P, Shah N, Pantelides CC (2001) Design of multi-echelon supply chain networks under demand uncertainty. *Indust. Engrg. Chemistry Res.* 40(16):3585–3604.

Vidal CJ, Goetschalckx M (1997) Strategic production-distribution models: A critical review with emphasis on global supply chain models. *Eur. J. Oper. Res.* 98(1):1–18.

Verification Letter

Tim McVittie, Sr. Director, Global Logistics, USG Building Systems, 550 West Adams Street, Chicago, Illinois 60661, writes:

"As Sr. Director, Global Logistics at USG Building Systems, I hereby certify that the project described in the attached manuscript, "USG Uses Stochastic Optimization to Lower Distribution Costs," was conducted by my department in cooperation with Amy David and Urmila Diwekar of University of Illinois at Chicago, and all information contained therein is accurate to the best of my knowledge."

Amy David joined Purdue University's Krannert School of Management as a clinical assistant professor in August 2014, after completing her PhD in industrial engineering and operations research at the University of Illinois at Chicago. She teaches the faculty-directed student project course, as well as undergraduate and graduate courses in production planning and control and supply chain management. Previously, Amy worked as logistics planning and process development manager at USG and a logistics analyst at Medline Industries. She holds an MBA from Lake Forest Graduate School of Management, a BS in general engineering from the University of Illinois at Urbana-Champaign, and the APICS CPIM designation.

David Farr is manager, logistics planning at USG Corporation. He is responsible for leading network optimization planning for USG's North American operations. He has worked at USG for more than 30 years in various assignments. He holds an MBA from Kennesaw State University and a BS in accounting from the State University of New York at Geneseo.

Ross Januszyk is senior manager, logistics planning at USG Corporation. He is responsible for leading sales and operations planning, demand planning, and network optimization planning for USG's North American operations. Prior to joining USG, he was a management consultant with PRTM and he started his career as an electrical engineer at Motorola. He holds an MBA from University of Michigan and a BS in electrical engineering from University of Illinois at Urbana–Champaign.

Urmila Diwekar, PhD, is currently president of the Vishwamitra Research Institute (VRI, www.vri-custom.org), a nonprofit research organization that she founded in 2004 to pursue multidisciplinary research in the areas of optimization under uncertainty and computer-aided design applied to energy, environment, and sustainability. She is also affiliated with Department of Bio Engineering and Industrial Engineering of University of Illinois at Chicago. From 2002–2004, she was a professor in the Departments of Chemical Engineering, Bio Engineering, and Industrial Engineering, and in the Institute for Environmental Science and Policy, at the University of Illinois at Chicago (UIC). From 1991–2002 she was on the faculty at Carnegie Mellon University, with early promotions to both the associate and the full professor level.