# QUANTIFYING SUPPLY CHAIN DISRUPTION RISK
# USING MONTE CARLO AND DISCRETE-EVENT SIMULATION

Amanda J. Schmitt
Mahender Singh

Center for Transportation and Logistics
Massachusetts Institute of Technology
Cambridge, MA 02142, USA

## ABSTRACT

We present a model constructed for a large consumer products company to assess their vulnerability to disruption risk and quantify its impact on customer service. Risk profiles for the locations and connections in the supply chain are developed using Monte Carlo simulation, and the flow of material and network interactions are modeled using discrete-event simulation. Capturing both the risk profiles and material flow with simulation allows for a clear view of the impact of disruptions on the system. We also model various strategies for coping with the risk in the system in order to maintain product availability to the customer. We discuss the dynamic nature of risk in the network and the importance of proactive planning to mitigate and recover from disruptions.

## 1    INTRODUCTION

As supply chains expand globally, their risk of disruption also grows. Catastrophic events, such as Hurricane Katrina, have brought focus to the analysis of disruption risk and development of mitigation plans to cope with that risk. In this paper, we model a complex supply chain using simulation software and analyze the impact of disruptions within the network. We use both discrete-event simulation (Arena) and Monte Carlo simulation (@Risk) to capture the behavior of the system and model the disruption risk. Combining these approaches allows us to evaluate the disruption risk in the entire network. The model presented was developed for a large consumer-products company to assess their level of risk in terms of the impact disruptions could have on product availability. We discuss the insights gained in developing the model and analyzing the system.

While having a large, complex supply chain may make a company more susceptible to disruptions, it can also act as an advantage if the network is utilized to mitigate those disruptions. It can mean the supply chain has more ability to adapt when a disruption occurs. A small supply chain with a single manufacturing plant may have no recourse if that plant is shut down. However, in a large supply chain, if a firm has invested in back-up capabilities or redundancy across its network, then one location may have the ability to back another up if the first is disrupted.

In this project, we consider a multi-echelon supply chain where disruptions may occur at any stage. Due to the complexity of the interactions in a such a system, especially when multiple mitigation tactics are considered, simulation was deemed the best modeling approach for this project. Jain and Leong (2005) discuss the advantages of constructing supply chain simulations to "stress test" a system to see how it performs under extreme conditions. Since disruptions can certainly be classified as potential extreme conditions, this aligns with our intent of testing the system when they occur to see what the impact is to the customer.

The sponsoring company (hereafter referred to as company ABC) is concerned with maintaining very high levels of customer service, or case fill rate, even if a disruption occurs. They are a consumer packaged goods (CPG) company that considers brand loyalty its key success factor. As a result, maintaining product availability is critical which in turn drives higher inventory levels. At the same time, being a CPG company, product margins are small and inventory costs must be minimized. Thus the company is torn between these conflicting objectives. Additionally, the customers of ABC inject another layer of complexity into the model, since their behavior in the face of a stock-out is not known with certainty. Based on past experience, the customer response in case of a stock-out can range anywhere from total lost-sales to total back-ordering

of demand. A simulation model allows us to effectively test these various customer behaviors while evaluating the trade-off between service level and inventory investment. To this end, we designed and tested a model that incorporated multiple complex issues in ABCs supply chain and provided meaningful results. Specifically, the project was designed to address the following objectives:

1. Assess the current level of supply chain disruption risk in the system
2. Test different mitigation strategies
3. Provide a tool that can be used in the case of a disruption to validate recovery steps before putting them into action
4. Identify redundancy in the system that can be removed without affecting the risk level

While the output of the first objective will be focused on operations-level performance, the remaining objectives focus on strategic-level decision making to mitigate disruptions. Tomlin (2006) discusses various strategies for coping with disruptions at a single facility: inventory control (proactive ordering and stocking to cover a potential shortage), sourcing (contingency strategies for backing up specific suppliers or links in the supply chain), and acceptance (doing nothing because the cost of mitigation outweighs the benefits). We designed the simulation to test each of these strategies and view the impact on customer service.

Literature on supply chain disruptions is growing, but applications in multi-echelon supply chains are scarce due to the complexity of those systems. Snyder and Shen (2006) develop simulation models for several networks comparing stochastic demand and supply disruptions. They conclude that the two different sources of stochasticity have very different impacts on optimal supply chain design. Schmitt et al. (2008) expand on this, developing an analytical model comparing inventory placement policies in a distribution system. They show that while stochastic demand provides risk pooling benefits (allowing for lower costs in a centralized inventory setting), risk diversification provides lower cost variance in a decentralized setting and generally protects better against disruptions. Hopp and Yin (2006) examine an assembly system with the potential for disruptions in capacity, and discuss capacity protection and inventory as mitigation methods. Their numerical studies show that increasing disruption occurrence or variability upstream tends to shift the location of protection upstream. Our model incorporates the systems considered in most of these models; we have individual locations that can use inventory for mitigation, as well as back-up location options. The system itself consists of both an assembly and a distribution system, with inventory held throughout the network.

Classification of risks in a supply chain and development of risk profiles is also a growing research trend. For example, Deleris and Erhun (2005) present a Monte Carlo simulation that they use to evaluate risk levels in the supply chain. Genentech, a global biopharmaceutical company, also uses Monte Carlo simulation to assess their network risk. Steckel (2008) discusses work performed at Genentech to quantify their disruption risk and make inventory-stocking decisions. They used insurance data and peer inputs to estimate the frequency and duration of disruptions, and developed a Monte Carlo simulation to estimate the total number of weeks out of a year that they could potentially be disrupted. They then set inventory levels accordingly so that they can maintain material flow using inventory even if disruptions occur. Part of our project also entails a Monte Carlo simulation to quantify the risk in the supply chain that we study, with the objective of developing distributions for input to a large-scale discrete-event simulation that captures the entire network flow.

The steps for this modeling project were broken out as follows:

1. Background Analysis
2. Base Model Development (Arena)
3. Risk Profile Model Development (@Risk and Arena)
4. Risk Simulation Model Development (Arena)
5. Risk Quantification Testing

Step 1 is discussed in Section 2 below. Steps 2 and 3 were performed concurrently and are discussed in Sections 3 and 4. The results of the Base Model and Risk Profile Model are combined in Step 4, which we discuss in Section 5. We present results achieved for Step 5 in Section 6 and summarize our observations in Section 7.

## 2 BACKGROUND ANALYSIS

### 2.1 Project Scope

The scope of the project was clearly defined with company ABC prior to initiating the project. The company team members chose a product family that they felt was representative, since it has a supply chain that captures the complexity of their
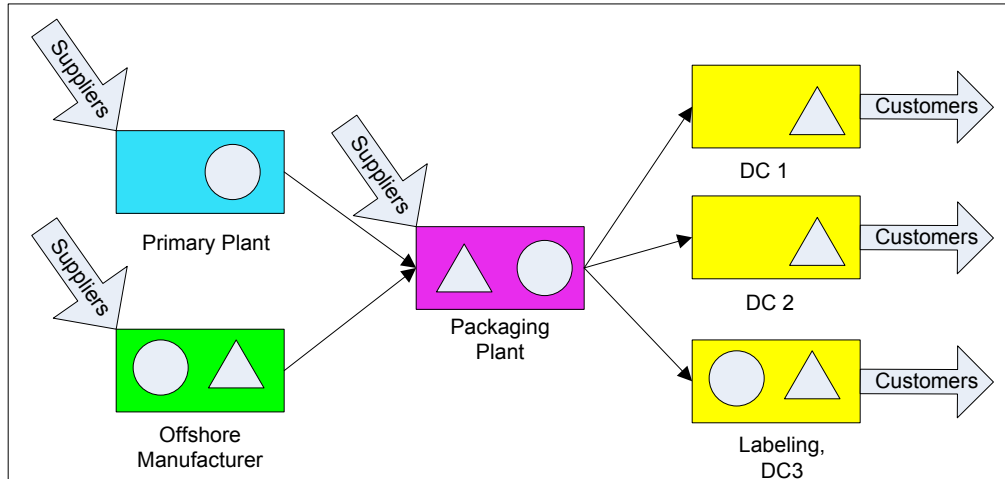
Figure 1: Supply Chain Flowchart. Circles represent processes, and triangles represent inventory locations.

network but is not too large in size. One deliverable of the project is the actual final simulation model, so that ABC can extend this analysis to other product lines if they see fit. Within the product family, two separate products were to be modeled, each with their own inventories and demand streams. We will refer to these as the high-volume and low-volume products in the future, since the low-volume product's sales are approximately 12% of the high-volume product's sales. Additionally, the scope was restricted to North America's market only, though the product is sold worldwide; this limited the supply chain scope mostly to US manufacturing and distribution sites, but one off-shore contract manufacturer is still part of the supply chain.

To evaluate system performance, ABC decided to focus on customer service levels (i.e., case fill rate and/or missed cases). While costs are considered relevant and can be discussed generally when comparing different mitigation strategies, they are not specifically within the scope of the supply chain model. Ultimately, ABC wants to assess the impact of disruptions on their customer service.

## 2.2 Product Supply Chain

Once conversations were begun with ABC experts on the product family, the supply chain flow took shape. Figure 1 presents the flow of the system to be modeled. The figure uses circles and triangles to represent operations and inventory locations, respectively.

Most raw materials are single-sourced. Each product consists primarily of two parts that are assembled and packaged at an outsourced packaging plant. For the high-volume product, both of those parts are manufactured at the primary plant; for the low-volume product, one of the parts is manufactured at the primary plant, and the other is manufactured at the off-shore contract manufacturer and then shipped to the U.S. Both products are distributed through all three DCs (distribution centers).

Large quantities of inventory are held at each of the locations labeled in Figure 1. According to ABC production planners, the target on-hand WIP (work in process) inventory levels at the packaging plant are approximately 2 weeks' worth of demand, and at the DCs are 4 weeks' worth of demand. The off-shore manufacturer carries an additional 6 weeks' worth of demand in post-processed WIP. ABC justifies carrying so much inventory because they absolutely prioritize fill rate to the customer and will not risk running out of material. However, they indicate that these levels may vary throughout the year; for example, they could be increased if launch of another product is planned which would constrain manufacturing capability, or could be lowered if lean initiatives or accounting concerns drive inventory-reduction incentives.

## 3 BASE MODEL

The base model was developed to capture the existing flow (without disruptions) as depicted in Figure 1. The intent was to successfully model the system as it currently functions and validate the simulation model using actual data. Once the current network was successfully simulated and validated, potential disruptions and proposed mitigation tactics can be added to the model; we present that extended model in Section 5.

As the model was constructed, constant communication was maintained with ABC team members in order to check the validity of assumptions and make the best decisions for the project. We focused on capturing the level of detail needed to make strategic decisions in the system, without overloading the model with operational factors. For example, the model was built so that demand and production occur on a weekly basis; the team agreed that time increments smaller than weeks were not necessary and would not add value to the output.

## 3.1 Demand Fitting and Structure

We received 8 months' worth of demand data for the two products at each of the DCs, and discussed using empirical distribution or fitting theoretical distributions to the demand with the ABC team. We decided that theoretical distributions would allow more flexibility with changing the mean or other parameters to see their impact, so we fit the demand using Arena's Input Analyzer program. The majority of the demand is reasonably fit using a normal distribution, though some was extremely variable; for example, demand for the low-volume product out of one of the DCs was best fit by an exponential distribution.

For the purposes of the model, we also needed to know how to handle demand that could not be satisfied from stock. Since the actual case fill rate is normally very high, ABC did not have a clear answer for how missed cases are handled. Based on discussion with operations and supply chain personnel, the team generally felt that most demand would be lost if it could not be satisfied from stock, but that some customers might have a tolerance period that they would wait before canceling an order. Thus we decided to build flexibility of customer behavior into the model, including the options for partial lost sales and partial back-ordering in the following manner: When demand arrives at a DC, any quantity that can be filled from stock is immediately filled. For the remaining amount that cannot be filled, a certain percentage (call this $X_1\%$) is lost immediately. The rest is back-ordered for up to a fixed number of weeks (call this $T$), after which a second percentage ($X_2\%$) is lost. The remainder is back-ordered indefinitely. These three variables ($X_1$, $X_2$, and $T$) are inputs to the model that can be varied and tested for sensitivity. This modeling approach gives ABC the flexibility to test the impact of customer behavior on case fill rate.

## 3.2 Model Structure

Several inputs are required to run the model, and we constructed an Excel spreadsheet to read them into Arena so that someone who is not very familiar with Arena can still easily modify and test input sensitivity. The inputs for the base model are listed in Table 1.

Table 1: Base Model User Inputs

| *Category* | *Input* |
|---|---|
| Distribution Centers | Demand distribution information |
| | Number of weeks' worth of finished goods inventory held |
| Customers | Immediate and secondary lost sales percentages ($X_1$ and $X_2$) |
| | Waiting time tolerance before secondary loss ($T$) |
| Plants | Production time per product |
| | Maximum capacity for production |
| | Number of weeks' worth of pre-processed material inventory held |
| | Number of weeks' worth of post-processed inventory held |
| Transportation | Time to ship between plants or DCs |
| | Maximum capacity for shipments |

Figure 2 presents the Arena base model that captures the flow presented in Figure 1. The figure is too small to read the labels, but it conveys the size of the model and illustrates the flow. We color-coded the segments to make explanation of the overall structure easier. The colors also correspond to the colors of the locations in Figure 1. We assume some familiarity with Arena on the reader's part in the following explanation of the flow of the simulation.

- The purple block on the top of Figure 2 occurs at time zero. All inputs are read in from the Excel sheet, target inventory levels are calculated appropriately, and the model begins.
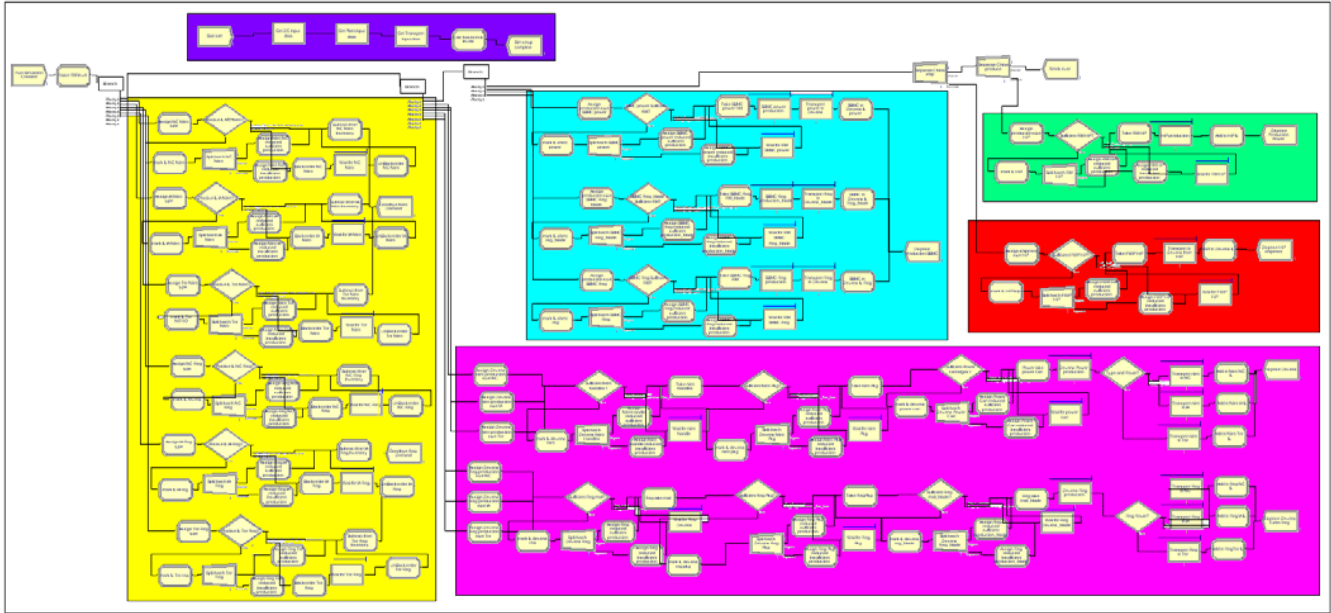
Figure 2: Arena Base Model Flow

- The logic that is uncolored, under the purple block, controls all other simulation functions. The order of logic is very important in this system; for example, production quantities are decided upon before demand is realized in a given week. Thus a single entity is created every week, and it branches and separates along the top of the Figure to create the other entities that process the logic for the simulation.
- The blue, green, and red colored blocks in Figure 2 represent the primary plant production, off-shore production, and off-shore shipment to the U.S., respectively. These occur independently and the quantity produced or shipped is equal to the target inventory level at the packaging plant, minus WIP and on-hand inventory, plus back-orders. The appropriate quantities are produced or shipped only if sufficient input inventory is available.
- The pink block represents the packaging plant; it checks whether it has sufficient inventory from three sources (the two parts of each product, plus packaging materials) before it processes its production batch.
- The yellow block represents the demand for each product at each DC. Back-ordered or lost demand is monitored as a simulation variable and statistics on those variables are generated in the output report.

## 3.3 Model Verification and Validation

We reviewed the structure, basic outputs, and animated flow of the simulation with ABC representatives multiple times. This helps ensure its validity and establish the credibility of the model with the company (Law 2007). We also verified that the outputs of the model made sense numerically by comparing them to their mathematical expectations. For example, the following formulation gives the expected inventory level at a DC:

$$\text{Expected Inventory Level} = \text{Target Inventory Level} - \text{Average Order Quantity} - \text{WIP} + \text{Back-orders}. \quad (1)$$

This formulation is straightforward, but the actual random values for inventory levels, order quantities, WIP, and Back-orders are calculated in several separate locations throughout the model, so verifying their accuracy was important. We evaluated this for multiple components at multiple locations and found that with some stochastic error, the relationship is valid. For example, for the high-volume product at the highest-volume DC, the percent error comparing the output of (1) to the average inventory level output from the simulation was 0.09%.

Additionally, since the amount of inventory on hand will have a direct impact on customer service in the case of a disruption, we validated those numbers as best as possible. We wanted to ensure that the actual inventory levels observed in the simulation matched the values observed in reality, so we requested data on historic inventory levels at the DCs. While

we had been told the target inventory levels should be 4 weeks' worth of demand, when we compared the historic values for inventory to the mean weekly demand, we found that the actual number of weeks on hand varied from 2.5 to 17.5 weeks' worth. Nonetheless, we used these actual weeks' worth of inventory as inputs in Excel for target inventory levels and then compared the on-hand inventory in the simulation to the actual historic values. We did this for all DCs and for the packaging plant, and Table 2 compares the actual versus observed results. Overall, the average of the absolute value of the percent error is 0.31%, which we felt was sufficiently accurate.

Table 2: Inventory Level Evaluation

| *Product* | *Location* | *Avg. % error* | *Avg. abs. value % error* | *Min. % error* | *Max. % error* |
|---|---|---|---|---|---|
| High-Volume | Pkg Plant | 0.13 % | 0.29 % | -0.17 % | 0.42 % |
| | DC 1 | 0.15 % | 0.28 % | -0.21 % | 0.60 % |
| | DC 2 | -0.09 % | 0.44 % | -1.48 % | 0.81 % |
| | DC 3 | 0.03 % | 0.53 % | -0.75 % | 1.62 % |
| Low-Volume | Pkg Plant | -0.01 % | 0.05 % | -0.06 % | 0.04 % |
| | DC 1 | -0.19 % | 0.30 % | -0.48 % | 0.45 % |
| | DC 2 | -0.13 % | 0.27 % | -0.51 % | 0.30 % |
| | DC 3 | -0.27 % | 0.32 % | -0.79 % | 0.19 % |
| All | All | -0.05 % | 0.31 % | -1.48 % | 1.62 % |

Having discussed these values as a team and walked through the simulation flow with ABC representatives, the base model was considered valid and ready for the inclusion of disruptions.

## 4 RISK PROFILE MODEL

At the same time as the base model was being constructed, we started collectin disruptions-related data and developing a Monte Carlo simulation to create a risk profile for every node or arc in the supply chain. For a single node or arc of the supply chain, the risk profile aggregates all the individual risks that that location may experience into one frequency and one duration distribution. These parameters are then used as an input to the Arena risk model. We assume disruptions occur independently at each arc or node (which is supported by the historic disruption information that we collected from ABC), but that the performance of any downstream nodes may be impacted by upstream disruptions. In this way, we capture the interaction of multiple locations in the supply chain, whereas Steckel (2008) focus on only a single node at a time. Yet we avoid end-to-end flows in the whole supply chain which assume that a disruption at a single element will mean the entire path cannot serve customers, as was the case in Deleris and Erhun (2005).

### 4.1 Data Collection

We developed a list of the relevant risks individually for each node and arc of the supply chain based on the records we collected of historic disruptions observed by the company as well as potential future risk sources. These included natural disasters (such as hurricanes or tornados), external disruptions (such as supplier bankruptcy), and internal disruptions (such as information system failure or fire). Some data on the occurrence of risks was collected externally; for example, frequency of natural disasters could be estimated using publically-available data on historic occurrences. Other inputs were estimated by ABC personnel, such as the likelihood of supplier bankruptcy or duration of downtimes caused by a given disruption. (Ideally, this information should come from insurance data, but it was not available for ABC.) Collection of the internally-estimated data was performed by an ABC team member who is responsible for business continuity planning. He spoke to local experts at each location, talking through each hypothetical risk on the list to get their estimate for frequency and duration. Frequency was reported in terms of expected number of years between the particular risk event, and duration was reported using a triangular distribution due to its intuitive nature (the experts estimated the minimum, most likely, and maximum time they could imagine a disruption lasting at their facility). This data was then used as an input to the Monte Carlo simulation.

Since the input data for the Monte Carlo simulation is mostly estimates provided by local experts, we subjected it to additional scrutiny. We validated the estimates by comparing similar locations in the network; for example, if one DC expert estimated that flooding would have a certain duration impact, we verified that this was not drastically different from estimates by experts at other DCs. Additionally, when input parameters are not known with certainty, Law (2007)
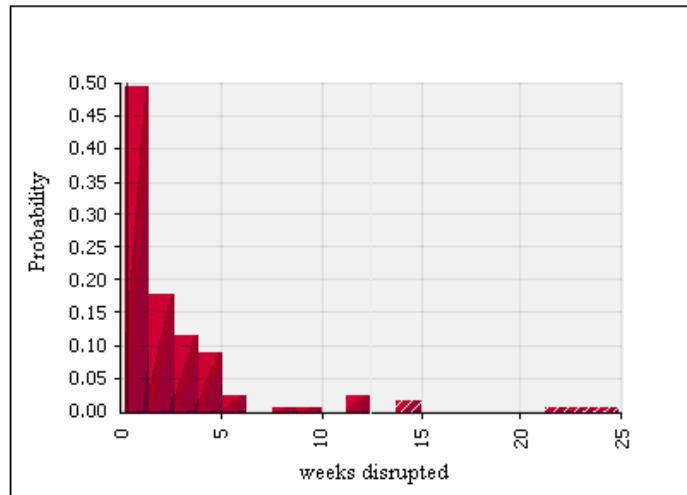
Figure 3: Example Disruption Duration Histogram

recommends conducting sensitivity analysis around the input values. Since the Risk Profile parameters output by the Monte Carlo simulation are inputs for the Arena Risk Simulation Model, we can easily vary them to test the sensitivity of the Risk Model to those parameters.

## 4.2 Monte Carlo Simulation

The Monte Carlo simulation was originally built in Excel using @Risk. Each simulation run generates random numbers that determine whether a particular risk occurs (using an exponential distribution for frequency calculated using the expert's estimate for interarrival time) and how long it lasts (using the triangular distribution provided). In order to develop a disruption-duration distribution, the model then takes the maximum observed disruption length for each run and tallies that value. We ran 10,000 replications for each location, had @Risk perform a distribution fit on the data (filtering out the cases where no disruption occurred), and output the parameters for the best distribution fit that can be read into Arena.

An example of a duration histogram for one of the facilities in the supply chain is given in Figure 3. A LogNormal distribution was the best fit for this histogram, with a shift of 0.14 days, mean of 3.19 days, and standard deviation of 8.64 days. Thus in Arena, the random value for a disruption duration is calculated using the following expression: $duration = 0.14 + LogN(3.19, 8.64)$ days. Most disruption durations were best fit using a lognormal distribution. Distributions for disruption inter-arrivals were fit using data generated in Arena. The results presented in Section 6 are based on this Arena-generated Monte Carlo output.

## 5  RISK SIMULATION MODEL

In this section, we first present a small example of how disruptions can be modeled using Arena. We then discuss the extension of the Base Model to include the risk profiles output from the Monte Carlo simulation and capture mitigation logic in Section 5.2.

## 5.1 Small-Scale Example

In order to illustrate how we accomplished disruption logic using Arena, we discuss a small example with a single product, single customer demand location, and single primary production plant. Additionally, there is a single source of disruptions, and a single back-up production facility. Figure 4 presents the Arena flow for this example.

When a disruption is not present, then a production entity is created every period and begins processing a single batch at Plant A. That batch finishes one period later, is added to inventory, and the production entity is disposed. Similarly, a demand entity is created every period. If there is inventory on hand, it subtracts a batch and is then disposed. If inventory is not available, it waits for material to be produced.
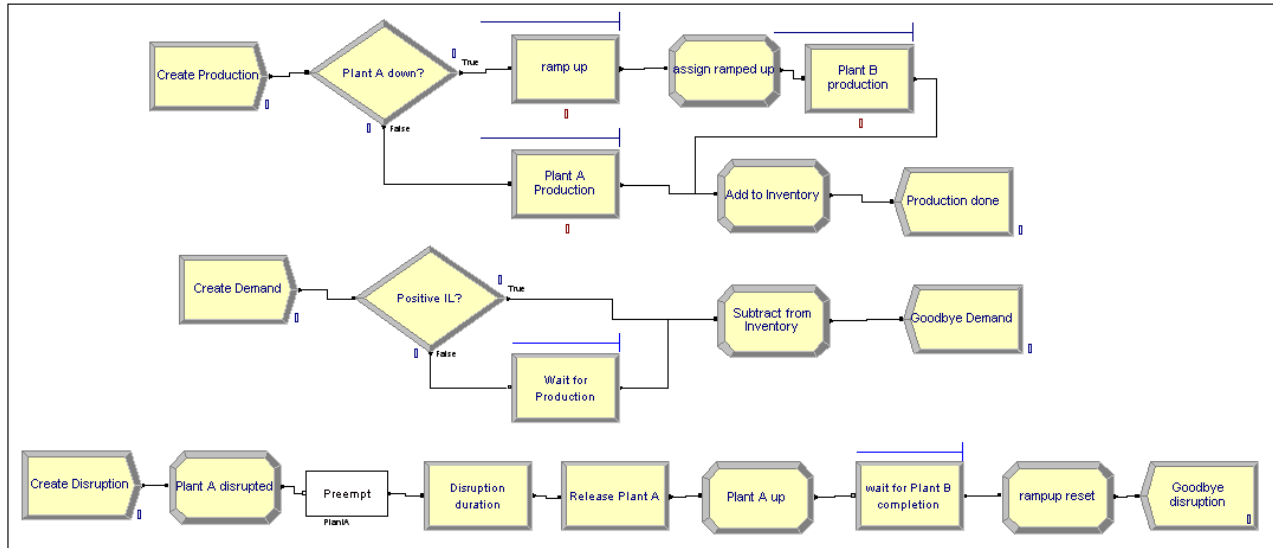
Figure 4: Example of Disruption Mitigation Flow

Disruptions arrive randomly in the simulation using a "Create" block that schedules disruption-entities with an exponential inter-arrival time. When a disruption arrives, it changes a global variable on Plant A's status to indicate that Plant A is disrupted, preempts any batch currently being processed at Plant A, and adds a stochastic delay. In the mean time, any production orders that arrive discover that Plant A is disrupted and are diverted to the "ramp up" delay node on the top of the figure. This ramp up period represents the time for Plant B to be able to produce material in Plant A's place. After the ramp up completes, Plant B continues to produce material for as long as the disruption lasts. When the disruption entity reaches the end of its stochastic delay, it changes the global variable on Plant A's status to indicate Plant A is no longer disrupted, waits for Plant B to finish any back-up processing it is conducting, resets the system so that a ramp-up would again be required for the next disruption, and is disposed.

This is a simple model, but is useful in demonstrating how disruptions work in the full model and the impact of input parameters. For example, if disruptions are short in duration and the starting inventory level is set very high, then the output statistics will show no impact to customer service. However, if inventory levels are low and disruptions are long, significant deterioration of customer service can be observed by high numbers of demand customers waiting in queue for long periods of time. In that sense, the risk profile of a supply network is dynamic in nature, changing every time the inventory levels change in the system.

## 5.2 Arena Logic for the Full Model

For the full risk simulation model, additional inputs need to be read in from Excel. A list is given in Table 3 and explanations are given below. The disruption parameters are the outputs of the Risk Profile Model, as discussed in Section 4. Other inputs describe the logic for ABC's recovery response plans for each node and arc in the supply chain.

The logic for these recovery response plans was estimated by ABC team members. While most personnel had an opinion on how an emergency back-up would be handled, for most locations in the supply chain that policy was not officially documented. We constructed back-up logic in Arena for every location that had at least a hypothetical mitigation plan. For example, if one DC is disrupted, a second DC will begin to serve the first DC's customers. Since the response would probably not be immediate (especially for the locations that have no formal policy documented), we created a parameter for a ramp-up time that is read into Arena, similar to the ramp-up discussed in Section 5.1.

Another input is the capacity of any back-up location, or the time until 100% back-up can be performed. For example, one of the parts manufactured at the primary plant could potentially be supplied by a similar plant in ABC's network in another region of the world. However, since the back-up plant already has a high level of utilization, ABC personnel estimated that only 10% of normal demand volumes could be supplied by this plant after an initial ramp-up of a few weeks. If needed, after several months, the back-up plant's capacity could be expanded so that it could provide 100% of the needed capacity.

Table 3: Risk Model Additional User Inputs

| Category | Input (per location) |
|---|---|
| Disruptions | Frequency |
| | Duration distribution type |
| | Duration distribution parameters |
| Recovery | Ramp-up time for recovery |
| | Capacity available after ramp-up |
| | Time to 100% capacity |
| | Duration threshold to initiate mitigation |

The final back-up input parameter is the time-to-initiation. Not every disruption is long enough to warrant mitigating, so we generated additional logic to see whether a back-up should be initiated or not.

Within Arena, the arrival of disruptions was constructed similarly to the small example discussed in Section 5.1, using the parameters that were obtained from the Monte Carlo simulation. When disruptions arrive, they shut down a production or transportation resource, or stop service from a DC. Several disruption creation blocks were added, with each generating disruptions that have impacts to different locations in the models. Back-up mitigation techniques were also constructed in the same manner as the small-scale example, with additional logic to check whether a disruption warrants mitigating based on the initiation-time inputs.

## 6 RESULTS

Step 5 of the project, Risk Quantification Testing, involved testing both the steady-state performance of the system, as well as specific mitigation strategies. The model proves to be a valuable tool for evaluating mitigation strategies and response time, and quantifying the impact that disruptions have on the system based on its operational state. The primary performance metric that ABC wanted to evaluate is case fill rate, or the percentage of new demand that is satisfied immediately from stock.

### 6.1 Steady-state Analysis

The steady-state analysis of the system is run using very long replications (200 years, which graphical output showed was the time for performance metrics to reach steady levels), without initializing the system between replications. This steady-state evaluation addresses objective 1 introduced in Section 1, assessing the current level of risk in the system. It can also be used to address objective 4, since we will test the sensitivity to input parameters and may be able to identify where unnecessary redundancy exists.

We focused on the impact of the operational state of the system, specifically in terms of target inventory levels, on how the system copes with disruptions. (Recall that ABC generally targets carrying 4 week's worth of inventory at its distribution centers.) We varied this target and compared the fill rates for the base model (with no disruptions) and the risk model (with disruptions turned "on"). Table 4 presents numerical comparisons of the output, providing the average fill rate, halfwidth of a 95% confidence interval on the fill rate, and minimum fill rate observed during the entire run.

Table 4: Model Output: Fill Rates

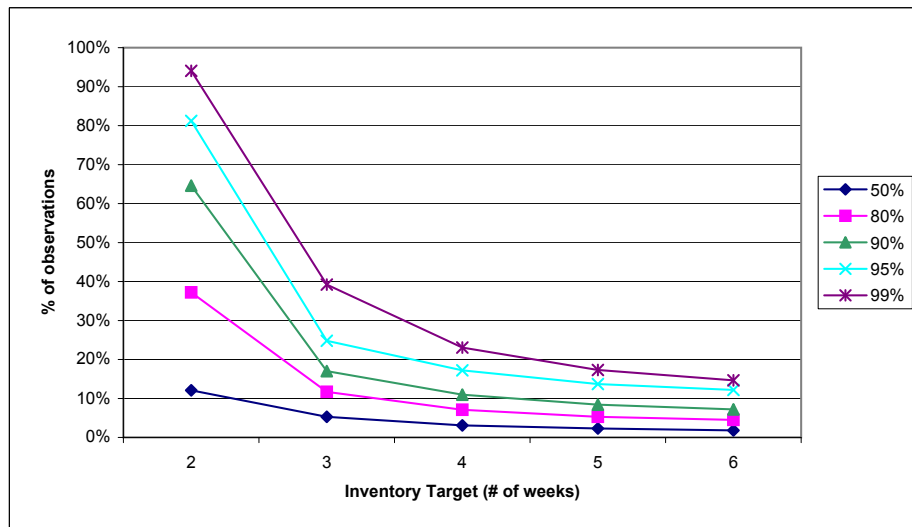| Inventory Target | Model | Average Fill Rate | 95% C.I. halfwidth | Minimum observed |
|---|---|---|---|---|
| 6 weeks | Base | 99.99 % | 0.00 % | 94.16 % |
| | Risk | 96.62 % | 0.24 % | 0.00 % |
| 4 weeks | Base | 99.93 % | 0.00 % | 89.95 % |
| | Risk | 95.16 % | 0.37 % | 0.00 % |
| 2 weeks | Base | 87.68 % | 0.09 % | 32.92 % |
| | Risk | 78.00 % | 0.42 % | 0.00 % |

Figure 5: Percent of Observations Less Than a Given Service Level

The average performance with 4 week's worth of inventory (95% fill rate for the risk model) may not seem that bad, but it represents a significant drop in service for ABC (a drop of 4.77% represents thousands of sales units). Also, ABC's target service level is in the 98% range, so 95% is not satisfactory. The minimum is also clearly significantly different; regardless of the starting inventory level, all runs of the risk model experienced at least one week where 0% of the incoming orders were satisfied.

Interestingly, we observed a slight difference in performance between the two product types modeled in the simulation. While the low-volume product had an inferior service level in the base model (due to its longer, less-responsive supply chain and higher demand coefficient of variance), the high-volume product frequently had an inferior service level in the risk model. This occured because the low-volume product has smaller batch sizes, which means that after a disruption, batches could be completed and shipped more promptly. This implies that recovery from a disruption can be aided by flexible production sizing and scheduling rules.

To further illustrate the at-risk performance of the system, we developed histograms of the risk model's fill rates for various target inventory levels. We used these to calculate what percentage of the observations were below a given service level, and we present those results in Figure 5. For example, the top curve shows the 99% fill rate; when the inventory target is 6 weeks, only 15% of the observations had a fill rate less than 99%. However, when the inventory target is 2 weeks, then 93% of the observations had less than a 99% service level. Figure 5 highlights the dependence on inventory to protect the system from disruption risk, and shows that decisions on target inventory levels should certainly take supply risk into account. While lean initiatives are important, money-saving practices, a company whose inventory levels are too lean can put itself at risk from upstream disruptions in the long run.

## 6.2 Stress-test Analysis

Objective 2, testing different mitigation strategies, was evaluated using what we call stress-testing; using scenario suggestions from ABC, we tested individual disruptions of deterministic length at a given facility to measure their direct impact on the system (with all other disruptions turned off), and graphed the fill rate over time following a disruption. We ran 30 replications for each scenario and tracked the average recovery time, time until the fill rate reaches a steady level, and impact time, number of weeks that the fill rate was below 95% (where 95% was decided upon with ABC team members). Figure 6 presents one example, showing the fill rate over time for the first 5 replications (we limit ourselves to 5 since more than that crowds the figure) where a disruption occurs at the primary plant at time zero in the graph. This graph demonstrates that the recovery can be very different depending on the state of the system (while all replications have the same starting target of 4 weeks, the stochastic demand is different across all replications). While the first replication, in navy blue, climbs steadily from 0% fill rate to 100% between weeks 10 and 18, the second replication, in pink, is much less steady but actually recovers faster, in about 16 weeks.
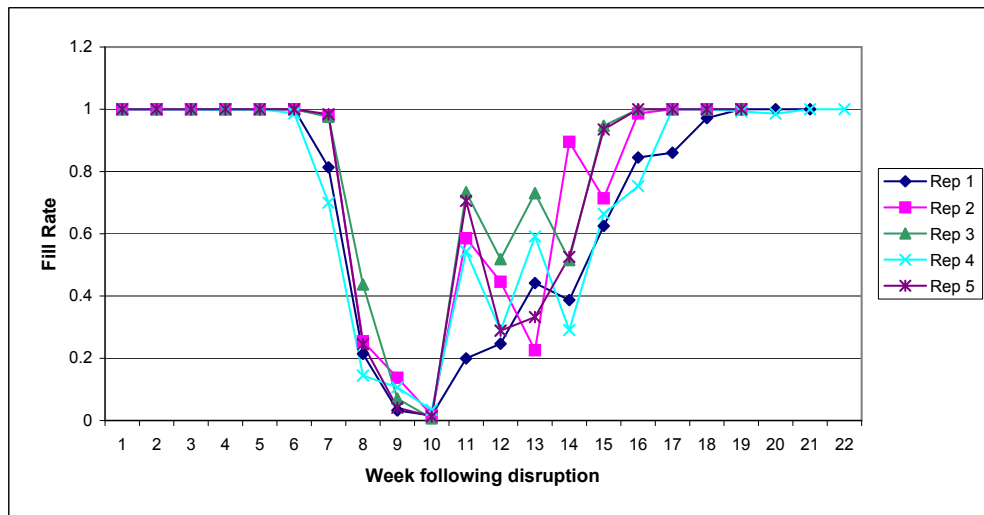
Figure 6: Fill Rate Following a 6-week Disruption at the Primary Plant

Note that the recovery time for Figure 6 typically exceeded the 6-week disruption length. This result was surprising to ABC and was consistently observed for long disruptions upstream in the network. When a long disruption occurs upstream, the system is constrained by capacity at each of its downstream processing locations and takes significant time to fulfill back-orders and replenish the inventory to normal levels. Disruptions in a multi-echelon network clearly have complicated effects, which is why a simulation model is a valuable tool to capture the network interactions.

We ran multiple stress-tests to evaluate disruptions at different locations and with different back-up plans. The output clearly showed that reducing response times or increasing capacity of back-up methods were able to reduce the impact of a disruption and speed recovery. ABC can now use the model to test the efficacy of different back-up options and evaluate whether the investments required are worth the reduced impacts of disruptions that could be achieved. They can also use it in real-time if a disruption occurs to evaluate alternate recovery steps based on the current system status, satisfying objective 3.

## 7 CONCLUSIONS

This project was initiated to evaluate the level of supply chain disruption risk that company ABC currently faces and evaluate mitigation strategies. We constructed a large simulation model in Arena which uses inputs generated by a Monte Carlo simulation to accomplish that task. Through our analysis, we quantified the disruption threat at various locations in the supply chain and developed risk profiles for each. We then developed the means to quantify the impact of disruptions at each of those locations to customer service using our Arena model. We built flexibility into the model through variable inputs that allow analysts or company members to test the impact of different parameters, such as customer behavior, disruption probabilities, back-up response times, etc. Overall, construction of this model using simulation allows for flexible testing of disruption risk in this supply chain, and quantification of the impact of that risk.

The process of the project itself resulted in many useful observations. For example, collecting the data for the Monte Carlo simulation made ABC realize that they lack a database with historic distribution information. Such a database would be useful both for risk quantification and for training purposes.

Another observation was ABC's lack of formal strategic-level mitigation procedures. We can use the model to demonstrate the importance of different response times to urge them to formalize those plans. Additionally, some of the informal back-up plans, which provide a measure of confidence in continuity of the system, were determined to be operationally infeasible. For example, we were initially told that the back-up plant could fully substitute for the primary plant. In reality, the back-up plant lacks some specialized equipment, so it cannot produce all of the components of the product. The model's output shows the importance of documenting and ensuring the feasibility of this strategic-level recovery planning for the network.

Finally, the model output demonstrates the sensitivity of the system performance to the system state. The level of customer service is entirely dependent on the inventory levels at the start of a disruption, and inventory levels change

continually due to stochastic demand or variations in operations planning. Therefore a single label cannot accurately capture the level of risk in the system; risk is dynamic in nature and must be continually evaluated based on operational changes in the network.

## ACKNOWLEDGMENTS

## REFERENCES

Deleris, L. A., and F. Erhun. 2005. Risk management in supply networks using monte-carlo simulation. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 1643–1649. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hopp, W. J., and Z. Yin. 2006. Protecting supply chain networks against catastrophic failures. Working paper, Dept. of Industrial Engineering and Management Science, Northwestern University, Evanston, IL.

Jain, S., and S. Leong. 2005. Stress testing a supply chain using simulation. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 1650–1657. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Law, A. M. 2007. *Simulation modeling & analysis*. Fourth ed. Boston, MA: McGraw-Hill. pp. 253–264.

Schmitt, A. J., L. V. Snyder, and Z. J. M. Shen. 2008. Centralization versus decentralization: Risk pooling, risk diversification, and supply uncertainty in a one-warehouse multiple-retailer system. Working paper, P.C. Rossin College of Engineering and Applied Sciences, Lehigh University, Bethlehem, PA.

Snyder, L. V., and Z. J. M. Shen. 2006. Supply and demand uncertainty in multi-echelon supply chains. Working paper, P.C. Rossin College of Engineering and Applied Sciences, Lehigh University, Bethlehem, PA.

Steckel, M. 2008, December. Measurement, management, and communication of network risk. In *Supply Network Risk Management Symposium Proceedings*. Cambridge, MA: MIT Center for Transportation and Logistics.

Tomlin, B. 2006. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science* 52 (5): 639–657.

## AUTHOR BIOGRAPHIES

**AMANDA J. SCHMITT** is a Postdoctoral Associate at the Center for Transportation and Logistics at MIT. Her research focuses on designing and managing supply chains to protect against supply disruption risk. She received her Ph.D. in Industrial Engineering in 2008 from Lehigh University. Prior to pursuing her Ph.D., she worked for Ethicon, a Johnson & Johnson medical device company, as a Design Quality Engineer in R&D. Her email address is <aschmitt@mit.edu>.

**MAHENDER SINGH** is a Research Director for the MIT Supply Chain 2020 Project, a multi-year research effort looking into the future of logistics and supply chain management, at the Center for Transportation and Logistics at MIT. His research focuses on operations and supply chain management, with particular interest in exploring the underlying structure of complex supply chains. He received his Ph.D. from the University of Tennessee, Knoxville, and has over ten years of experience in the field of supply chain management. His email address is <msingh@mit.edu>.