Proceedings of the 2020 Winter Simulation Conference
K.-H. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, eds.

INVENTORY MANAGEMENT WITH DISRUPTION RISK

Canan Gunes Corlu

Metropolitan College
Boston University
Boston, MA 02215 USA

Bahar Biller

Analytics Center of Excellence
SAS Institute, Inc.
Cary, NC 27513 USA

Elliot Wolf

Risk Management
The Chemours Company
1007 Market Street 12097
Wilington, DE 19899 USA

Enver Yücesan

Technology and Operations Management Area
INSEAD
1, Ayer Rajah Avenue
138676 SINGAPORE

ABSTRACT

Initiatives like lean manufacturing, pooling, and postponement have been effective in mitigating the cost-service trade-off by maintaining high levels of service while reducing system inventories. However, such initiatives exacerbate supply chain disruptions during a catastrophic event, thereby creating a new trade-off between robustness during disruptions and efficiency during normal operations. We evaluate stocking decisions in the presence of operational disruptions, which represent different risks from those associated with demand uncertainties as they stop production flow and typically persist longer. Operational disruptions can therefore be much more devastating though their likelihood of occurrence may be low. Using stochastic simulation, we combine the newsvendor and order-up-to models capturing demand uncertainty costs with catastrophe models capturing not only the cost of supply disruption, but also the cost of recovery, to obtain insights for managing inventory under disruption risk.

1 INTRODUCTION

Man-made or naturally occurring, catastrophic events impact not only human lives and assets, but also perturb the functioning of the economic system resulting in indirect losses (Hallegatte 2014). While catastrophe modeling of property damage is well developed, modeling of operational disruption and the resulting business interruption appears to suffer from the crude nature of functional relationships in catastrophe models that translate property damage into operational disruption. Moreover, estimating business losses is more complicated as it depends largely on decisions during recovery that mitigate losses by using remaining resources more efficiently to maintain business continuity and to accelerate recovery. Hashemi et al. (2015) propose a probabilistic framework for modeling business interruption and reputational losses for process facilities. Rose and Huyck (2016) introduce an approach for improving hazard loss estimation from an insurance perspective.

Inventory management models, on the other hand, are well developed for managing the trade-off between customer satisfaction and the cost of service. Guided by these models, innovative operational initiatives such as lean manufacturing, pooling, and postponement have been successfully introduced to mitigate the cost-service trade-off by maintaining high levels of service without significantly increasing system inventories. However, such initiatives have also reduced the buffers that a firm could fall back on in the event of a disruption, exacerbating the costly effects of operational disruptions. However, as indicated by Pournader et al. (2020), there needs to be a clear definition of the term disruption, which is typically understood to be an adverse event that disrupts the flows of goods or services in supply chains. Under this definition, a minor machine breakdown and a pandemic fall into the same category. However, there should be a clear differentiation between different types of disruptions with varied levels of impact/likelihood, as they require different responses, including different kinds of recovery and resilience planning.

The objective of this paper is therefore to explicitly incorporate the low-probability/high-impact operational disruptions in inventory models. To this end, using stochastic simulation, we combine the newsvendor and order-up-to models capturing demand uncertainty costs with catastrophe models capturing not only supply uncertainty costs, but also business disruption and recovery costs to obtain insights for managing inventory under disruption risk.

1.1 Literature Review

Operational disruptions are fundamentally different from the risks arising from demand uncertainties as they completely stop the production flow and typically persist longer (Kleindorfer and Saad 2005); as a result, the impact of supply chain disruptions can be much more devastating while their likelihood of occurrence is very low. Unfortunately, reinsurance companies report that frequency of natural hazards is on the rise, further increasing the cost associated with supply chain disruptions.

Supply chain disruptions have therefore started receiving increasing attention in the literature. In an empirical study, Hendricks and Singhal (2005) document the negative impact that supply chain disruptions inflict on a firm's financial performance and the slow recovery from such shocks. A comprehensive review of models is provided by Snyder et al. (2014). An equally comprehensive list of approaches for mitigating supply chain disruptions is discussed by Tang (2006). These approaches could be anchored on classical inventory models as in Parlar (1997) and Arreola-Risa and DeCroix (1998), on network design principles as in Aydin et al. (2011), Lim et al. (2013) and Gao et al. (2019), and on big data analytics as in Araz et al. (2020). Most of these studies, however, adopt a high-level perspective in modeling operational disruptions whereby a facility (typically depicted as a node in a network) becomes unavailable following an underlying stochastic process. For example, Snyder and Shen (2006) model disruptions through a two-state Markov chain in which the UP and DOWN states represent the non-disrupted and disrupted states, respectively. To capture the frequency of various categories of disruptions, Schmitt and Singh (2012) use a triangular distribution whose parameters are estimated through discussions with business continuity personnel. Our work differs from the above literature in that we explicitly model the impact of catastrophic events that result not only in down time, but also in loss of inventory, and incur significant clean-up costs.

Catastrophic events typically result in significant harm to humans and substantial damage to the environment along with possible loss of the firm's ability to operate (Lohmann and Yue 2011). Manufacturing firms are uniquely susceptible to these categories of loss due to the diverse range of interconnected supply chains. For example, as reflected in *100 Largest Losses 1978-2017* compiled by Marsh, the accumulated value of the hundred largest losses in the energy sector is more than \$34 billion and the single largest loss is approximately \$1.8 billion due to an explosion at Piper Alpha in the North Sea. Hallegatte (2014) estimates the economic impact of Hurricane Katrina in Louisiana. MacKenzie et al. (2012) show the impact of the Fukushima disaster on production and logistics. There exists a sizeable literature in catastrophe modeling of both natural and man-made disasters. As a detailed review is beyond the scope of the current paper, we mention Embrechts et al. (2012) and Chavez et al. (2016) as basic references for natural extreme events. Lohmann and Yue (2011) offer a simulation-based approach to catastrophe modeling. Watson and Johnson (2004) describe the components of hurricane catastrophe modeling and approximation techniques to estimate loss severity. Ehlen et al. (2012) evaluate large-scale chemical plant and supply disruptions caused by a hurricane with homeland security implications using agent-based simulation methods. Kappes et al. (2012) outline concepts for extending single-hazards risk analysis to multi-hazard modeling and discuss the challenges of these approaches.

As for man-made disruptions, Deleris et al. (2004) offer a conservative probabilistic estimate for the supply chain risk. Industry-specific models can be found, for example, in Rossetti and Bright (2019) for oil and gas, Elkins et al. (2007) for automotive, and in Hashemi et al. (2015) for process industries. Finally, Barker and Haimes (2009) discuss the impact of extreme events on interdependent infrastructure sectors. Similarly, Barker and Santos (2010) extend the dynamic inoperability input-output model to assess productivity degradations due to disasters; in particular, they evaluate the impact of inventories on the resilience of disrupted interdependent systems.

1.2 Our Contributions

Research streams on inventory management and catastrophe modeling have evolved largely independently from each other. This paper is an attempt to intertwine them. More specifically, this paper studies inventory policies in the presence of operational disruptions caused by natural disasters (e.g., floods or earthquakes) or by man-made disasters (e.g., fire or site contamination). We focus on two settings: a single-site, single-period setting with heavy-tailed customer demand and a single-site multi-period model with light-tailed customer demand. Using stochastic simulation, we combine the newsvendor and the order-up-to models capturing the cost of demand uncertainty with catastrophe models capturing the cost of disruption/recovery to obtain insights for managing inventory under disruption risk. Motivated by the growing evidence in the literature that heavy tailed demand can be observed in practice, our single-period model takes into account the heavy-tail behavior of the demand distribution and derives optimal inventory levels capturing heavy-tail demand uncertainty costs and disruption/recovery costs. It is important to note that the inclusion of the recovery costs in the newsvendor model and the order-up-to model makes it impossible to compute the optimal inventory levels and the associated expected total costs. This is where the use of Monte Carlo simulation becomes critical. Similarly, the consideration of multiple recovery days from the disruption calls for discrete-event simulation.

We find that, while for light-tailed demand operational disruptions drive the optimal inventory levels lower with a corresponding degradation in customer service levels, optimal inventory levels increase in the presence of heavy-tailed demand. In all cases, the functional form of the cost of recovery, which captures the cost of unmet demand due to lost inventory as well as the costs associated with cleaning up the facility and resuming production, is the key driver of the expected cost.

The remainder of this paper is organized as follows. Section 2 introduces the notation used throughout the paper. Section 3 presents a mathematical analysis for a newsvendor model with heavy-tailed demand under multiple sources of uncertainty. Section 4 extends the newsvendor model to multiple periods. Section 5 presents a comprehensive simulation study that relaxes some of the key assumptions associated with the analytical models. Section 6 concludes with future research directions.

2 NOTATION

Our focus is on inventory stocking problems in the presence of disruptions caused by random events such as natural disasters or man-made disturbances. Our models combine the well-known newsvendor model for the single-period setting and the base stock policies for the multi-period setting capturing the cost of demand and supply mismatch with catastrophe models capturing the cost of business disruption and recovery. To this end, we will use the following notation in the remainder of the paper:

Random variables

X, demand at a location with cdf F
p, probability of a disruptive event at a location
γ, fraction of inventory lost due to the disruptive event at a location

Model parameters

c, unit procurement cost of inventory
h, unit inventory holding cost at a location
s, unit shortage cost at a location
$R(K)$, cost of recovering K units of inventory following the disruptive event at a location
α, the discount rate

Decision variables

K, quantity of on-site inventory at a location
y, the order-up-to (base stock) level

3 SINGLE-SITE, SINGLE-PERIOD MODEL WITH HEAVY-TAILED DEMAND

Among several distributions that can represent heavy-tailed demand, power-law distribution is shown to fit well to empirical data. For example, among others, Chevalier and Goolsbee (2003) show that the empirical distribution of book demand at Amazon follows a power-law distribution. Similarly, Gaffeo et al. (2008) investigate book demand in Italy and conclude that power-law is a good model to capture the right-tail behavior of book demand. Bimpikis and Markakis (2016) analyze the demand of a large number of movies at Netflix and find that the number of ratings per movie follows a power-law distribution. Natarajan et al. (2018) fit the demand data for 36 spare parts SKUs of a European automobile manufacturer to 17 demand distributions and find that the best fit is obtained by power-law distribution, extreme value or t-distribution. Motivated by these findings in the literature, we also use a power-law distribution in this paper to capture the heavy-tailed demand behavior.

Our goal in this section is thus to characterize the optimal on-site inventory minimizing the mean of the total cost under the demand that follows a power-law distribution. Total cost function is composed of inventory holding cost, demand shortage cost, and the cost of recovering the inventory due to the disruptive event. Similar to Lodree and Taşkin (2008) and Biller et al. (2019), the expected total cost I as a function of the on-site inventory K is given by

$$\begin{aligned} I(K) &= (1-p)\left\{\int_0^K h(K-x)f_X(x;\Psi)dx + \int_K^\infty s(x-K)f_X(x;\Psi)dx\right\} \\ &+ p\left\{\int_0^{(1-\gamma)K} h((1-\gamma)K-x)f_X(x;\Psi)dx + \int_{(1-\gamma)K}^\infty s(x-(1-\gamma)K)f_X(x;\Psi)dx + R(K)\right\} \end{aligned}$$

The first-order optimality condition satisfies

$$\frac{\partial I(K)}{\partial K} = (1-p)(h+s)F_X(K;\Psi) + p(h+s)(1-\gamma)F_X((1-\gamma)K;\Psi) + p\frac{\partial R(K)}{\partial K} - s(1-p\gamma). \tag{1}$$

Plugging in power-law CDF

$$F_X(K;\beta) = 1 - \left(\frac{K}{x_{min}}\right)^{1-\beta}$$

into Equation (1) yields

$$(1-p)(h+s)\left[1-\left(\frac{K}{x_{min}}\right)^{1-\beta}\right] + p(h+s)(1-\gamma)\left[1-\left(\frac{(1-\gamma)K}{x_{min}}\right)^{1-\beta}\right] + p\frac{\partial R(K)}{\partial K} - s(1-p\gamma).$$

Assuming a linear recovery cost $R(K) = \gamma K$ and taking $x_{min} = 1$ w.l.o.g., we obtain the following optimal inventory level:

$$K^* = \left(\frac{h-\gamma ph+p\gamma}{(h+s)(1-p+p(1-\gamma)^{2-\beta})}\right)^{\frac{1}{1-\beta}}, \;\; \beta > 1. \tag{2}$$

If $p = 0$ (i.e., there is no disruptive event), then (2) simplifies to

$$K^* = \left(\frac{h+s}{h}\right)^{\frac{1}{\beta-1}}, \ \beta > 1.$$

This formula is consistent with the optimal newsvendor quantity derived in Bimpikis and Markakis (2016). We observe that the optimal inventory level K^* is decreasing in the parameter β.

If we assume that $\gamma = 1$ (i.e., in the case of a disruptive event all inventory is lost), then the optimal inventory level derived in (2) simplifies to:

$$K^* = \left(\frac{(h+s)(1-p)}{h(1-p)+p}\right)^{\frac{1}{\beta-1}}, \ \beta > 1.$$

Consider the following setting where demand is distributed according to a power-law with parameters $\beta \in \{4, 6, 10\}$. Notice that as β increases, the tails become less heavy. In all experiments, we set $x_{min} = 1$ and we normalize the mean to 20. The likelihood of a disruptive event, $p \in \{0.01, 0.05\}$ while the fraction of inventory lost, $\gamma \in \{0, 0.20, 0.50, 0.80\}$. The cost of recovery takes a linear $(R(K) = \gamma K)$, a cubic $(R(K) = \gamma K^3)$, and an exponential $(R(K) = \gamma e^K)$ form. Table 1 presents the optimal inventory levels K^*, and the expected costs E[Cost], for varying levels of γ and different recovery cost functions for $p = 0.01$ while Table 2 does the same for $p = 0.05$. In both tables, expected costs have been generated through a Monte Carlo simulation under the optimal order-up-to levels K^*. One million independent replications were run in Matlab for each setting to estimate the expected total cost with common random numbers across the different settings.

Table 1: Optimal inventory levels with $p = 0.01$.

			$R(K) = \gamma K$		$R(K) = \gamma K^3$		$R(K) = \gamma e^K$	
β	p	γ	K*	E[Cost]	K*	E[Cost]	K*	E[Cost]
β = 4	0.01	0.0	18.0961	71.3454	18.0961	71.3454	18.0961	71.3454
	0.01	0.2	18.1409	71.5238	18.1354	83.2717	18.1374	1.48E+05
	0.01	0.5	18.3028	72.3681	18.2886	102.4865	18.2938	4.35E+05
	0.01	0.8	19.4882	74.5986	19.4601	132.6269	19.4709	2.25E+06
β = 6	0.01	0.0	19.2180	40.2125	19.218	40.2125	19.2180	40.2125
	0.01	0.2	19.2800	40.5353	19.2774	54.6501	19.2782	4.65E+05
	0.01	0.5	19.7806	41.7861	19.7735	79.8642	19.7757	1.91E+06
	0.01	0.8	28.5946	97.0597	28.5553	280.5000	28.5717	2.02E+10
β = 10	0.01	0.0	19.6831	21.4363	19.6831	21.4363	19.6831	21.4363
	0.01	0.2	19.7932	21.9057	19.7920	37.1830	19.7923	7.78E+05
	0.01	0.5	22.6698	35.6644	22.6649	93.0221	22.6665	3.44E+07
	0.01	0.8	49.3786	292.778	49.2820	1.24E+03	49.3126	2.06E+19

We note the following behavior:

- For fixed β and p, the optimal inventory level K^* and the expected cost both increase with the proportion of inventory lost (γ). The increase is more dramatic as we move from $\gamma = 0.5$ to $\gamma = 0.8$. This observation holds for all types of recovery functions.
- As we move from a linear recovery function $R(K) = \gamma K$ to an exponential recovery function $R(K) = \gamma e^K$, we observe that there is no significant change in the optimal inventory levels, K^*.

Table 2: Optimal inventory levels with $p = 0.05$.

β	p	γ	$R(K) = \gamma K$ K*	E[Cost]	$R(K) = \gamma K^3$ K*	E[Cost]	$R(K) = \gamma e^K$ K*	E[Cost]
β = 4	0.05	0.0	18.0961	71.3454	18.0961	71.3454	18.0961	71.3454
	0.05	0.2	18.3193	72.2012	18.2908	133.0473	18.3012	8.85E+05
	0.05	0.5	19.1035	76.7853	19.0210	247.8090	19.0521	4.68E+06
	0.05	0.8	23.8251	98.1055	23.5557	617.1808	23.6657	7.56E+08
β = 6	0.05	0.0	19.2180	40.2125	19.2180	40.2125	19.2180	4.02E+01
	0.05	0.2	19.5226	41.7749	19.5090	115.6301	19.5132	2.97E+06
	0.05	0.5	21.5920	51.1445	21.5433	299.7098	21.5607	5.76E+07
	0.05	0.8	38.7671	191.6859	38.2672	2.42E+03	38.4514	1.99E+15
β = 10	0.05	0.0	19.6831	21.4363	19.6831	21.4363	19.6831	21.4363
	0.05	0.2	20.1939	23.9513	20.1874	105.7895	20.1891	5.84E+06
	0.05	0.5	26.4040	68.7767	26.3623	524.5559	26.3785	7.12E+09
	0.05	0.8	59.2396	381.274	58.4386	8.33E+03	58.5699	1.09E+24

However, expected costs increase dramatically. This is expected because the value of the recovery cost function that is added to the expected cost keeps increasing when we move from $R(K) = \gamma K$ to $R(K) = \gamma K^3$ and to $R(K) = \gamma e^K$ as there is no significant change in K^* values.

- For fixed p and γ, we observe that expected costs generally decrease with increasing β. Two cases where this observation does not hold are when $\gamma = 0.8$ and $R(K) = \gamma e^K$.
- Optimal inventory levels and expected costs increase with increasing p.

4 SINGLE-SITE MULTI-PERIOD MODEL

The natural extension of the newsvendor model in a single-site, single-period setting onto a single-site, multi- period setting is the order-up-to (base stock) policy. To this end, let x_t represent the inventory level at the beginning of period t. A positive value for x_t indicates that x_t units of inventory were carried from the previous period while a negative value for x_t indicates that a backlog of $-x_t$ units is carried from the previous period. Let $y_t - x_t \geq 0$ denote the size of the replenishment order in period t, resulting in a procurement cost of $c_t(y_t - x_t)$ and an increase in inventory level to y_t. To keep the exposition simple, let us assume that the replenishment is instantaneous. As a result, note that $x_{t+1} = y_t - D_t$. The optimization problem over a horizon of T periods can then be written as:

$$C_1(x_1) = min_{y_t \geq x_t} E\big[\sum_{t=1}^{T} \alpha^{t-1} c_t(y_t - x_t) + G_t(y_t)\big] + \alpha^T c_{T+1}, \tag{3}$$

where $G_t(y_t)$ is the newsvendor loss function. Under the assumption of stationarity (i.e., $c_t \equiv c, h_t \equiv h, s_t \equiv s$) with independent and identically distributed (IID) demands across all periods, the finite-horizon discounted costs converge when the discount rate $\alpha < 1$. The optimization problem in (3) can then be written as

$$C(x) = min_{y \geq x}\{c(y-x) + G(y) + \alpha E[C(y-D)]\},$$

where the newsvendor loss function is given by

$$G(y) = (1-p)\{hE[(y-D)^+ + sE[(D-y)^+]\} + p\{hE[((1-\gamma)y - D)^+ + sE[(D-(1-\gamma)y)^+] + R(y)\}.$$

Recall that $R(y)$ is the cost of recovering y units of inventory following the disruptive event. In this setting a *myopic* policy would order $(y^m - x)^+$ units, where y^m minimizes the current cost

$$C_m = c(1-\alpha)y + G(y). \tag{4}$$

Surprisingly, under the mild assumption that demand takes only non-negative values, the myopic policy turns out to be optimal. We then find the optimal order-up-to (base stock) level, y^*, by differentiating (4):

$$\frac{\partial C_m}{\partial y} = c(1-\alpha) + (1-p)(h+s)F(y) + p(1-\gamma)(h+s)F((1-\gamma)y) - s(1-p\gamma) + p\frac{\partial R(y)}{\partial y} \equiv 0.$$

We follow the experimental setting in Section 3 except that we consider the case where demand is IID across periods under a normal distribution with a coefficient of variation of 0.05 and 0.5. Table 3 shows the optimal order-up-to levels for various settings. The last two columns in the table, which show the expected cost and the expected fill rate, respectively, have been generated through a Monte Carlo simulation under the optimal order-up-to levels y^*. For each setting, one million independent replications were run in Matlab to estimate the expected total cost with common random numbers across the different settings.

Table 3: Optimal order-up-to levels with demand CV of 0.05.

p	γ	$R(y)$	y^*	E[Cost]	E[fill rate]
0.01	0.00	y	19.7981	99.77	0.9746
	0.20	y	19.8056	100.32	0.9730
	0.50	y	19.8007	101.25	0.9700
	0.80	y	19.7957	102.22	0.9668
0.01	0.00	y^3	19.7981	99.77	0.9746
	0.20	y^3	19.7977	115.88	0.9728
	0.50	y^3	19.7807	139.66	0.9693
	0.80	y^3	19.7637	164.57	0.9658
0.01	0.00	e^y	19.7981	99.78	0.9746
	0.20	e^y	8.5631	220.91	0.4272
	0.50	e^y	7.6425	230.82	0.3802
	0.80	e^y	7.1682	235.95	0.3556
0.05	0.00	y	19.7981	99.78	0.9745
	0.20	y	19.8373	102.56	0.9666
	0.50	y	19.8116	107.30	0.9510
	0.80	y	19.7857	112.09	0.9353
0.05	0.00	y^3	19.7981	99.78	0.9746
	0.20	y^3	19.7959	179.83	0.9654
	0.50	y^3	19.7072	298.41	0.9480
	0.80	y^3	19.6167	414.30	0.9302
0.05	0.00	e^y	19.7981	99.77	0.9746
	0.20	e^y	6.9422	238.61	0.3436
	0.50	e^y	6.0039	249.42	0.2926
	0.80	e^y	5.5114	255.20	0.2647

We note that while both the likelihood of a disruptive event, p, and the fraction of inventory lost, γ, reduce the optimal order-up-to levels, y^*, it is the functional form of the recovery cost that has the most significant impact on the optimal inventory investment. We also note that lower order-up-to levels also result in drastically lower fill rates. This, in turn, implies that, to maintain the same service level, the firm would need to hold additional inventory at an independent secondary site.

In the next section, we relax some of the fundamental assumptions underlying the base-stock models through more detailed discrete event simulations where we calibrate our models with industry data.

5 SIMULATION STUDY

In this section, we develop a discrete-event simulation, which additionally accounts for disruptions that last for multiple periods. A high-level description of this model is presented in Section 5.1 and preliminary results are provided in Section 5.2.

5.1 Description

The simulation mimics the flow of inventory through a distribution center where customer orders are received on a daily basis. The upstream warehouse from which the distribution center replenishes its inventory is considered not to be subject to any capacity constraints. We further assume a negligible replenishment lead time. We simulate this inventory flow, illustrated in Figure 1, using a discrete-event simulation model developed in SAS Simulation Studio (Hughes et al. 2018).

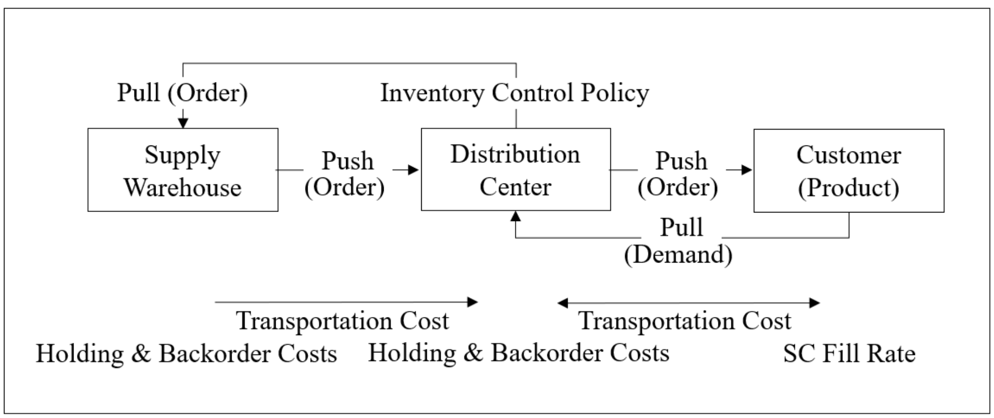

Figure 1: Inventory flow simulated by SAS Simulation Studio.

When the length of the simulation is restricted to one day, the discrete-event simulation setting of this section reduces to the single-site, single-period model of Section 3 (Table 2), where p = 0.05, γ = 0.8, $R(K) = \gamma K$, and the daily average customer order size is 20 units with β = 10. However, the discrete-event simulation extends the length of the simulation to 30 days with the possibility of a disruptive event on each day of the 30-day horizon. It is for this reason that we set the initial inventory level at simulation time zero to 60 units given that K^* is presented as 59.2396 in Table 2. At the beginning of each day, the inventory position (i.e., units of inventory on hand – daily average customer order size – units of backlogged customer orders) is checked to determine whether to place a replenishment order to bring the inventory position to the order-up-to level, assumed to be 60 units. Because the replenishment lead time is assumed to be zero, the replenishment decision is driven by the records of on-hand inventory, the daily average customer order size, and the number of units backlogged. The simulation is designed to observe the size of the customer order two hours into the day; however, it is possible to experience disruption before the realization of the demand on any day.

Consistent with the modeling assumptions made in Sections 3 and 4, we set the daily cost of holding unit inventory, h to 10 dollars and the unit back-ordering cost, s and the unit procurement cost, c to 15 dollars and 5 dollars, respectively. However, the key distinguishing feature of the simulation model is the incorporation of the number of days it would take for the facility to recover from the disruptive event.

While we have the flexibility to represent stochastic recovery durations, we design the experiments of the following section by assuming a deterministic recovery duration of one day, two days, and five days for each of the disruptive events that could realize. More specifically, we load the simulation with outage entities, each of which is defined to have the following attributes: time of arrival, identifiers of product and distribution center, likelihood (i.e., p), percentage inventory loss (i.e., γ), type of the recovery cost function (which is linear, cubic or exponential), and the number of days to recovery. Thus, our model possesses the flexibility of representing multiple distribution centers each of which may be exposed to different types of outages at different times with their own specific attributes. It is also possible that the system suffers from another disruption before fully recovering from the previous disruption.

If it takes only one day to recover from the disruption, then the only effect of this event would be the loss of a γ fraction of inventory on hand, which could result in partial fulfillment of the incoming customer orders, incurring back-ordering cost on the day of the disruption only. However, if it takes more than one day to recover from the disruption, then the distribution center would still place a replenishment order, but be unable to receive it immediately during the duration of recovery, in addition to losing a fraction of its on-hand inventory. However, a replenishment order is never lost and any customer order received at the distribution center is eventually fulfilled. Those replenishment orders that cannot be fulfilled instantaneously during the recovery period, despite the zero lead-time assumption, are tracked and utilized in the calculation of the inventory position (i.e., units of inventory on hand + units of outstanding replenishment orders – daily average customer order size – units of backlogged customer orders) at the beginning of each day of recovery.

We conduct 1,000 independent replications and record the output data to enable the calculation of unit fill rate and total cost as the key performance indicators. Specifically, the unit fill rate is computed for each day of the simulation horizon as the ratio of the number of customer-order units fulfilled to the size of the customer order received. The mean unit fill rate presented in the following section is the unit fill rate averaged across the 30 consecutive days simulated. The total cost is, on the other hand, the sum of time-averaged inventory holding cost, back-ordering cost, cost of procuring the replenishment orders and cost of recovering from all the disruptions that the system has been exposed to.

5.2 Preliminary Results

We run the discrete-event simulation for five different scenarios (Table 4). The scenarios differ from each other by the form of the cost function and the duration of recovery. Table 4 presents the fill-rate predictions for the first three scenarios where the form of the cost function is linear but the duration of recovery increases from 1 day to 5 days. We observe that increasing the recovery duration from 1 day to 5 days reduces the unit fill rate from 98% to 75%. In addition to this significant impact on the average fill rate, we also observe a considerable increase in the standard deviation of the fill rate.

Table 4: Fill-rate prediction: mean and standard deviation.

Experiment	Cost Function	Duration	Mean Fill Rate	Standard Deviation
Scenario 1	linear	1 day	98.42%	0.92%
Scenario 2	linear	2 days	90.88%	15.74%
Scenario 3	linear	5 days	74.95%	25.54%

Table 5 shows the impact of the form of the recovery cost function on the mean and standard deviation of the total cost. We observe, from Table 5, that the significant effect of the recovery function's form, observed in Section 3 and Section 4, continues to hold while being even more strongly felt in the case of simulating the inventory flow across multiple periods. Table 6 further decomposes the average total cost into its four components associated with holding inventory, back-ordering, procurement, and recovery. As we switch from Scenario 1 to Scenario 3, we observe the total recovery cost to decrease even though the duration of recovery increases. However, it is reasonable to expect the total recovery cost to be higher under

Scenario 3 than under Scenario 1. The explanation of this observation is that as recovery from disruption lasts longer, it takes us longer to receive incoming inventory from the warehouse. Therefore, the amount of on-hand inventory decreases. When a second disruption is realized, we lose a fraction of the already reduced amount of inventory. Because the recovery cost function is linear under Scenario 3, the impact on total recovery cost appears to decrease, masking the seriousness of how severely the inventory flow is affected by the disruptive events. In this case, a better indicator of the situation is the sharp increase in the back-order cost incurred as we fail to meet an increasing number of customer orders on time.

Table 5: Total cost prediction: mean and standard deviation.

Experiment	Cost Function	Duration	Mean Total Cost	Standard Deviation
Scenario 1	linear	1 day	$ 20,884	$ 186
Scenario 2	linear	2 days	$ 23,747	$ 15,493
Scenario 3	linear	5 days	$ 32,512	$ 26,363
Scenario 4	cubic	5 days	$ 412,694	$ 593,051
Scenario 5	exponential	5 days	$ 8.964E+77	$ 2.835E+79

Table 6: Mean total cost decomposition.

Experiment	Inventory Cost	Back-order Cost	Procurement Cost	Recovery Cost
Scenario 1	$ 17,027	$ 108	$ 3,653	$ 123
Scenario 2	$ 15,626	$ 4,572	$ 3,486	$ 115
Scenario 3	$ 12,876	$ 16,575	$ 3,018	$ 96
Scenario 4	$ 12,876	$ 16,575	$ 3,018	$ 487,509
Scenario 5	$ 12,876	$ 16,575	$ 3,018	$ 1.149E+78

Significant increase in variability tabulated in this section from Scenario 1 to Scenario 5 arises from the differences of sample paths experiencing small exposure to the disruption risk and those experiencing large exposure to the disruption risk. The simulation model can be further utilized to extract rules about how to mitigate the impact of disruption depending on the type of the sample path realized. Therefore, integration of this simulation with predictive models of disruption as well as with the simulation optimization techniques suitable for sequential decision making would be of great value in the management of the disruption risk. This is the subject of the ongoing work.

6 CONCLUSION

Using stochastic simulation, we combine the newsvendor and the order-up-to models capturing the cost of demand uncertainty with catastrophe models capturing the cost of disruption/recovery to obtain insights for managing inventory under disruption risk. Motivated by the growing evidence in the literature that heavy tailed demand can be observed in practice, our single-period model takes into account the heavy-tail behavior of the demand distribution and derives optimal inventory levels capturing heavy-tail demand uncertainty costs and disruption/recovery costs. We observe that, while for light-tailed demand operational disruptions drive the optimal inventory levels lower with a corresponding degradation in customer service levels, optimal inventory levels increase in the presence of heavy-tailed demand. In all cases, the functional form of the cost of recovery, which captures the cost of unmet demand due to lost inventory as well as the costs associated with cleaning up the facility and resuming production, is the key driver of the expected cost.

An alternative to adjusting the inventory levels in the face of catastrophic risk is to have an emergency plan to evacuate the inventory ahead of the potentially catastrophic event –provided that there is sufficient time to do so. For instance, it is common for chemical manufacturers to have an emergency response plan, which prescribes the relocation of chemicals during a Category 3 or greater storm. Simulation models as presented in this work, which considers the frequency of extreme weather events, equipment reliability, disruptions severity, and resource availability, can inform optimal policy decisions which can be used for strategic scenario planning.

REFERENCES

Araz, O. M., T. M. Choi, D. L. Olson, and F. S. Salman. 2020. "Role of Analytics for Operational Risk Management in the Era of Big Data". *Decision Sciences*. DOI=https://doi.org/10.1111/deci.12451.

Arreola-Risa, A., and G. A. DeCroix. 1998. "Inventory Management Under Random Supply Disruption and Partial Backorders". *Naval Research Logistics* 45(7):687–703.

Aydin, G., V. Babich, D. R. Beil, and Z. Yang. 2011. "Decentralized Supply Risk Management". In *The Handbook of Integrated Risk Management in Global Supply Chains*, edited by P. Kouvelis, L. Dong, O. Boyabatli, and R. Li, 389–424. New York: Wiley.

Barker, K., and Y. Y. Haimes. 2009. "Assessing Uncertainty in Extreme Events: Applications to Risk-based Decision Making in Interdependent Infrastructure Sectors". *Reliability Engineering and System Safety* 94(7):819–829.

Barker, K., and J. R. Santos. 2010. "Measuring the Efficacy of Inventory with a Dynamic Input-Output Model". *International Journal of Production Economics* 126(1):130–143.

Biller, B., E. Wolf, and E. Yücesan. 2019. "Inventory Management Under Disruption Risk". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 1918–1929. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Bimpikis, K., and M. G. Markakis. 2016. "Inventory Pooling Under Heavy-Tailed Demand". *Management Science* 62(6):1800–1813.

Chavez, N., M. Ghil, and J. Umitia-Fucugauchi. 2016. *Extreme Events: Observations, Modeling, and Economics*. Hoboken, New Jersey: Wiley and Sons.

Chevalier, J., and A. Goolsbee. 2003. "Measuring Prices and Price Competition Online: Amazon.com and BarnesandNoble.com". *Quantitative Marketing and Economics* 1(2):203–222.

Deleris, L. A., D. Elkins, and M. E. Pate-Cornell. 2004. "Analyzing Losses from Hazard Exposure: A Conservative Probabilistic Estimate Using Supply Chain Risk Simulation". In *Proceedings of the 2004 Winter Simulation Conference*, edited by R. G. Ingalls, M. D. Rosetti, J. S. Smith, and B. A. Peters, 1384–1391. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Ehlen, M. A., A. C. Sun, M. A. Pepple, E. D. Eidson, and B. S. Jones. 2012. "Chemical Supply Chain Modeling for Analysis of Homeland Security Events". Technical Report SAND2012-2956 J, Sandia National Laboratories, Albuquerque, New Mexico.

Elkins, D., A. C. Lafleur, E. Foster, J. Tew, B. Biller, and J. R. Wilson. 2007. "Clinic: Correlated Inputs in an Automotive Paint Shop Fire Risk Simulation". In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. Tew, and R. Barton, 250–259. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Embrechts, P., C. Klüppelberg, and T. Mikosch. 2012. *Modeling Extremal Events for Insurance and Finance*. 9th ed. Berlin, Germany: Springer.

Gaffeo, E., A. E. Scorcu, and L. Vici. 2008. "Demand Distribution Dynamics in Creative Industries: The Market for Books in Italy". *Information Economics and Policy* 20(3):257–268.

Gao, S. Y., D. Simchi-Levi, C.-P. Teo, and Z. Yan. 2019. "Disruption Risk Mitigation in Supply Chains: The Risk Exposure Index Revisited". *Operations Research* 67(3):831–852.

Hallegatte, S. 2014. "Modeling the Role of Inventories and Heterogeneity in the Assessment of the Economic Cost of Natural Disasters". *Risk Analysis* 34(1):152–167.

Hashemi, S. J., S. Ahmed, and F. Khan. 2015. "Probabilistic Modeling of Business Interruption and Reputational Losses for Process Facilities". *Process Safety Progress* 34(4):373–382.

Hendricks, K. B., and V. R. Singhal. 2005. "An Empirical Analysis of Supply Chain Disruptions on Long-run Stock Price Performance and Equity Risk of the Firm". *Production and Operations Management* 14(1):35–52.

Hughes, E., J. Pratt, and B. Biller. 2018. "Solving Business Problems with SAS Analytics and OPTMODEL". In *Technology Workshop Presented at INFORMS Analytics Conference*. April 15$^{\text{th}}$-17$^{\text{th}}$, Baltimore, Maryland.

Kappes, M. S., M. Kailer, K. von Elverfeldt, and T. Glade. 2012. "Challenges of Analyzing Multi-hazard Risk: A Review". *Natural Hazards* 64(2):1925–1958.

Kleindorfer, P., and G. Saad. 2005. "Managing Disruption Risks in Supply Chains". *Production and Operations Management* 14(1):53–68.

Lim, M. K., A. Bassamboo, S. Chopra, and M. S. Daskin. 2013. "Facility Location Decisions with Random Disruptions and Imperfect Estimation". *Manufacturing and Service Operations Management* 15(2):239–249.

Lodree, E. J., and S. Taşkin. 2008. "An Insurance Risk Management Framework for Disaster Relief and Supply Chain Disruption Inventory Planning". *Journal of the Operational Research Society* 59(5):674–684.

Lohmann, D., and F. Yue. 2011. "Correlation, Simulation and Uncertainty in Catastrophe Modeling". In *Proceedings of the 2011 Winter Simulation Conference*, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, 133–145. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

MacKenzie, C. A., J. R. Santos, and K. Barker. 2012. "Measuring Changes in International Production: Case Study of the Japanese Earthquake and Tsunami". *International Journal of Production Economics* 138(2):293–302.

Natarajan, K., M. Sim, and J. Uichanco. 2018. "Asymmetry and Ambiguity in Newsvendor Models". *Management Science* 64(7):3146–3167.

Parlar, M. 1997. "Continuous Review Inventory Problem with Random Supply Interruptions". *European Journal of Operational Research* 99(2):366–385.

Pournader, M., A. Kach, and S. Talluri. 2020. "A Review of the Existing and Emerging Topics in the Supply Chain Risk Management Literature". *Decision Sciences*. DOI=https://doi.org/10.1111/deci.12470.

Rose, A., and C. K. Huyck. 2016. "Improving Catastrophe Modeling for Business Interruption Insurance Needs". *Risk Analysis* 36(10):1896–1915.

Rossetti, M. D., and J. Bright. 2019. "Analyzing Pre-positioning within a Disrupted Bulk Petroleum Supply Chain". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. Bae, S. Lazarova-Molnar, C. Szabo, P. Haas, and Y.-J. Son, 1813–1824. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Schmitt, A. J., and M. Singh. 2012. "A Quantitative Analysis of Disruption Risk in a Multi-echelon Supply Chain". *International Journal of Production Economics* 139(1):22–32.

Snyder, L. V., Z. Atan, P. Peng, Y. Rong, A. J. Schmitt, and B. Sinsoysal. 2014. "OR/MS Models for Supply Chain Disruptions: A Review". Technical report, Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, Pennsylvania.

Snyder, L. V., and Z. J. M. Shen. 2006. "Supply and Demand Uncertainty in Multi-Echelon Supply Chains". Technical report, Lehigh University, Bethlehem, Pennsylvania.

Tang, C. S. 2006. "Robust Strategies for Mitigating Supply Chain Disruptions". *International Journal of Logistics* 9(1):33–45.

Watson, C., and M. Johnson. 2004. "Hurricane Loss Estimation Models, Opportunities for Improving the State of the Art". *Bulletin of the American Meteorological Society* 85(11):1713–1726.

AUTHOR BIOGRAPHIES

CANAN GÜNEŞ ÇORLU is Associate Professor in the Administrative Sciences Department of Metropolitan College at Boston University. She holds a Ph.D. in Operations Management from the Tepper School of Business at Carnegie Mellon University. Her research interest is in the area of the design and analysis of stochastic simulations with applications in operations management. Her email address is canan@bu.edu.

BAHAR BILLER is a member of SAS Analytics Center of Excellence. She develops data-driven and scalable solutions for complex and dynamic business processes. Previously, she was a Senior Scientist at General Electric's Global Research Center and an Associate Professor in the Tepper School of Business at Carnegie Mellon University. She holds a Ph.D. degree in Industrial Engineering and Management Sciences from Northwestern University. Today at SAS, she operationalizes healthcare and supply chain simulations through machine learning and artificial intelligence for supporting decisions in real time. Her e-mail address is bahar.biller@sas.com.

ELLIOT WOLF is a Global Process Safety Leader at the Chemours Company with experience in chemical process safety, supply chain, operations, research and development, and engineering. His interests include application of data analytics and risk analysis methodologies to protect people, the environment, and businesses. He holds an undergraduate degree and Masters degree in Chemical Engineering from the University of Pittsburgh and the University of Delaware, respectively, and an MBA from the Tepper School of Business at Carnegie Mellon University. His email address is elliot.wolf@chemours.com.

ENVER YÜCESAN is the Abu Dhabi Commercial Bank Chair in International Management and a professor in the Technology and Operations Management Area at INSEAD. He earned his Ph.D. in Operations Research from Cornell University. Using simulation, he addresses problems in agriculture, supply chain management, and social media. His email address is enver.yucesan@insead.edu.