

Detecting Criminal Disruption of Supply Chains Study: Phase II: Medical Devices and Critical-Care Supply Chains

CARLOTTA DOMENICONI, ABHISEKH RANA, RAJ PATEL, JAMES JONES, HAMDİ KAVAK,
SEAN LUKE, AND VLADIMIR MENKOV

George Mason University

FRED ROBERTS, ALOK BAVEJA, WEIWEI CHEN, DENNIS EGAN, AMAN GOSWAMI, RONG
LEI, PETER MARSH, BENJAMIN MELAMED, AND VISWANATH NARAYAN

Rutgers University

ANDREW CROOKS

University at Buffalo

June 15, 2023

Executive Summary This report describes the results of Phase II of a project that involves the development of supply chain models, methods by which criminal organizations may disrupt supply chains effectively, indicators and warnings regarding these disruptions, and ultimately approaches to mitigate such disruptions. The focus of this phase is on supply chains for medical devices deployed for healthcare applications. Such supply chains are part of the national critical infrastructure since the products they deliver are vital in maintaining public health and safety in hospitals, clinics, long-term care facilities, and home health-care. Consider what happened during the pandemic when there was a shortage of PPE and ventilators. Non-availability of medical devices poses significant risk to the safety and health of society. Further, these medical devices due to their electronic components that are connected to the internet, are especially vulnerable to attacks from criminals worldwide. Disruptions to medical device supply chains have increased rapidly, with some sophisticated attacks compromising the availability and operation of such critical devices. This report (1) describes a generic supply chain for the medical device industry and the process by which it was collaboratively developed and vetted, (2) presents a wide variety of plausible disruption scenarios based on extensive study of past events and discussions with subject matter experts, and (3) discusses required capabilities of criminal organizations to carry out such disruptions. This report also presents a step-by-step,

detailed implementation of the supply chain model using a hybrid discrete event and agent-based simulation. We illustrate ways to model the application of criminal organization capabilities to disrupt the supply chain, and present stochastic optimization of simulated approaches. The report lays out how such a supply chain was analyzed through an extensive series of simulation experiments, provides a summary of insights and conclusions gained from this analysis, and makes recommendations for further analysis. The approach described here can successfully identify disruptions and suggest promising mitigation strategies, and lays the groundwork for the analysis of other supply chains. The approach can potentially provide a framework for the Department of Homeland Security, other government agencies, and the private sector to build models of supply chains, to pose *what-if* questions regarding optimal ways to disrupt them given specific amounts and types of resources and know-how, to examine how to make supply chains more robust against attacks, to identify and mitigate attacks as they are ongoing or even as they are being planned, and to repair damage after the fact. This phase of the project offers new insights into the significance of coordination between two sub-supply chains (electronic equipment and delivery systems), importance of maintenance/repair of the electronic equipment, critical role of Contract Manufacturing Organizations, and the direct impact of unavailability of devices on patients. Finally, this work highlights the critical importance of building reactive speed capability to respond to emerging threats in medical device supply chains that complement proactive disruption management strategies such as building safety stocks throughout the supply chain.

1 Introduction

Criminal organizations have previously disrupted and manipulated critical supply chains for financial gain and other reasons, and they will continue to do so. Such disruption and manipulation may take different forms including blocking one or more elements of a supply chain to demand ransom, causing damage to the target, creating delay or uncertainty, motivating a redirection to alternative suppliers, injecting adulterated material into the supply chain, or removing genuine materials, to name a few. Since this project began, there has been a dramatic increase in attention paid to supply chains, their impact on the global economy, and their exposure to disruptions. During a time of global crisis, the effects of supply chain disruptions or manipulations are magnified as already fragile systems and populations are under stress.

The Command, Control and Interoperability Center for Advanced Data Analysis (CCICADA) and the center for Criminal Investigations and Network Analysis (CINA), two Department of Homeland Security (DHS) Centers of Excellence, are developing novel methods and tools to model, detect, and mitigate active,

pending, or past criminal manipulation or disruption of a supply chain. Our approach is to model both supply chains and criminal actors, then merge those models to establish potential disruption or manipulation scenarios, associated indicators, and mitigation strategies. The resulting analysis for the specific supply chains aims to identify likely attack points, develop indicators and warnings (I&W) that can serve to alert authorities and supply chain operators about a pending, active, or past attack, and provide recommendations to mitigate the identified vulnerabilities and reduce attack impacts. The resulting methodology can then be applied to other supply chains as needed.

The CCICADA team brings strong experience and expertise in supply chain modeling and analytics, including the identification and mitigation of vulnerabilities in supply chains. The CINA team brings strong experience and expertise in criminal operations, simulation, and machine learning, including the modeling, detection, and disruption of criminal organization structure and activities. The two centers are working together in cross-functional and cross-organizational teams, leveraging their combined expertise and strengths to develop a methodology and apply it to specific supply chains.

In phase II of this DHS-funded project, we have studied a supply chain for medical devices critical to medical care. Medical devices play a vital role in the healthcare industry, contributing to the diagnosis, treatment, and monitoring of various medical conditions, particularly in critical-care settings where patients with life-threatening conditions require immediate and precise medical intervention. Medical devices encompass a wide range of equipment from simple tools like thermometers and blood pressure monitors to complex systems such as magnetic resonance imaging (MRI) machines and ventilators. The development, production and deployment of medical devices require a robust supply chain to ensure their availability in critical-care facilities. Such a supply chain typically involves a complex network of suppliers, manufacturers, distributors, and healthcare providers. The recent COVID-19 pandemic has significantly stressed the critical-care supply chain for medical devices, with a sudden surge of demand for devices like ventilators and personal protective equipment (PPE) [1, 2, 3]. Another reason for choosing the medical device supply chain in the second phase of this project relates to the fact that medical devices can involve the use of electronic components. The raw materials for such electronic materials depend on global sources making the device industry especially vulnerable to criminal attacks and disruptions. Further, electronic components, due to their connectivity to the internet, are candidates for remote global criminal attacks, increasing the vulnerability of this supply chain [4].

This report explains how we modeled a generic supply chain in this domain, simulated it, modeled criminal organization capabilities, and combined the two models to gain insight into the effects, indicators, and warnings of supply chain disruptions.

This project demonstrates that observable data, detailed understanding of a licit supply chain and of a criminal organization's capabilities, may be used to detect the planned, current, or past disruption of the supply chain by criminal organizations. We used simulation to identify possible ways that criminal organizations might be able to disrupt or damage important supply chains, such as those for medical devices, and compromise the service capability of healthcare providers. We examined approaches to provide warnings, and to identify possible ways to prevent, mitigate, or repair a disruption. We sought to gain understanding by modeling specific licit supply chains and criminal organization operations and capabilities, and, through both, to identify potential data sources of value and associated indicators and warnings.

We apply a combination of agent-based modeling, discrete event simulation, and stochastic optimization to model a supply chain and its interactions along with ways that criminal organizations might attempt to disrupt it. Our approach models possible attack points; develops indicators that can alert authorities and supply chain operators about a pending, active, or past attack; and offers what-if analysis to consider ways to discover and mitigate vulnerabilities to reduce attack impacts. The project uses both subject matter expertise and automated distributed optimization techniques to develop and calibrate the models.

We have developed a variety of scenarios for supply chain disruption based on the experience of our industrial partners and examples of real-world attacks, and identified specific vulnerabilities of the medical devices and critical-care supply chain being studied and potential mitigations against attempted disruptions by criminal organizations. We have studied a variety of types of attacks and identified which components of the supply chain may be targeted.

Constructing a detailed generic supply chain map for medical devices for critical-care is an important first step in understanding where disruptive vulnerabilities may exist. Based on a general methodology for building such generic supply chains, the development of such a map has been completed and vetted by partners in the industry, and a wide variety of disruption scenarios and required capabilities of criminal organizations have been identified. The supply chain and disruptions have been implemented and integrated in a simulation environment; the effects of the disruptions have been measured, and early warning signals and mitigation strategies have been studied and developed. Automated model calibration strategies, via distributed stochastic optimization methods, have been integrated with the simulation framework to analyze worst-case disruption scenarios and related optimal mitigation strategies.

A model developed earlier to understand criminal organization capabilities required to accomplish different disruptions has been applied to optimize the organization's actions to produce impactful disruptions. The point of optimizing the criminal organization's actions to produce the most impact is not that the criminal organization will have comprehensive information about a supply chain or the ability to try thousands of attacks to find the optimal use of their resources, as our simulation does. Rather, the point is to assess what happens if they *do* choose an optimal allocation, inadvertently or deliberately. That is, what does the worst case look like, even if it hasn't yet happened or even if an existing criminal organization can't execute such a disruption? This approach allows us to mitigate these potential disruptions before they happen and to develop indicators and warnings to alert us that such a disruption is pending.

Section 2 of this report describes the medical devices and critical-care supply chain model, the types of disruptions identified and studied, and the types of mitigations we have studied and that should still be studied. Section 3 goes into detail about the simulation, implementation of disruptions in the simulation, the criminal organization model and its implementation, and the implementation of mitigations in the simulation. It also discusses the deployed evaluation metric to quantify the impact of disruptions. We have used our simulation of the supply chain and our criminal organization model to run a wide variety of experiments to study the impact of different disruptions and the effectiveness of different mitigations. Section 4 describes these experiments and discusses some of the main results. In Section 5 we summarize knowledge gained about indicators and warnings that a supply chain disruption is pending, active, or has taken place. This project has given us many insights about the medical device and critical-care industry, in particular and about supply chains in general. Section 6 summarizes some of the key insights arising from our analysis. This project has only begun to scratch the surface of what a combination of supply chain expertise, expertise about criminal organization capabilities, modeling, and simulation can teach us. In Section 7, we discuss some of the things that we were unable to do given the time and resources available to us, and that we recommend doing in the future.

We believe that the results of this project will advance research into how DHS, other government agencies, academic institutions and the private sector can collaboratively build models of supply chains; pose what-if questions regarding optimal ways to disrupt them given specific amounts and types of resources and know-how; examine how to make supply chains more robust against attacks; identify and mitigate attacks as they are ongoing; and repair damage as rapidly and effectively as possible after the fact.

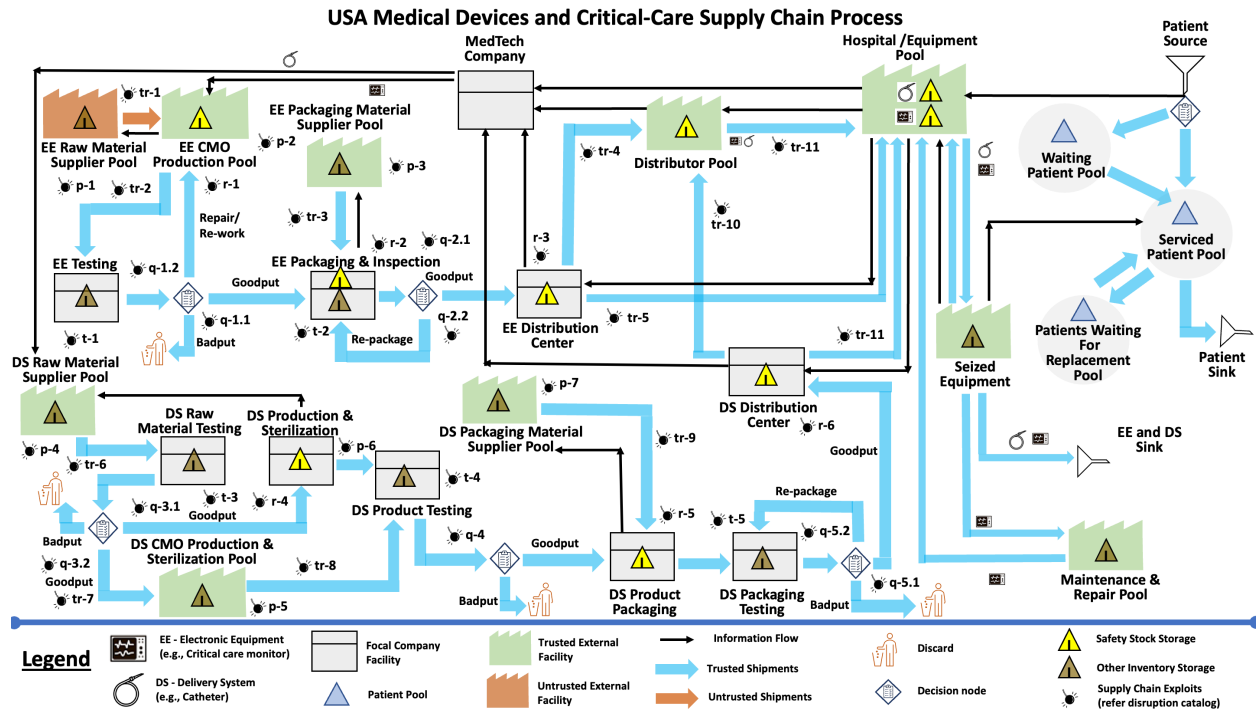


Figure 1: Layout of MedTech SC.

2 MedTech Supply Chain

2.1 Description of the Supply Chain Model

This section describes a simulation model of a medical devices and critical-care supply chain, dubbed *MedTech SC*, which was constructed in consultation with industry experts, and was programmed in the MASON software tool. The logical layout of MedTech SC is depicted in Figure 1.

MedTech SC models the supply chain for two categories of medical devices, namely, electronic equipment (e.g., medical monitor) and delivery system (e.g., catheter). The electronic equipment (EE) and delivery system (DS) are produced and transported concurrently without interaction. EE units are reusable, but have some random user-specified lifetime, while DS units are not reusable. MedTech Company is the focal company (FC) that manufactures EE and DS. The facilities of the FC are shown as grey icons, trusted suppliers and service providers as green icons, untrusted suppliers as red icons, and customer (patient) pools as grey circles. In Figure 1, the black arrows represent information flows between nodes, while blue and red arrows represent material flows between nodes, with or without delivery delay. Next, we describe the components of MedTech SC.

2.1.1 Patient Arrivals

Patients arrive according to an inter-arrival distribution. Upon each arrival, a request for a kit (1 EE unit and 1 DS unit) is sent to the Hospital/Equipment Pool. If both an EE unit and DS unit are available, the patient is routed to the Serviced Patient Pool. Otherwise, if a kit is unavailable, the patient waits in the Waiting Patient Pool until a kit becomes available, whereupon the patient is moved to the Serviced Patient Pool. The patient stays in the Serviced Patient Pool for a user-defined service duration, after which it exits the system. If a patient in the Serviced Patient Pool needs a replacement for a failed EE unit, the patient is moved to the Patients Waiting for Replacement Pool; when the replacement becomes available, it is moved back to the Serviced Patient Pool.

2.1.2 Hospital Equipment Management

On receiving a request for a kit of 1 EE unit and 1 DS unit, the Hospital/Equipment Pool tries to satisfy the request from its inventory. If a kit is available, it is moved to the Seized Equipment Pool; otherwise, the request is backlogged until a kit becomes available. Once the patient service duration terminates, the DS unit of the kit exits the system, while the EE unit of the kit returns to the Hospital/Equipment inventory. Recall that EE units can be reused by multiple patients, and DS units are only used once before being discarded. If an EE unit breakdown occurs or maintenance is required while at the Seized Equipment Pool, then that EE unit is moved to the Maintenance & Repair Pool. A replacement EE unit, if available, is then sent from the Hospital/Equipment Pool to the Seized Equipment Pool; otherwise, the kit waits for a replacement in the Seized Equipment Pool, and the patient is moved from the Serviced Patient Pool to the Patients Waiting for Replacement Pool. When the replacement arrives, The patient returns to the Serviced Patient Pool and resumes service. If the EE unit at the Seized equipment exhausts its lifetime (namely, it becomes broken beyond repair), that EE unit exits the system, and the Hospital/Equipment Pool is notified to update its inventory of equipment.

The Hospital/Equipment Pool is replenished from the following suppliers: Distributor Pool (trusted external supplier), DS Distribution Center (trusted, internal to FC), and EE Distribution Center (trusted, internal to FC). The replenishment policy at the Hospital/Equipment Pool is Make-to-Stock (MTS), separately for DS and EE, as follows:

1. When the level of any inventory (EE or DS) hits or down crosses the corresponding reorder point, the Hospital/Equipment Pool computes a replenishment size that will bring the inventory level to the corresponding target level.
2. For each unit type (EE and DS) to be replenished, the Hospital/Equipment Pool picks one of the suppliers that carries that unit type with prescribed probabilities. If the order size cannot be fully satisfied by that supplier, the supplier ships what it has at the time and logs the remainder as a backorder. When the supplier is replenished, that backorder is processed like an order (the supplier again ships what it has, and so on, until the original order is fully satisfied.)
3. The Hospital/Equipment Pool places an order with the FC to initiate the production of that order.
4. For each unit type (EE and DS) to be replenished, the FC then places corresponding production orders of the same replenishment sizes with the corresponding raw material suppliers.

2.1.3 Focal Company

The FC receives replenishment orders of EE and DS units from the Hospital/ Equipment Pool, the Distributor Pool, the EE Distribution Center, and the DS Distribution Center. The FC starts the production process of an order by relaying order information for EE and DS supplies (raw material and packaging material). For any Distributor Pool replenishment or Hospital/Equipment Pool EE order, the FC sends the corresponding orders to the corresponding EE or DS Raw Material Supplier Pool.

2.1.4 EE Production

The FC completely outsources its EE production to Contract Manufacturing Organizations (CMOs), specifically to the EE CMO Production Pool, while packaging and testing are carried out by the FC. When the EE CMO Production Pool receives an EE order from the FC, it starts a batch production using raw materials in its safety stock, while the raw materials are supplied by the EE Raw Material Supplier Pool following an MTS policy. Each completed EE batch is sent to EE Testing for testing. After testing, an EE unit can be discarded (so-called badput unit), repaired/re-worked or pass through (so-called goodput unit). Each batch of EE goodput units is sent to the EE Packaging & Inspection node for packaging using the packaging material (e.g., boxes) obtained from the EE Packaging Material Supplier Pool. An inventory of safety stocks of packaging material is maintained and replenished using the MTS policy. The packaged EE unit is tested,

and is either discarded as badput, or sent back to EE Packaging & Inspection for re-packaging, or sent as goodput to EE Distribution Center.

2.1.5 DS Production

The FC produces its DS units internally or outsources them to the DS CMO Production Pool, while packaging and testing are carried out by the FC. The FC sends DS orders to the DS Raw Material Supplier Pool, which sends the raw material for the order to the FC's Raw Material Testing facility as a single shipment. The incoming order is tested as a single batch using sampling, regardless of the order size. The resultant goodput DS batches are split between the DS Production & Sterilization node and the DS CMO Production & Sterilization Pool node. The DS Production & Sterilization node has a safety stock, managed by an MTS policy, and it uses the raw materials in the safety stock to produce the requested DS units. The DS CMO Production & Sterilization Pool node produces batches of DS units from incoming DS raw material batches. Regardless of the production venue, each DS unit is tested in the DS Product Testing before proceeding as batches of goodput units to DS Product Packaging. The packaged DS unit is then tested at the DS Packaging Testing node, and the tested DS units are either discarded as badput, or sent for re-packaging, or pass through as goodput to the DS Distribution Center.

2.2 Modeling of Disruptions

This section describes the disruptions' model. Five classes of disruptions are modeled in the simulation: disrupted production, disrupted testing, disrupted transportation, disrupted replenishment lead time, and disrupted product quality, to be further explained in the ensuing sections. A single model captures all disruption classes, and consists of two parts: a disruption timeline, and associated disruption impacts on a given disruption-specific parameter of the node where the disruption occurred. The disruption timeline consists of a sequence of three abutting time periods, called *inactive period*, *detection period*, and *restoration period*, respectively, at the end of which, the impacted parameter value may change to a user-specified value (possibly random). The time periods are as follows:

1. **The inactive period.** During an inactive period, the disruption is dormant. This period terminates at a disruption arrival time, at which point the disruption impact takes effect (the disruption-specific parameter changes) and the second period is inaugurated. The parameter change typically corresponds to reduced performance such as reduced production capacity, higher percentage of badput, or longer transportation times.

Table 1: Types of Disrupted Production.

<i>Disruption Location</i>	<i>Disruption</i>		<i>Parameter Impacted</i>
	<i>Code</i>	<i>Unit</i>	
EE Raw Material Supplier Pool	p-1	Batch	mean batch production delay
EE CMO Production Pool	p-2	Batch	
EE Packaging Material Supplier Pool	p-3	Batch	
DS Raw Material Supplier Pool	p-4	Batch	
DS CMO Production Pool	p-5	Batch	
DS Production & Sterilization	p-6	Batch	
DS Packaging Material Supplier Pool	p-7	Batch	

2. **The detection period.** During a detection period, the disruption is active but undetected. This period terminates at a disruption detection time, at which point the disruption-specific parameter may change again, and the third period is inaugurated. The parameter change may reflect a quick partial repair of the disruption’s damage or there may be no change.
3. **The restoration period.** During a restoration period, the disruption is active and is being worked on. This period terminates at a disruption restoration time, at which point the disruption-specific parameter may change again, and a new inactive period (if any) is inaugurated. The parameter change may reflect a final repair of the disruption’s damage that restores the original value of the impacted parameter, or any other “terminal” value due to incomplete restoration or possibly an upgrade.

For each disruption period, the period duration distribution and the impact on the disruption-specific parameter are user-specified. Furthermore, distinct disruptions specified by a user can occur in a stream, and the number of occurrences (finite or infinite) is also user-specified. (In our analysis, we also considered pre-disruption and post-disruption mitigations that modified these distributions, e.g., allowing for faster or slower disruption detection time.)

We next proceed to list all disruptions by category in the tables below, which specify disruption location (node in the Figure 1 diagram), code (internal code used to conveniently refer to disruptions), product unit (batch or end-unit), and the impacted parameter.

2.2.1 Disrupted Production

Disrupted production is the class of disruptions in production nodes that increase the production time of a product unit, thereby lowering the production capacity. Table 1 lists these disruptions.

Table 2: Types of Disrupted Testing.

<i>Disruption Location</i>	<i>Disruption</i>		<i>Parameter Impacted</i>
	<i>Code</i>	<i>Unit</i>	
EE Testing	t-1	End Unit	mean batch testing delay
EE Packaging & Inspection	t-2	End Unit	
DS Raw Material Testing	t-3	Batch	
DS Product Testing	t-4	End Unit	
DS Packaging Testing	t-5	End Unit	

Table 3: Types of Disruptions in Transportation.

<i>Disruption Location</i>	<i>Disruption</i>		<i>Parameter Impacted</i>
	<i>Code</i>	<i>Unit</i>	
From EE Raw Material Supplier Pool to EE CMO Production Pool	tr-1	Batch	mean batch transportation delay
From EE CMO Production Pool to EE Testing	tr-2	Batch	
From EE Packaging Material Supplier Pool to EE Packaging & Inspection	tr-3	End Unit	
From EE Distribution Center to Distributor Pool	tr-4	Batch	
From EE Distribution Center to Hospital/Equipment Pool	tr-5	Batch	
From DS Raw Material Supplier Pool to DS Raw Material Testing	tr-6	Batch	
From DS Raw Material Testing to DS CMO Production Pool	tr-7	Batch	
From DS CMO Production Pool to DS Product Testing	tr-8	End Unit	
From DS Packaging Material Supplier Pool to DS Product Packaging	tr-9	End Unit	
From DS Distribution Center to Distributor Pool	tr-10	Batch	
From DS Distribution Center to Hospital/ Equipment Pool	tr-11	Batch	
From Distributor Pool to Hospital/ Equipment Pool	tr-12	Batch	

2.2.2 Disrupted Testing

Disrupted testing is the class of disruptions in testing nodes that increase the testing time of a product unit, thereby lowering the testing capacity. Table 2 lists these disruptions.

2.2.3 Disrupted Transportation

Disrupted transportation is the class of disruptions in transportation that increase the transportation time of a product shipment, thereby lowering the transportation capacity. Table 3 lists these disruptions.

2.2.4 Disrupted Replenishment Lead Time

Disrupted replenishment lead time is the class of disruptions in replenishment that increase the lead time of a product replenishment. Table 4 lists these disruptions.

2.2.5 Disrupted Production Quality

Disrupted production quality is the class of disruptions in production that decrease the probability of goodput at a testing node, thereby reducing the product throughput. Table 5 lists these disruptions.

Table 4: Types of Disruptions in Replenishment.

<i>Disruption Location</i>	<i>Disruption</i>		<i>Parameter Impacted</i>
	<i>Code</i>	<i>Unit</i>	
EE Raw material safety stocks at EE CMO Production Pool	r-1	Batch	mean replenishment lead time
EE Packaging Material safety stocks at EE Packaging & Inspection	r-2	Batch	
EE unit safety stocks at EE Distribution Center	r-3	Batch	
DS Raw Material safety stocks at DS Product & Sterilization	r-4	Batch	
DS Packaging Material safety stocks at DS Product Packaging	r-5	Batch	
DS unit safety stocks at DS Distribution Center	r-6	Batch	

Table 5: Types of Disruptions in Production Quality.

<i>Disruption Location</i>	<i>Disruption</i>		<i>Parameter Impacted</i>
	<i>Code</i>	<i>Unit</i>	
EE Testing (Goodput)	q-1.1	End Unit	mean percentage of goodput
EE Testing (Repair-Rework)	q-1.2	End Unit	
EE Packaging (Goodput)	q-2.1	End Unit	
EE Packaging (Re-package)	q-2.2	End Unit	
DS Raw Material Testing (Goodput to DS Production and Sterilization)	q-3.1	Batch	
DS Raw Material Testing (Goodput to DS CMO Production Pool)	q-3.2	Batch	
DS Product Testing (Goodput)	q-4	End Unit	
DS Packaging Testing (Goodput)	q-5.1	End Unit	
DS Packaging Testing (Re-package)	q-5.2	End Unit	

The five classes of disruptions defined above can model a broad variety of disruptions (exploits) based on the disruption’s location and type of impact. For example, for an intentional destruction of the DS Production & Sterilization facility, one can model such an exploit by defining and entering a p-6 disruption code into the MedTech SC simulation program, and specifying the resulting production delay parameter that reflects the severity of the disruption, including its complete shutdown.

2.3 Mitigation Strategies

Managing disruptions in supply chains requires developing and deploying a portfolio of mitigation strategies that address vulnerabilities but are often associated with resource and cost investments. As part of this phase of the project, we identified a number of complementary pre-disruption and post-disruption mitigation strategies and studied their efficacy in sustaining supply chain performance when confronted with disruption(s).

Typically, mitigation strategies of supply chain disruptions can be put into a number of different categories:

1. Building Safety Inventories
2. Improving Visibility and Transparency of the Supply Chain

3. Diversifying Supply Base and Identifying Backup Suppliers
4. Emergency Contingency Planning Through Training and Simulated Exercises
5. Identifying, Categorizing and Prioritizing Vulnerabilities of the Supply Chain

We selected a variety of pre- and post-disruption mitigations to study, and focused on the speed of detection, recovery times, and impact on customers. Specifically we considered questions such as - How much did they shorten time to discover a disruption? How much did they shorten time to recover to a “normal” production level? What was the impact of the disruption on customer deliveries? We also considered additional mitigations that would be very interesting to study in the future, but might require more extensive development of our model.

2.3.1 Building Safety Inventories

A great deal of our work is focused on safety inventories. While not as prominent in medical device and critical-care supply chains as in pharmaceutical supply chains, safety stocks are still of critical importance in medical device supply chains. In our modeling, the safety stocks are held at a variety of nodes (e.g., EE production, DS Production, EE Packaging, and DS Packaging) and are of a variety of types (e.g., EE Production has a stock of raw material, EE Packaging has stocks of the EE and packaging materials, etc). Based on input from industrial partners, in our modeling, the stocks have a target level, replenished using a specified reorder point, target inventory level and a replenishment lead time with a triangular distribution. The items in each safety stock are perishable, with an expiration date after which they are discarded.

Using the simulation model, we were able to experiment with different sizes of safety stocks to understand the impacts of disruptions, for example, to study the effect of different initial sizes and different target levels of safety stock at different nodes. We studied the effect of modifying the sizes of safety stocks that trigger reorder, of invoking different replenishment protocols, and of using post-disruption strategies to draw down safety stocks at an increased speed after an anomaly is detected. Because cost limits the use of safety stock mitigation strategy, we studied different allocations of safety stocks under different budget limitations, and developed ways to “optimize” safety stock allocation. This optimization exercise provided an insight into ways to prioritize different safety stock buffers based on their mitigation efficacy.

2.3.2 Improving Visibility and Transparency of the Supply Chain

Visibility is often an important first step in managing risk within complex supply chains. Building a detailed supply chain map, as described in Section 2, includes mapping out of the different parts of the network and their inter-linkages, and leads to understanding the significance of measuring important parameters of each of the network components. Supply chain visibility requires detailed monitoring of information/data [5]. This is particularly important in the medical device critical-care sector where the production of two components, EE and DS, needs to be coordinated and synchronized. It can be achieved by using a collaborative strategy, and by the deployment of technology for relevant data sharing across the supply chain network. Other measures for improving visibility include deployment of additional/enhanced “sensors” throughout the supply chain. This can be achieved by measures such as increased testing of raw materials, more thorough and ongoing vetting of suppliers and their processes, more inspectors, improved security protocols, requiring tier 1 suppliers to provide ongoing reports on the resiliency of their processes and suppliers, adding penalties in contracts for failure to implement improved security, and the judicious use of track-and-trace systems that use smart technology.

While we did not directly incorporate specific additional sensors in the current model, we simulated the effect of additional sensors by decreasing time to detect a disruption. Of particular emphasis has been the addition of sensors at the two raw material supplier pools which would reduce the detection time for a disruption. The model provides the capability of gathering information from improved visibility under different disruption scenarios for decision makers to test the efficacy of different mitigation strategies. Increasing visibility and transparency in medical device supply chain can help improve reactive speed which is critical in responding to new, previously unknown disruptive threats.

2.3.3 Diversifying Supply Base and Identifying Backup Suppliers

“Redundancy” is a well-known strategy to achieve supply chain resilience [6]. Our models can incorporate this redundancy by increasing the number of CMOs either at the outset or through the post-disruption mitigation strategy of having additional back-up CMOs lined up in advance, with a contract in place, to add to the supply chain in case a disruption is discovered. Our simulation studies different strategies for calling upon such backup CMOs should a certain level of disruption be detected. We do recognize that these back-up CMOs would incur higher costs and increase the complexity of the supply chain, and recommend a comprehensive cost-benefit analysis as a direction for future research.

An alternative mitigation strategy is to utilize backup suppliers when the regular suppliers of EE or DS do not have enough supply on hand to meet the demand. While these backup suppliers will likely be more expensive, they could be local and called up fast. Again, these backup supplier provide reactive speed capability that is important for responding to disruptions.

2.3.4 Training and Simulated Exercises

Training and simulated exercise plans is a key part of mitigating disruptions. This is simulated through a decrease in disruption detection time in our simulation model. Another post-disruption mitigation is to increase threshold for passing tests at selected nodes once an anomaly is found. This is a topic for future research. In building our supply chain model, we placed a lot of emphasis on identifying important parameters and their approximate values/distributions, with the aid of subject matter experts from the medical device industry. These prevention mitigations involve a change of parameters, which can be examined in order to see their impact on the operation of the supply chain. Decision makers involved in training exercises could change model parameters and get an idea of the resulting impact under various "what-if" scenarios. Parameter modification to safety stock size, reorder points, and replenishment protocols either in advance or in response to detection of an anomaly is a major mechanism of mitigation to disruptions.

By measuring time until the end of a disruption in the simulation under various scenarios, we can help decision-makers understand the time and extent of recovery. If such times are unsatisfactory, this would suggest the need for alternate recovery mechanisms that are within an acceptable range for the decision-makers.

2.3.5 Identifying, Categorizing and Prioritizing Vulnerabilities of the Supply Chain

Structural and operational complexities in a supply chain make it very difficult to identify, categorize and prioritize vulnerabilities of the supply chain. Experiments that allow one to pinpoint vulnerabilities can be extremely valuable in identifying and prioritizing mitigation strategies. We conducted multiple experiments aimed at identifying which network nodes created maximum disruption in the system, thereby being high-impact, making them targets of criminal organizations. Our work also focused on identifying downstream nodes (from the disrupted node) that were most affected by the disruption. These cascading effects of disruptions are not obvious and often occur at nodes much further downstream. The simulation model offers the capability of identifying and quantifying these system-wide impacts.

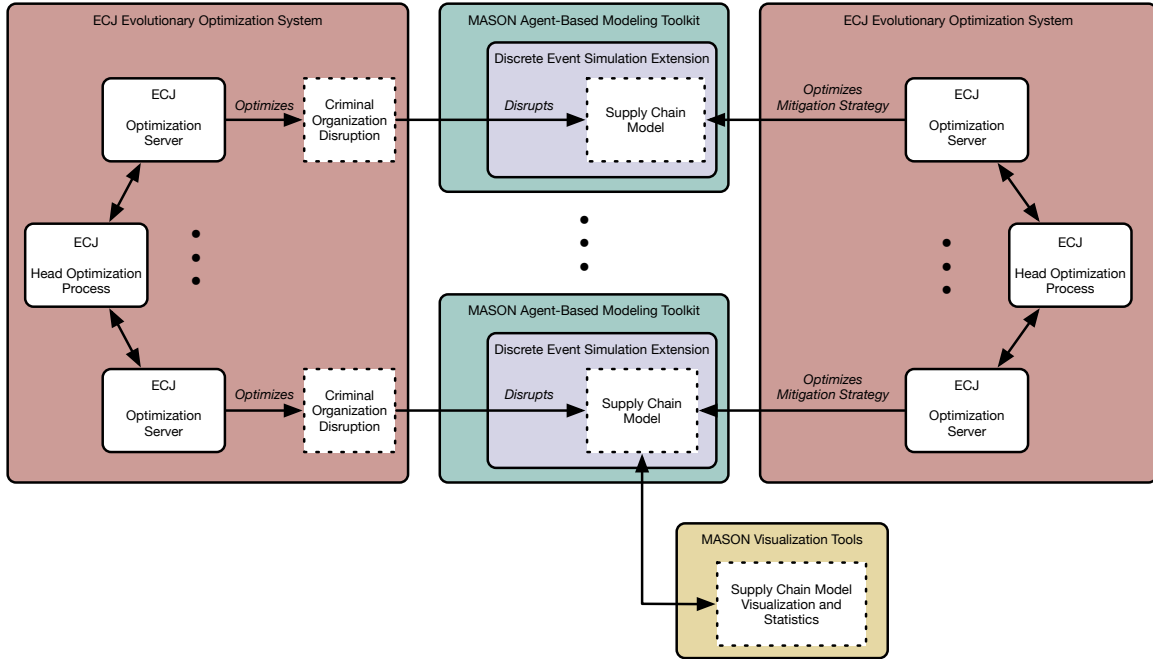


Figure 2: Overview of the Interaction between ECJ, MASON, and the Model. ECJ is used both as a potentially massively distributed optimization framework (left) to search for potential disruptions to the model using parallel supply chain model instances running on separate MASON processes (center), and is also used to optimize possible mitigation strategies (right). Models run in a custom Discrete Event Simulation (DES) environment inside MASON, and can be visualized using specialized visualization tools.

3 Simulation

3.1 Modeling and Optimization Tools Used and Developed for the Project

MASON is an open source agent-based modeling toolkit developed by Co-PI Luke [7]. It is widely used for simulation in many areas, and particularly in the social sciences. MASON is designed to run efficiently for large and complex models, or to run many models at once in parallel, and it has many visualization and statistics tools.

However, the supply-chain model we have developed is a DES model, not an agent-based model. In a DES model, various processes provide resources to other processes in real time: for example, the EE Distribution Center may provide EE to the Hospital/Equipment Pool. In an agent-based model, the processes (now called *agents*) have much more flexibility to move about and interact with one another: for example, various gangs may organize themselves dynamically in order to disrupt a supply chain in real-time. The two simulation approaches are related, and so it is possible to develop DES on agent-based modeling tools with some elbow grease. To make this much simpler, for this project we have developed an open source DES simulation toolkit

on top of MASON [8]. This toolkit allows us not only to easily build DES models in MASON, but to combine them with agent-based models in the same simulation. Figure 2 gives an overview of our system.

The supply chain model is implemented using a number of Java classes from the MASON DES simulation toolkit as its building blocks. Specifically, we have used the `Queue` class to model places where products are stored (pools and intermediate storage location between processing stages), and the `Delay` and `SimpleDelay` classes for modeling production, transportation, and testing (QA) stages of the supply chain. The model implementation is fully configuration controlled, with a configuration file or a corresponding Java object providing the information about the business rules controlling the system.

Optimization In our current approach, a criminal agent has some set of *resources*, appropriate to the type of agent, with which it may disrupt the supply chain. The question for the criminal agent is how to maximally disrupt the supply chain with the facilities available to it. This is an optimization question, and to solve it we apply a stochastic optimization method. We start with some (N) randomly-generated disruption approaches and test each one for efficacy in disrupting the supply chain model. Once we have assessed them all, we then use the optimization algorithm to produce a new generation of (N) more disruption approaches (the “children”). This is done by sampling from the space in the region of the original approaches (the “parents”), tending to sample more in the vicinity of the more successful ones. We then test the children for efficacy, and repeat. With time, each next generation improves until it reaches a global or local maximum of disruption ability.

There are many stochastic optimization algorithms we might apply: we have chosen to start with the Covariance Matrix Adaptation Evolutionary Strategy (or CMA-ES), a well regarded technique. The optimization software we use is ECJ, a popular stochastic optimization library for which MASON was designed to dovetail: indeed we have developed special open source tools which make it easy to marry ECJ and MASON. ECJ has facilities for CMA-ES and many other stochastic optimization algorithms, and like MASON it is easy to extend and modify.

To test each disruption approach we must run a separate supply chain model, which may take a few seconds: but if we like we can run them in parallel on many background MASON simulations using distributed evaluation in ECJ. Specifically, a master ECJ process builds the candidate disruption techniques (“disruption scenarios”), then ships them off to remote ECJ server processes for assessment. Each ECJ server process fires up a supply chain model on its own MASON simulation, applies the disruption, runs the model,

and assesses its performance, then reports back to the master ECJ process. This can scale to very large (many thousands) of machines in parallel if need be.

Once we have determined an optimal disruption approach, we then ask what techniques can be used to mitigate this disruption. The experimenter may set up mitigation methods manually, and may also benefit from applying an optimization algorithm to find the settings of a mitigation method which lead to optimally recover from disruption. To do this, we can again apply a stochastic optimization algorithm, and ECJ, to search for these settings. ECJ would repeatedly set up each MASON supply-chain model as before, but now it will configure each with the mitigation settings to test rather than the disruption as before: the disruption now will be hard-set to the previously discovered optimal disruption.

Competitive Coevolution An alternative to the two-part optimization approach described here is to use so-called *competitive coevolution* to optimize the disruptions and mitigation strategies simultaneously. The idea is to use two separate evolutionary processes, one optimizing disruptions and one optimizing mitigations. We test a child from one process by iteratively selecting another child from the other process, then pitting them against one another in a simulation. We then assess the performance of the child as the average, minimum, or maximum performance over simulation tests against M opposing children. The opposing optimization process assesses its children in exactly the same way. Thus as the mitigation strategies are improving, so too are the disruption approaches to challenge them. We expand on our use of this method as Future Work in Section 7.6.

3.2 Disruption Implementation

A simulation in our MASON-based application can be run jointly with a *disruption scenario* object, which specifies the list of disruptions that can occur during the simulation. Each disruption is characterized by its type, location, starting time, and magnitude. The type, at present, can be one of the following: `HaIt` (stopping the operation of a certain production nodes or information flow between nodes for a certain amount of time), `ShipmentLoss` (loss of the shipments sent between certain supply chain nodes during a certain period of time), `Adulteration` (an increase in the probability that the product batches produced on a certain day at a certain production node will be faulty), or `Depletion` (the destruction of certain amount of the product stored at a specified location). The location specifies the supply chain node affected by the disruption, e.g. the EE Raw Material Supplier Pool or the DS Production and Sterilization Pool. The meaning of the

magnitude is type-specific: for halts and shipment losses it represents the number of days the disruption lasts; for adulterations, an increase in failure rate; and for depletions, the amount of the product destroyed.

Once the disruption scenario has been loaded into the system, it affects the behavior of the targeted supply chain nodes accordingly, e.g. the DS Production and Sterilization node will shut down for a while, or will produce more bad product batches.

Table 6 outlines the different disruptions implemented in this report, along with their descriptions. 46 disruptions were implemented and studied, out of 53 designed ones.

3.3 Mitigations

At present, the main mitigation tools available to our supply chain model is the maintenance of adequate levels of extra inventory (safety stocks) and access to backup suppliers and CMOs in case of procurement or production disruptions. However, as we note in Section 2.3, other mitigations (e.g., adding more sensors), are studied through a variety of changes such as modification of parameters.

Ordering stock via backup suppliers In the normal (baseline) simulation, the focal company receives orders for the EE and DS from the Hospital Pool. It then orders the amount of materials required to fulfill these orders from its regular suppliers. In the case of a disruption, or for any other reason, where the materials delivered by the regular supplier cannot meet the required demand, the focal company has the option to order from a backup supplier. This backup supplier is considered to be local, and thus more expensive, but able to fulfill the shortfall in required materials at a faster rate. The way this mechanism works is that at each node in the production chain the company has a target level for materials required to meet the production demand. If for any reason the target level is not met and thus there is insufficient inventory to meet demand, it orders the shortfall in materials from the backup supplier. In this phase of our project we carry out experiments with and without these backup suppliers in our model in order to measure their efficacy at ameliorating the damage from disruptions to the supply chain.

Outsourcing production to backup CMOs Similar to the backup suppliers the focal company also has access to backup CMOs to which it can outsource production in case of a disruption to its own production facilities. The backup CMOs are also considered to be local, however they have a lower production capacity than the focal company's own production facilities. A key difference is that these backup CMOs can only be utilized if the focal company is able to detect a disruption is in effect at one of its production facilities. Thus, a larger detection delay will lead to a delay in the outsourcing of production and will give rise to a correspondingly larger disruption. In this phase of our project we carry out different experiments with these

Scenario	Disruption Code	Disruption Description	Exploit Description
1	D1	EE Raw Material Supplier Pool	Cyber
2	D2	EE Raw Material Supplier Pool	Physical Attack
3	D3	EE Raw Material Supplier Pool	Transportation Theft
4	D4	EE CMO Production Pool	Cyber
6	D6	EE CMO Production Pool	Transportation Theft
7	D7	EE Testing	Testing Failure
8	D8	EE Packaging Material Supplier Pool	Cyber
10	D10	EE Packaging Material Supplier Pool	Transportation Theft
11	D11	EE Packaging Material Safety Stocks	Cyber
12	D12	EE Packaging Material Safety Stocks	Physical Attack
13	D13	EE Packaging and Inspection	Testing Failure
14	D14	DS Raw Material Supplier Pool	Cyber
16	D16	DS Raw Material Supplier Pool	Transportation Theft
17	D17	DS Raw Material Testing	Testing Failure
18	D18	DS Production and Sterilization	Cyber
19	D19	DS Production and Sterilization	Physical Attack
20	D20	DS Raw Material Safety Stocks	Cyber
22	D22	DS CMO Production and Sterilization Pool	Cyber
23	D23	DS CMO Production and Sterilization Pool	Physical Attack
24	D24	DS CMO Production and Sterilization Pool	Transportation Theft
25	D25	DS Product Testing	Testing Failure
26	D26	DS Packaging Material Supplier Pool	Cyber
28	D28	DS Packaging Material Supplier Pool	Transportation Theft
29	D29	DS Packaging Material Safety Stocks	Cyber
30	D30	DS Packaging Material Safety Stocks	Physical Attack
31	D31	DS Product Packaging	Cyber
33	D33	DS Packaging Testing	Testing Failure
34	D34	EE Distribution Center to Hospital Pool	Transportation Theft
35	D35	EE Distribution Center to Distributor Pool	Transportation Theft
36	D36	DS Distribution Center to Distributor Pool	Transportation Theft
37	D37	DS Distribution Center to Hospital Pool	Transportation Theft
38	D38	DS Distributor Pool to Hospital Pool	Transportation Theft
39	D39	Hospital to Distributor Pool	Cyber
40	D40	Hospital to MedTech Company	Cyber
41	D41	Hospital to EE Distribution Center (DC)	Cyber
42	D42	Hospital to DS Distribution Center (DC)	Cyber
43	D43	EE Packaging	Cyber
45	D45	Hospital to EE Distributor Pool (DP)	Cyber
46	D46	HEP to DS Distributor Pool	Cyber
47	D47	DC to MedTech (both EE and DS orders)	Cyber
48	D48	EE DC Depletion	Physical
49	D49	EE DP Depletion	Physical
50	D50	EE HEP Depletion	Physical
51	D51	DS DC Depletion	Physical
52	D52	DS DP Depletion	Physical
53	D53	DS HEP Depletion	Physical

Table 6: Disruptions Implemented in the Report

backup CMOs in place and not in place, and we also experiment with different detection delay times, which model the impact of inaccurate or missing sensors at these production facilities.

Safety Stock Inventories The focal company also maintains extra inventory (safety stocks) of the finished EE and DS products at the EE DC (Distribution Center), DS DC, EE DP (Distributor Pool), and DS DP nodes. The Hospital Pool also maintains its own inventory of safety stocks. These safety stocks allow the supply chain to protect itself from disruptions and to also meet sudden temporary surges in demand. These levels are maintained by having a target level for products at each of these nodes that are larger than the actual demand. In this phase of the project we carry out experiments with different target levels at these nodes. We optimize the target levels to identify the minimum values necessary to mitigate worst-case disruption scenarios. This process leads to a robust cost-effective solution to protect the supply-chain against performance degradation.

3.4 Evaluating Fitness

In order to optimize the disruption caused by a criminal organization, we need a measure of the disruption of a particular attack the organization might choose. This measure, known as a *fitness function*, will be used to assess a given attack method, and the optimizer applies this function in its search for effective attacks.

Our criminal organization's fitness function is simply the daily average number of patients that could not be treated immediately and had to wait for treatment during the simulation period. It is computed as the total number of patients waiting for treatment during the simulation period divided by 2000 (i.e. the number of simulation days). This metric measures the effect of a disruption's propagation through the supply chain on the end customer. In a normal baseline simulation (a simulation without disruptions) all patients can be immediately treated once they enter the waiting patient pool. Every disruption action that causes a production halt, or a destruction or an adulteration of some product, contributes to an overall increase in the number of patients waiting for treatment in this pool. Thus the goal of the CMA-ES optimization for the criminal agent (outlined in Section 3.1) is to maximize the average number of patients who had to wait for treatment during the simulation period.

Figure 3 displays the effect of 2 different disruption scenarios, one located at the EE facility (D4) and one located at the DS facility (D33), at the waiting patient pool for a specified criminal budget. This figure illustrates the effect of different disruption scenarios on the waiting patient pool and shows that the impact of a disruption varies by scenario type.

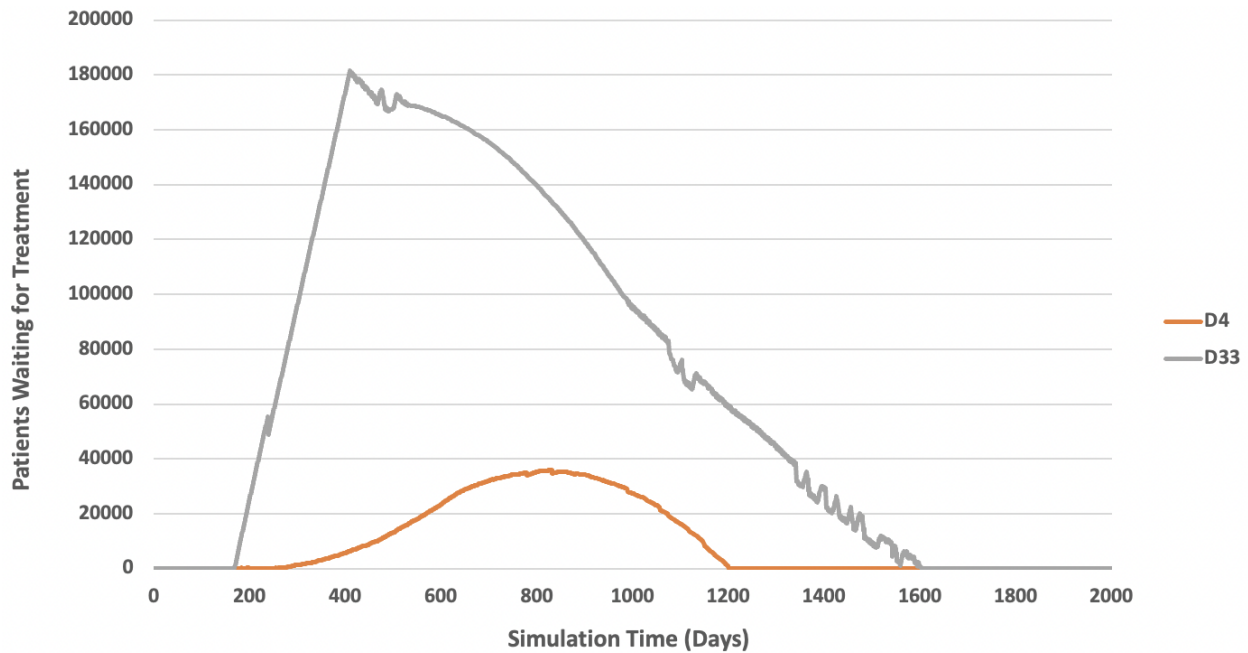


Figure 3: Time Series Data of the waiting patient pool for a disruption at the EE facility (D4) and a disruption at the DS facility (D33).

4 Experiments and Results

Using the supply chain model, simulation, and a criminal organization capability model, we ran a large number of experiments to understand which disruptions might have the most effect, which mitigations might best counter those effects, and what criminal capabilities could cause the most harm. In what follows we describe a selection of representative experiments. All of the experiments led to the various insights and conclusions described in Section 6.

4.1 Single Disruptions

In our first set of experiments, we ran simulations of the 46 disruptions (outlined in Section 3.2) individually. Each of the simulations was run for a total of 2000 days. For consistency the criminal agent is allocated the same amount of resources. We measure disruption strength based on how many days they last, and in these experiments we chose 400 days of disruption as a consistent measure of the resources expended. For each single disruption run, the disruptions were all carried out 400 days after the start of the simulation. We analyzed the effect of the damage done, by each of the 46 individual disruptions, on the supply chain performance. This effect was quantified by the fitness (as described in Section 3.4) measured at the Waiting




Figure 4: Fitness for different disruption scenarios that target the EE network of the supply chain.

Patient Pool. Figure 4 gives the effects of the different disruption scenarios that specifically target the EE network of the supply chain.

The results in Figure 4 clearly demonstrate that some types of disruptions have a much larger impact than others. Disruptions that target the EE CMO Production Pool (D4 and D6) or the EE Packaging and Inspection facility (D43) resulted in large increases in the fitness and disruption to the number of patients that can be treated. The most impactful disruptions were found to be those that target the EE DC (Distribution Center) and EE DP (Distributor Pool) facilities (D48 and D49), resulting in fitness values of over 4700 for both these scenarios. In Figure 5 we report similar results for disruptions to the DS network of the supply chain.

The results in Figure 5 show that, though there are fewer disruptions that propagate through the DS network to affect the Waiting Patient Pool and thus the fitness value, when an effective disruption scenario occurs its magnitude is significantly larger than the effect of a disruption on the EE network. For example the fitness for the most impactful EE disruption scenario (D49) is 5,567 while the corresponding most impactful disruption scenario for the DS (D51) is significantly larger at 120,000. Additionally, the average fitness for disruptions to the DS network that do have an effect is 54,708 while the corresponding number for the EE network is 2,565. These results clearly underscore the fact that the DS network of the supply chain is more susceptible to large scale disruptions that directly affect the patients. This is due to the fact that DS products are single use and must be discarded after every use, while the EE is re-usable with an average lifetime of 7 years. Thus, a disruption to the supply of EE has a far lower impact on the availability of the medical device




Figure 5: Fitness for different disruption scenarios that target the DS network of the supply chain.

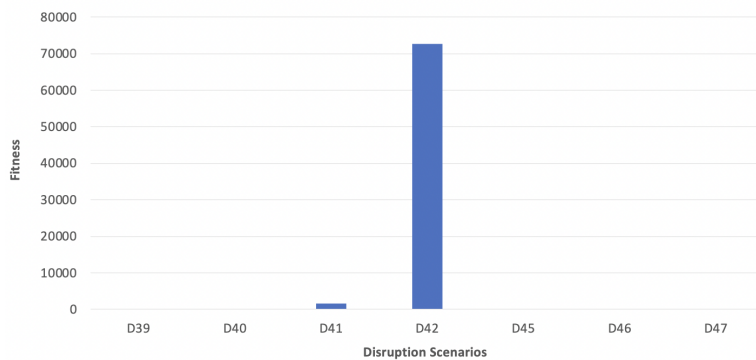


Figure 6: Fitness for different disruption scenarios that target the information flow between the Hospital Pool and the Distribution network of the supply chain.

for patients since the stock at hand can be re-used for a long period of time before new stock is required. This result was not obvious without the simulation experiments conducted in this project.

Finally we report the fitness values for disruptions that target the information flows between the Hospital Pool and Distributor Network of the supply chain. These disruptions are potential cyber exploits that prevent order and re-order information passing accurately throughout the network. The results are detailed in Figure 6.

The results in Figure 6 show that only disruption scenario D42 has any real effect on the fitness. This is because this disruption targets the information flow between the Hospital Pool and the DS Distribution Center. As noted earlier the DS network of the supply chain is most vulnerable to large scale disruptions.




Figure 7: Fitness for disruption scenario D4 with mitigation via backup CMO in place using different detection delay times.

Furthermore, most of the orders sent to the Focal Company are fulfilled by the DS Distribution Center, while the DP centers are rarely utilized by the EE or the DS network. The corresponding disruption to the EE Distribution Center (D41) has a much smaller impact, further highlighting the fact that the DS network of the supply chain is the part of the supply chain most susceptible to large scale disruptions.

4.2 Mitigation of Single Disruptions via Backup CMOs

In our next set of experiments we studied the efficacy of having backup CMOs and suppliers added to the model, in an effort to contain the damage of disruptions to the supply chain. We also investigated the effect of delays in detection due to inaccurate or missing sensors.

We first studied the impact of the disruption scenario D4, which halts production at the EE CMO Production node of the supply chain. This disruption occurs on day 1000 of the simulation. As discussed in Section 3.3 we added an additional backup CMO to the model, which can be utilized by the supply chain once it has detected a disruption. This backup CMO being local is able to ship goods faster than the regular supplier, though it has only 70% of the production capacity of the FC. We carried out experiments with different detection delay times and investigated how this affects the backup CMO's ability to mitigate the impact of this disruption. The results are reported in Figure 7.

The results in Figure 7 clearly demonstrate the advantage of having a backup CMO in place in case of a disruption. Even with a backup CMO which has only 70% production capacity of the FC's own production facility, this mitigation is able to negate the impact of a disruption to the facility if the disruption is detected with no delay. However as the delay detection time increases the magnitude of the disruption, measured via




Figure 8: Fitness for disruption scenario D12 with and without mitigation via backup supplier in place using 0 detection delay time.

the fitness function, also increases. This result illustrates the importance of not only having backup CMOs in place at the production facilities of the supply chain but also the importance of having accurate and reliable sensors in place at these points, in order to speedily detect and mitigate disruptions .

4.3 Mitigation of Single Disruptions via Backup Suppliers

We next investigated the impact of having additional backup suppliers (as discussed in Section 3.3) to mitigate the effect of disruptions that may occur to the regular supplier. Specifically, we focused on disruption scenarios D12 and D30, which represent physical attacks on the Packaging Materials Safety Stock stored at the EE and DS Packaging Materials facilities respectively. We execute both these disruptions at day zero of the simulation and investigate the efficacy of access to backup suppliers, with zero disruption detection delay time, to mitigate the impact of these disruptions. The results of these two experiments are detailed in Figures 8 and 9, which gives the fitness at the Waiting Patient Pool.

The results in Figures 8 and 9 clearly demonstrate that, if the mitigations are deployed with no detection delays, the mitigations via additional backup suppliers are able to protect the supply chain from these two disruptions. As with backup CMOs, these findings highlight the importance of having backup suppliers in place in case of disruptions to the primary supplier, and the importance of accurate and reliable sensors to identify disruptions as soon as they occur. In summary, the efficacy of backup planning is contingent on




Figure 9: Fitness for disruption scenario D30 with and without mitigation via backup supplier in place using 0 detection delay time.

early detection of disruptions. Therefore, early detection and backup planning can synergistically offer a robust mechanism to mitigate disruptions.

4.4 Optimizing Disruptions

We next applied stochastic optimization techniques to the set of disruptions, in order to learn particularly “effective” combinations of disruptions. Similar to the single disruption case, we started from day 400, and applied 400 days of disruptions. However, instead of running one disruption for 400 days, we distributed the days among disruptions (for example, run D1 for 2 days and D4 for 5 days and so on, totalling 400 across the disruptions). We refer to these day limitations as the disruption budget. The goal here was to learn what combinations of disruptions have the most adverse impact subject to budgetary constraints.

We optimized the disruptions with CMA-ES using the ECJ toolkit. The algorithm was run for 200 generations and with 35 individuals at each generation. The highest fitness function values at the Waiting Patient Pool across generations are shown in Figure 10.

After optimization, we analyzed our optimized set of disruptions. In our experiment, the optimized set of disruptions had a fitness function value of 214,328.98. We display the disruptions that last the longest in our optimized set in Table 7. The optimized disruption set focuses on the Hospital, depleting its equipment and interfering with its communication with the MedTech company, preventing the Hospital from ordering




Figure 10: Fitness over 200 Generations of Most Impactful (“Best”) Set of Disruptions

<i>Disruption Code</i>	<i>Disruption Location</i>	<i>Day Allocation</i>
D50	EE HEP Depletion Physical Attack	80
D40	Hospital to MedTech Cyber	77
D28	DS Packaging Material Supplier Transportation Theft	17

Table 7: Top Disruptions in the Optimized Set.

more equipment. The combination of these disruptions is important; we note that the disruptions with the most adverse impact found through this approach do not perform well as individual disruptions. This indicates that when the actions of weak individual disruptions are combined and properly coordinated, their collective disruptive effect on the supply chain can be greatly strengthened. In addition, our optimized set of disruptions “outperforms” (i.e. creates a higher impact) by a large margin over the case where the entire budget is allocated to any single disruption. This suggests that combinations of disruptions can be particularly dangerous to the supply chain. This finding is consistent with the real-life situations where criminal organizations capitalize on an existing natural disruption by injecting counterfeit products or materials.




Figure 11: Fitness over 200 Generations of Most Impactful (“Best”) Set of Target Levels

4.5 Optimizing Mitigations

The motivation for studying disruptions is to learn the best way to protect the supply chain against their adverse impact. Therefore, we next focus on mitigating effective disruptions using optimization. To this end, we repeated CMA-ES with the same settings as before, but this time optimizing the ten target levels for on-hand supplies (safety stocks), using 20 individuals per generation. We used a budget of the total sum of target levels in our default configuration, and the goal was to find a reallocation of these resources in order to best mitigate the optimized disruption. We also changed the reorder point to half of the target value, and initial values to the target value, as these values need to be adapted to match the new target values. We show the “best” fitness function value (in this case, a lower average number of waiting patients is better because we are optimizing mitigations) across 200 generations in Figure 11. We show the optimal target level allocations in Table 8. For comparison, we also include the original target levels as well, which were chosen during the design of the supply chain, and used in the other experiments.

The optimized target levels are very effective at mitigating the optimal disruption set, reducing the average patients waiting from 214,328.9 (the fitness based on the original target values, reorder points and

Location	Original Target Level	Optimized Target Level
EE Manufacturing Safety Stock	10362	41266
EE Packaging Safety Stock	1500	20496
EE Distribution Center	1000	55565
EE Distributor Pool	1036	28221
EE Hospital/Equipment Pool	1000	26062
DS Manufacturing Safety Stock	65268	2792
DS Packaging	65268	42996
DS Distribution Center	70707	20056
DS Distributor Pool	7071	2032
DS Hospital/Equipment Pool	72528	56253

Table 8: Optimized Target Levels

% of Original Target Budget	Fitness Value Against Most Impactful Disruption Set
100 %	308.40
85 %	412.89
75 %	659.28
50 %	1228.20
25 %	4563.56
10 %	7720.03

Table 9: Optimized Reduced Target Levels

initial values) to 308.40, a sizable improvement. The optimal target levels reallocate a large amount of stock from DS to EE sections of the supply chain. This makes sense, as the optimal combination of disruptions targets the EE side of the Hospital, so it makes sense that the model would reinforce this area to compensate. This demonstrates the effectiveness of optimizing target levels to reduce disruption damage.

4.6 Reducing Target Levels

Under budget constraints for resources and safety stocks, it's very important to be able to allocate the limited available resources to best protect the supply chain, and in general to understand the cost-benefit trade-offs under different constraint-level scenarios. In this last set of experiments we tackled this issue. Besides using optimization techniques to figure out where best to allocate resources, we repeated the optimization process on reduced target levels, to see how well we can mitigate the optimized disruption set with fewer resources. The results of these experiments are shown in Table 9.

The results demonstrate that, while reducing the budget has a large impact on the overall fitness value, optimized allocation leads to a much stronger fitness value than not. Even 10% of the target budget optimized has a much lower average number of waiting patients than a disrupted supply chain with the default, non-optimized target values. This demonstrates the importance of optimized allocation, both in

terms of protecting the supply chain and in potentially saving resources. This generally makes sense, as the disruption set was optimized against the original target levels, and so is expected to be effective against them. This experiment can also be expanded with coevolution (see Section 7.6 for details), where the set of disruptions are evolved to be more effective as the target levels are also evolving. With coevolution, testing limited target levels on a variety of improving disruptions could show how far reduced budgets can go to protect the supply chain.

5 Indicators and Warnings

One purpose of this report is to demonstrate the ability to generate Indicators and Warnings (I&W) that could be used to detect past, ongoing, and especially future supply chain disruptions by criminal organizations. For the purposes of this study, we distinguish between I&W from the viewpoint of the supply chain operator and externally observable I&W. This section focuses on externally observable I&W. Detection of disruptions from the supply chain operator's point of view has been discussed previously and elsewhere.

5.1 Indicators and Warnings Based on External Observables

A key idea behind this study is that the modeling and simulation efforts allow us to identify supply chain vulnerabilities and the potential damage that a particularly resourced criminal organization might inflict on that supply chain. The most damaging or otherwise interesting simulated disruptions can be traced back to specific criminal organization capabilities and activities, which form the basis for associated I&W.

Each of the 46 disruptions considered in this study can be associated with specific criminal organization capabilities in the form of geographic footprint, financial and other resources, expertise, manpower, etc. These capability requirements are listed below for the selected disruptions and were previously developed and delivered to the DHS Office of University Programs as the document "Criminal Organization Capabilities and Disruptions".

In this section, we consider potential I&W for the six most impactful disruptions identified in the experiments. For each scenario, we identify the disruption type (see section 3.2 and Table 6) and its impact (as a fitness score, described previously, at the Waiting Patient Pool), the criminal organization capabilities required for the disruption, and I&W examples for that disruption.

5.1.1 Scenario 1: Depletion (destruction in storage)

Disruptions D48 and D49 (against the Distribution Center and Distributor Pool of the EE component of the supply chain) and D51 (against the DS component of the supply chain) produced disruption fitness of about 5k, 5k, and 120k respectively (see Figure 4 and Figure 5). These disruptions involve the physical destruction of product stored at a specified location. A criminal organization would require the following capabilities in order to effect such a disruption:

- *Geographic footprint (size, area name)*: local, near storage locations
- *Financial resources (scale)*: modest
- *Shipping resources (range, scope/amounts)*: N/A
- *Explosives and weapons (weapon, delivery range)*: explosives; immediate area or greater range
- *Cyber expertise (type, compromise details)*: N/A
- *Manpower (scale)*: minimal
- *Manufacturing (product expertise, people, equipment)*: minimal, some, none
- *Local government relationship (to org, to US)*: not directly corrupt but supportive/friendly, neutral

Considering these capability requirements and the nature of physical destruction of stored product, possible I&W related to this disruption focus on a local criminal organization with knowledge of the storage location, awareness of the value of the stored product to the supply chain, enough explosives or similar to destroy the stored product, manpower to effect the theft destruction. Specific I&W examples follow:

- Theft or disclosure of building plans or physical security mechanisms; exposure of same in the public domain or other venues (e.g., dark web).
- Surveillance or physical security probes of storage locations.
- Purchase of or access to explosives.
- New facility employees, especially without sufficient background checks; recently terminated facility employees.
- Lack of other criminal organization activity in the area.
- Local government bribe activity and/or lack of local enforcement presence.

5.1.2 Scenario 2: Halt (information flow)

Disruption D42 (cyber attack against the information flow between the Hospital and DS Distribution Center) produced a disruption fitness of about 70k (see Figure 6). A criminal organization would require the following capabilities in order to effect such a disruption:

- *Geographic footprint (size, area name):* N/A
- *Financial resources (scale):* minimal
- *Shipping resources (range, scope/amounts):* N/A
- *Explosives and weapons (weapon, delivery range):* N/A
- *Cyber expertise (type, compromise details):* N/A
- *Manpower (scale):* minimal
- *Manufacturing (product expertise, people, equipment):* minimal, some, none
- *Local government relationship (to org, to US):* not directly corrupt but supportive/friendly, neutral

Considering these capability requirements and the nature of a cyber exploit for disruption purposes (which is different than a cyber exploit for process manipulation or similar purposes), possible I&W related to this disruption focus on criminal organization cyber expertise and their process and system knowledge and access. Specific I&W examples follow:

- Acquisition of or access to cyber expertise (potentially as affiliate to an established ransomware gang or similar).
- Acquisition of or access to knowledge of critical communication paths.
- Cyber reconnaissance against systems.
- Probes or attacks against similar systems elsewhere.
- Existing or recent system compromise, IT infrastructure exposure, password/data leak, or similar.
- History of complicity/inaction by criminal organization's relevant authorities.

5.1.3 Scenario 3: Halt (product packaging - cyber)

Disruption D31 (cyber attack against the DS product Packaging Pool) produced a disruption fitness of about 75k (see Figure 5). A criminal organization would require the following capabilities in order to effect such a disruption:

- *Geographic footprint (size, area name):* N/A
- *Financial resources (scale):* minimal
- *Shipping resources (range, scope/amounts):* N/A
- *Explosives and weapons (weapon, delivery range):* N/A
- *Cyber expertise (type, compromise details):* remote denial of service, remote compromise
- *Manpower (scale):* minimal
- *Manufacturing (product expertise, people, equipment):* minimal, some, minimal
- *Local government relationship (to org, to US):* not directly corrupt but supportive/friendly, neutral

Considering these capability requirements and the nature of a cyber exploit for disruption purposes (which may involve a cyber exploit since the packaging system is likely protected), possible I&W related to this disruption focus on criminal organization cyber expertise and their process and system knowledge and access. Specific I&W examples follow:

- Acquisition of or access to cyber expertise (potentially as affiliate to an established ransomware gang or similar).
- Acquisition of or access to knowledge of process or systems.
- Cyber reconnaissance against systems.
- Probes or attacks against similar systems elsewhere.
- Existing or recent system compromise, IT infrastructure exposure, password/data leak, or similar.
- History of complicity/inaction by criminal organization's relevant authorities.

5.1.4 Scenario 4: Halt (product packaging - testing)

Disruption D33 (attack against the DS product Packaging Testing Pool) produced a disruption fitness of about 35k (see Figure 5). A criminal organization would require the following capabilities in order to effect such a disruption:

- *Geographic footprint (size, area name):* local, near packaging testing facility
- *Financial resources (scale):* modest
- *Shipping resources (range, scope/amounts):* N/A
- *Explosives and weapons (weapon, delivery range):* N/A
- *Cyber expertise (type, compromise details):* remote denial of service, N/A
- *Manpower (scale):* minimal
- *Manufacturing (to org, to US):* advanced, moderate, minimal
- *Local government relationship (to org, to US):* not directly corrupt but supportive/friendly, neutral

Considering these capability requirements and the nature of disrupting the testing process for delivery system packaging via an intrusion, supply injection, or an insider, possible I&W related to this disruption focus on a local criminal organization with knowledge of the facility and testing process, some modest financial resources, and the expertise to disrupt the package testing process. Specific I&W examples follow:

- Theft or disclosure of building plans, physical security mechanisms, or testing process; exposure of same in the public domain or other venues (e.g., dark web).
- Surveillance or physical security probes of facility locations.
- Purchase of or access to testing equipment or consumables.
- New facility employees, especially without sufficient background checks; recently terminated facility employees.
- Change of vendor or shipping means for testing equipment or consumables.
- Lack of other criminal organization activity in the area.
- Local government bribe activity and/or lack of local enforcement presence.

5.1.5 Scenario 5: Halt (product production and sterilization - cyber)

Disruption D18 (cyber attack against the DS product Production and Sterilization Pool) produced a disruption fitness of about 40k (see Figure 5). A criminal organization would require the following capabilities in order to effect such a disruption:

- *Geographic footprint (size, area name):* N/A
- *Financial resources (scale):* minimal
- *Shipping resources (range, scope/amounts):* N/A
- *Explosives and weapons (weapon, delivery range):* N/A
- *Cyber expertise (type, compromise details):* remote denial of service, remote compromise
- *Manpower (scale):* minimal
- *Manufacturing (product expertise, people, equipment):* advanced, moderate, minimal
- *Local government relationship (to org, to US):* not directly corrupt but supportive/friendly, neutral

Considering these capability requirements and the nature of a cyber exploit for disruption purposes (which would likely involve a cyber exploit since the production and sterilization systems are likely protected), possible I&W related to this disruption focus on criminal organization cyber expertise and their process and system knowledge and access. Specific I&W examples follow:

- Acquisition of or access to cyber expertise (potentially as affiliate to an established ransomware gang or similar).
- Acquisition of or access to knowledge of process or systems.
- Cyber reconnaissance against systems.
- Probes or attacks against similar systems elsewhere.
- Existing or recent system compromise, IT infrastructure exposure, password/data leak, or similar.
- History of complicity/inaction by criminal organization's relevant authorities.

5.1.6 Scenario 6: Combined attack

A significant disruption was effected by sequentially combining disruptions D50, D40, and D28 (see Section 4.4, Table 7 and Figure 10). These particular experiments applied different disruptions to both the electronic equipment and the delivery system components of the supply chain. These disruptions involve physical, cyber, and transportation attacks. A criminal organization would require the following capabilities in order to effect such a combined disruption:

- *Geographic footprint (size, area name)*: multi-national, near facilities and routes
- *Financial resources (scale)*: significant
- *Shipping resources (range, scope/amounts)*: N/A
- *Explosives and weapons (weapon, delivery range)*: explosives; some range
- *Cyber expertise (type, compromise details)*: remote denial of service, remote compromise
- *Manpower (scale)*: significant
- *Manufacturing (product expertise, people, equipment)*: advanced, significant, moderate
- *Local government relationship (to org, to US)*: corrupt and supportive, not friendly to US

Considering these capability requirements and the nature of these combined attacks, possible I&W related to this disruption focus on a large criminal organization with significant funding, multiple skill sets, knowledge of the supply chain processes and flows, enough explosives or similar to destroy product and facilities, and manpower to effect the attacks. Specific I&W examples follow:

- Theft or disclosure of building plans or physical security mechanisms; exposure of same in the public domain or other venues (e.g., dark web).
- Surveillance or physical security probes of storage locations.
- Purchase of or access to explosives.
- New facility employees, especially without sufficient background checks; recently terminated facility employees.
- Lack of other criminal organization activity in the area; lack of this organization's activities elsewhere.

- Acquisition of or access to cyber expertise (potentially as affiliate to an established ransomware gang or similar).
- Acquisition of or access to knowledge of process or systems.
- Cyber reconnaissance against systems.
- Cyber or physical probes or attacks against similar systems elsewhere.
- Existing or recent system compromise, IT infrastructure exposure, password/data leak, or similar.
- History of complicity/inaction by criminal organization's relevant authorities.

6 General Insights and Conclusions from the Study

Medical device supply chains (MDSCs) are part of the national critical infrastructure since the products they deliver are vital in maintaining public health and safety in hospitals, clinics, long-term care facilities, and at-home. Supply or demand disruption to such a supply chain can limit the availability of critical medical devices, which can have serious and life-threatening consequences for patients. PPE and ventilator shortages experienced during the COVID-19 pandemic are examples of such shortages with severe health consequences. This study showed the relevance and importance of developing rigorous methodologies, tools, and processes for helping craft better-informed policies. Through the development of the supply chain map, classifying various types of disruptions, and overlaying these disruptions on the supply chain map, a simulation model was created. This simulation was parameterized with the help of industry experts, which allowed for conducting experiments under different baseline and disruption scenarios, analyzing the efficacy of detecting the disruptions, and analyzing the supply chain performance using multiple metrics. From mapping the supply chain to conducting and analyzing simulation experiments, this entire exercise helped us gain valuable insights. Importantly, many of these insights from studying MDSCs confirm the lessons learned from the Pharmaceutical Supply Chain work done earlier. Indeed, finding a reoccurring pattern of insights across the two supply chains helps validate the robustness of our findings, consistent with the methodology of building theories from case-study research [9].

1. One unique aspect of the medical device supply chain is the maintenance/repair of devices while in use. Given EE devices' reuse and relatively long life, maintenance and repair may be needed. The industry practice is to reactively send an EE device for repair once a fault or malfunction is detected,

often during deployment. This provides new opportunities for criminal disruption. We think that proactive preventive maintenance and regular software updates of the equipment can be a tremendous operational advantage in increasing the availability of properly operating EE devices and in minimizing opportunities for criminal disruptions such as cyber attacks. In fact, such a preventive maintenance operation can function much like a backup CMO since it ensures the availability of working equipment in stock. Given the relatively high cost of EE equipment, such preventive maintenance and updates can offer high benefit to cost value. On the other hand, DS components in medical device supply chains are often single-use, not requiring maintenance and repair. Therefore, in the case of DS, backup CMOs can offer considerable value by streamlining supply chain operations.

2. The simulation experiments we conducted showed the value of backup CMOs in reducing the impact of a disruption. Backup CMOs are called upon for a speedy response in emergencies, though at a cost. Our work showed that even when these backup CMOs had limited capacity (compared to the Focal Company's capacity), they still played an important role in reducing the severity and duration of the impact of the disruption. The simulation experiments from this phase of the study also showed that the efficacy of this backup planning (of backup emergency CMOs) depends on an early detection of the disruption itself. In other words, backup CMOs would not be as effective if there was a long delay in detecting a disruption and invoking the services of the backup CMO. Indeed, the backup CMO planning has to work in conjunction with early detection and communication systems. Therefore, early detection and backup planning can synergistically offer a robust mechanism to mitigate disruptions.
3. A key strategy for mitigating disruptions is maintaining high safety-stocks levels and end-product inventories [10]. The process of determining safety inventories for various products and materials is based on values and variability of factors such as lead times, costs, and number of suppliers. Accordingly, these safety inventories can vary considerably from node to node in the supply chain network. Typically inventories of some items, such as packaging, might be kept low, anticipating ease and speed of replenishment. Our current work on medical device supply chains shows an industry-wide emphasis on 'efficiency' (costs) and not as much on responsiveness (speed). Responsiveness in these supply chains depends mainly on inventories, which may not be appropriate for an unanticipated emergent need. The COVID-19 pandemic has shown that some of the serious policy issues facing supply chains during an emergency are ensuring adequate supply, equitable distribution/allocation, regulatory compliance, and oversight to ensure patients' health and safety. If we have learned anything

from the pandemic, it is the need to be able to rapidly produce existing and new medical devices (such as testing kits) and distribute these speedily and equitably. Indeed, there is a need for such supply chains to have a two-pronged supply chain strategy: an efficient supply chain that is used in normal regimes where supply and demand follow forecasts and a responsive supply chain ready to be deployed on short notice in emergent situations. Such a rapid-response supply chain should have the ability to tap into excess capacity, shorten lead times, quickly develop and deploy compliance protocols for new products, deploy new resources for manufacturing and compliance, adapt to shortages and disruptions through already developed contingency plans, and have a high degree of awareness and capabilities to rapidly detect and remove fraudulent/counterfeit products. For such planning to take place, there could be a need for a governmental policy that goes beyond inventories and incentivizes the speed and responsiveness of supply chains for such items as medical devices, pharmaceuticals, and other critical supplies. Furthermore, such a policy should have a preparedness plan process for identifying and allocating scarce resources effectively, equitably, and transparently.

4. Reliable safety inventories are important strategic areas of strength that ensure effective functioning of the medical device industry. Because of their strategic importance, compromising the availability or effectiveness of these strategic assets, individually and collectively, can be a source of significant disruption if they are impacted. Disruption to the availability of safety stocks could seriously impact the supply chain operations. While the results presented in this report do not directly model these disruptions, the simulation model we have developed does have the capability of studying the impact of multiple disruptions. Further enhancements to the model could consider investigating the impact of (partially or fully) compromised items in safety inventories. This “bad” stock of safety inventory could be due to a scenario where the existing testing protocols were not geared to detect a newly discovered (natural or criminally intended) cause.
5. High inventory levels and a primary focus on end-product inventories often mask early detection of disruptions, which is important in limiting the impact of disruptions as shown in our experiments. High inventory levels and primary attention to end-product inventory might lead to delays in recognizing early warnings. This suggests that medical device supply chain stakeholders implement policies, procedures, and technologies that enable early detection of disruptions through real-time and system-wide warning systems that trigger a rapid investigation and if needed, mitigation responses. One way this could be accomplished is by installing multiple “sensors” throughout the supply chain and

specifically placing measurement sensors directly in safety stocks, where their usage patterns and testing outcomes are monitored (in addition to the measurements in the main supply chain itself). More generally, there is significant value to continuous monitoring of key indicators such as inventory levels, waiting times, testing-node outputs, flow times, etc. Because effects of disruptions are carried forward to other nodes, having “sensors” of numerous kinds everywhere is important. The first anomaly detected may not be the most important one, but it is important as an early indicator and warning. The effect of not having sensors at early nodes in the supply chain can lead to lengthened time to detect anomalies at later nodes. Operationally, the installation and monitoring of these sensors will require upfront and ongoing commitment of resources. However, leveraging existing technology could keep the costs low, especially with an increasing scale of deployment, offering a high benefit-to-cost ratio.

6. The actual levels of safety stocks should not be considered in isolation but as holistic management of a strategic asset. Specifically, these safety stocks require the management of a variety of important levers, such as replenishment cycle times, reorder levels, and order quantities. Actively managing these levers as decision variables, individually and collectively, will allow medical device supply chain firms to respond effectively to various levels of disruptive threats.
7. Information flows are an essential and integral component of supply chain strategy, ensuring coordination and synchronization. The modeling and simulation we have carried out provide a deeper understanding of the importance of information flows within medical device supply chains. For example, to maintain the timely availability of finished goods to end customers, the EE/DS production processes must be synchronized and coordinated to produce a complete, usable medical kit (one EE and one DS per incoming patient). Indeed, the information flows in the focal company supply chain model support the physical/material flows and need to be reliable and accurate. Further, these information flows are critical in ensuring compliance with regulatory agencies such as the FDA. The continual availability of accurate information is especially critical to the effective functioning of medical device supply chains and, indeed, most modern supply chains. Accuracy of and trust in data are critical in monitoring a supply chain and in getting reliable indications and warnings of a disruption or planned disruption. A disruption that replaces accurate data with inaccurate data could be difficult to discover and lead to major disruptions. A disruption that creates distrust in underlying data could accomplish much the same thing. Therefore, the importance of reliable information flows in supply chains where compliance is paramount cannot be overstated.

8. Because of the importance of data and information, there is a need to ensure accurate and protected knowledge across disruption categories, locations, impacts, and mitigations. Such knowledge sharing can be challenging among companies that often compete with each other. Government agencies can play an important role by being neutral brokers for collaboration. We believe there is a need for a government-facilitated industry-wide collaborative platform that can rapidly share information, adapt, and respond in emergencies.
9. For the highly regulated medical device industry, significant effort and resources are spent in preventing, tracking, and remediating external failures (i.e., where the product has reached parts of the supply chain outside the organization, closer to or with the end customer). Accordingly, the developed medical device supply chain model incorporated extensive testing nodes throughout the process consistent with the industry practice. As with safety stocks, compromising the availability or effectiveness of testing competencies could be a source of significant disruption. Disruption to the reliability of testing could seriously impact the supply chain operations. Ongoing risk assessment and periodic checking of the efficacy of testing at various supply chain components are critical operational safeguards for medical device manufacturers to put in place. Such safeguards should not be limited to the focal company alone but should also involve a comprehensive risk assessment of external parties such as material suppliers and CMOs, which could be sources of significant vulnerabilities.
10. Making a supply chain more visible and transparent can be vitally important in identifying its potential vulnerabilities [11]. One of this project's key contributions was mapping the process flows and interlinkages in the medical device supply chain. The process linkages among the focal company, CMOs, suppliers, distributors, and customers in a visual representation were valuable to understanding flow patterns, bottlenecks, safety measures, and vulnerabilities. Visibility and transparency of external components of the supply chain are as critical as those of the internal components. This visibility and transparency can be achieved by strategic and purposeful strategic partnerships with external suppliers and CMOs leveraging collaborative and data-sharing platforms.
11. Our earlier report on pharmaceutical supply chains discussed the importance of developing anomaly detection tools. Our work on medical device supply chains further reinforced this earlier finding. Anomaly Detection leverages the multiplicity of sensors throughout the supply chain. An Anomaly Detection System (ADS) needs to be carefully empirically calibrated and adjusted, keeping in mind "known" and "knowable" threats. The efficacy of early warnings and alerts depends on the empirical

and real-time adjustment of ADS parameters. Such a real-time ADS system is very important for medical supply chains, where hospitals may depend on these devices to provide critical care to patients. Linking the ADS system to a catalog of threats provides a mechanism for the early detection of a variety of anomalies throughout the supply chain. We believe a built-in, integrated “learning” component in the ADS would be especially valuable to managers and a worthwhile direction for a future investigation.

12. Anomalous order levels should trigger review and possible intervention. Too often, these anomalous order levels are noted by sales teams in a company but not shared with the security teams in a timely way. Much like pharmaceutical companies, medical device companies and their supply chain partners would benefit from putting together a dedicated inter-organizational team that studies, understands, and continuously updates about the past, current, and emerging capabilities of criminal organizations. These teams could develop the capability of leveraging “Honey-pot” sting operation mechanisms to gather intelligence on emerging or planned disruptive threats. The team members could then be embedded in their respective organizations’ strategic and planning decision-making processes to help prevent, alert, and respond to disruptive threats from criminal organizations. The threats imposed by criminal organizations in the medical device industry are often enabled by the presence and usage of sub-par devices in developing countries where regulatory standards are not in place.
13. While much effort has focused on reducing risks associated with physical aspects of the supply chain, less attention has been paid to developing cyber-resilient supply chains. A cyberattack on the information infrastructure of a medical device supply chain could impact the physical flows and jeopardize regulatory compliance [12]. Given the importance of this problem, medical device supply chain stakeholders should develop automated systems that continuously monitor physical flows/inventory levels and compare them with digital data to uncover data disruptions and missing inventory (possibly stolen) as early as possible. In addition, given the increasing challenge of cyberattacks, cyber-defense agreements with external parties and regular monitoring of those agreements are critical.
14. Medical device supply chain stakeholders and policymakers could use standardized quality systems and disruption metrics to compare vulnerabilities, mitigation strategies, and the impact of structural changes in a supply chain. When integrated into simulation models, these metrics could help measure the impact of disruptions via controlled in-vitro experiments, identify critical vulnerabilities, efficacious mitigation strategies, and promising structural or strategic changes for preventing future disruptions.

In addition, combining the standardization of disruption classification and disruption metrics could help guide policy decisions more effectively and systematically. There is a need for developing industry-wide policy guidance on reliable metrics for measuring the impact of disruptions on medical device supply chains.

15. The simulation models developed for the medical device supply chains can be useful tools for understanding system-wide behavior, which can be a strategic benefit. For example, injecting disruptions in the EE system or DS system individually and then comparing them with a simultaneous disruption in EE and DS systems could help decision-makers understand the underlying dynamics and interactions of the two systems and their impact on the availability of the end product. This deeper human causal understanding of the dynamics is critical in developing proactive and reactive strategies for preventing and mitigating disruptions. A causal understanding is especially useful in finding leverage points in the supply chain where the allocation of limited resources can yield a disproportionate benefit.
16. As noted for pharmaceutical supply chains, supply chain managers concerned with maintaining continuity of production and maximizing sales revenues may make the organization less willing to accept the cost of false alarms, leading to “missing” early warnings of a disruption. Accepting more false alarms has a cost, but it is important to recognize that sometimes this cost is much less than missing a real disruption, which requires an organizational and managerial shift in mindset.

7 Future Work

Building and analyzing a supply chain for medical devices and the critical-care industry turned out to be both a rewarding and useful exercise, which gave us many important insights into defense against disruptions (see Section 6). However, it also suggested the need to take the work further. Here we summarize a few of the ideas that we feel are important to pursue, but which we were not able to pursue given the time and resources we had available to us.

7.1 Safety Stocks

The medical device industry has recognized the importance of safety stocks to achieve resilience and our work has reiterated the importance of safety stocks. In the MedTech supply chain, most production nodes have “virtual safety stocks” for their input materials. That is, there is no segregated “safety” stockpile of,

say, a raw material stored in addition to the “regular” inventory (fed by the normal supplier). Instead, a factory’s raw material warehouse has a business rule with two parameters, the target level and the reorder point. When it is detected that the inventory level has dropped below the reorder point (e.g. because of disruptions involving the “normal” supplier), a shipment is requested from an “alternative supplier”, in an amount necessary to bring the inventory level back to the target level.

Disruptions to safety stocks represent a real and serious risk to the operation of supply chains such as those for medical devices. Of particular interest are attacks not on the safety stocks themselves, but on the data available about size of safety stocks or on the trustworthiness of the data. The set of disruption types available in our model includes “DisableTrackingSafetyStock”. A disruption of this type, applied to a particular warehouse, prevents it from sending any requests to the alternative supplier during a specified time interval. Such a disruption may correspond to a breakdown in the warehouse’s own inventory tracking system, a breakdown of communication with the alternative supplier, or a breakdown at the alternative supplier itself, or a criminal attack on the tracking system itself. Future work should go further to enable and then investigate other types of disruptions to the safety stock or its data integrity.

We studied the idea of modifying the parameters that control safety stock replenishment (the target level and the reorder point). Because cost is a major issue limiting the use of safety stock mitigation strategy in particular, we studied different allocations of safety stocks under a budget limitation, and developed ways to “optimize” safety stock allocation. Results from this optimization exercise not only help in a judicious allocation of resources but provide an insight into the relative prioritization of the safety stock buffers based on their mitigation efficacy. However, further work with benefit-cost analysis and optimization of disruption-defense from the use of safety stocks will be a very promising direction for future research, especially with interactive capability for industry experts to explore and evaluate scenarios.

7.2 Sensors

Much of our analysis of the MedTech supply chain suggested the importance of different kinds of “sensors” that would help prevent disruptions or provide early warning of disruptions and, more generally, provide improved visibility and transparency of the supply chain. Such “sensors” could correspond to increased testing of raw materials, more thorough and ongoing vetting of suppliers and their processes, more inspectors, improved security protocols, requiring tier 1 suppliers to provide ongoing reports on the resiliency of their processes and suppliers, adding penalties in contracts for failure to implement improved security, and the judicious use of track-and-trace systems that use smart technology. We did not include explicit sensors in the

current MedTech model or simulation, though we did include testing at each stage of the process. Instead, we modeled the effect of additional sensors indirectly by decreasing time to detect a disruption. Not all of these notions of sensor would be easy to include in an enhanced supply chain model or simulation, but we believe some of them could be done, allowing for developing and testing a variety of additional mitigations.

The MedTech model includes the process for all nodes requesting and receiving shipments to monitor their arrival, and to match received shipments against orders. Each node can therefore detect that a particular order has not been filled by its expected arrival date, declare it as "missing", and then cancel the original order (to prevent accidental double-ordering) and place a new one. This is a fairly effective countermeasure against disruptions associated with shipment losses. Our framework enables one to experiment with the parameters of this reordering process (e.g. how soon should a delayed shipment be considered missing?) under different disruption regimens. More work is needed, however, from the point of view of the nodes that receive orders. An order that is unusually high or unusually low should trigger some response or at least a query. What kinds of responses are reasonable and fruitful and what defines "unusual" need to be studied in future work.

7.3 CMOs

Our model called for the use of contract manufacturing organizations (CMOs) at various stages of the production process. We included a number of CMOs as part of its normal operation; additionally, an experiment designer can switch on additional CMOs as a countermeasure in the case of a disruption that halts the operation of some normal production nodes. We did not have time to properly analyze the effect of adding such CMOs. Diversifying the supply base would serve to add resilience to the supply chain, and protect against disruptions to either some of the CMOs or the major components of the focal company's production process.

Another idea we discussed but didn't have time to analyze involved identifying backup suppliers. We began an analysis of a post-disruption strategy of having additional back-up CMOs lined up in advance, with a contract in place, to add to the supply chain in case a disruption is discovered. These would presumably be more expensive and complex, calling for a cost-benefit analysis, which is a direction that calls for future research. It is already a strategy in use by FEMA and other DHS components, and one that could be analyzed effectively with some modifications in MedTech.

An alternative to the post-disruption mitigation of having backup CMOs or backup suppliers, called upon after a disruption is detected, is to add redundancy by increasing the number of CMOs or suppliers

utilized from the beginning, before any disruptions occur. It would be of interest to be able to compare the resilience of medical devices and other supply chains with different numbers of CMOs or suppliers on hand. Our simulation should be readily able to do experiments that vary these numbers, but we were unable to run such experiments due to time constraints.

7.4 Other Mitigations

There are a wide variety of potential pre- and post-disruption mitigations outlined in Section 3.3, and many more that we have not described. Here we mention just one example: increasing threshold for passing tests at selected nodes once an anomaly is found (e.g., in terms of the rules for declaring an expected shipment as “missing”). This mitigation could be tested in an enhanced supply chain simulation that allowed for modifications in testing protocols and corresponding speed of disruption detection/disruption recovery.

Another direction of research on mitigations involves making some facilities more difficult to disrupt by adding an extra cost to disruptions at these points. Our work already analyzed the effect of such strengthening of nodes against disruptions by carrying out element-wise multiplication of the criminal agent resources allocation vector by an additional weight vector before passing the result to the CMA-ES optimizer. But we should also look at this through allocating strengthening to different nodes of the supply chain under a limited budget for strengthening, and then analyzing a resulting optimization problem.

Additionally, optimizing these mitigations, along with others and besides target levels (i.e. backup CMOs and suppliers) is a potential future direction. Optimization across different mitigation strategies can lead to more insight on which combinations of mitigations are most effective at thwarting disruptions.

7.5 Anomaly Detection

Anomaly detection is a key component of disruption detection. The process of anomaly detection leverages the multiplicity of sensors throughout the supply chain. An Anomaly Detection System (ADS) needs to be carefully empirically calibrated and adjusted, keeping in mind known and knowable threats. While in some situations offline ADS can be useful, the online deployment of adaptive ADS seems to offer the most robust mechanism for indicators and early warning. We believe that a built-in, integrated “learning” component in the ADS would be especially valuable, and a worthwhile direction for a future investigation.

Unlike the pharmaceutical supply chain analyzed in Phase 1 of the project, where most production facilities are operational at all times, in the MedTech supply chain they are shut down for long periods of time when there is no demand. For this reason, a sensor, to be able to detect a disruption and trigger

mitigations, needs to measure the daily production, and compare it against what needs to be produced on a given day. To achieve this, we would need sensors at all nodes measuring various metrics, material levels based on their type and other characteristics, and a rule-based system based on required production information available to the FC. One element of this system, already implemented in our model, is the procedure for the recipients of shipments to compare shipments' arrival time with the order records, and thus enabling the detection of delayed or missing shipments. We recommend pursuing the a comprehensive implementation and analysis of this extension as a direction for future work.

7.6 Expanded Criminal Model and Coevolution

Our present criminal model assumes a criminal organization which, given some fixed resources and constraints specified by the modeler, produces a maximally damaging disruption. This is done through optimization. We then can use a follow-on optimization process to find the best mitigation strategy for that particular disruption.

However we expect that in many cases different disruption patterns will require different mitigation strategies. Indeed in an expanded criminal organization model, we may assume that the criminal organization can adapt to mitigation strategies, or anticipate them, and so change its disruption approach to make plausible mitigations ineffective or more difficult. We intend to handle this situation using a different optimization method commonly known as *competitive coevolution*. ECJ supports competitive coevolution among many other coevolution methods.

In competitive coevolution, we use two separate optimization processes running in parallel: one process for the disruption approaches and one for the mitigation strategy. To assess the performance of a mitigation strategy, we select N current disruption approaches from the other optimization process. We then test the mitigation strategy against each of these approaches in a separate simulation. Similarly, to assess the performance of a disruption approach, we test it against N current mitigation strategies in separate simulations.

The idea is that while the mitigation strategies are attempting to optimize against a *variety* of disruption methods, the disruption methods are optimizing to find hidden corners in the space not yet handled by the mitigation strategies with which to challenge them. At the end of the day, we are likely primarily interested in the optimal mitigation strategies, and so could discard the disruption approaches when appropriate.

To assess the performance of a mitigation strategy, we assess against N disruption approaches (and vice versa), and then combine them to make a final assessment. We could, of course, use the *mean* result over the

N results as our assessment: but we could also use the *minimum* (worst) or *maximum* (best) result. In the case of minimum, we are emphasizing the discovery of solutions which are good against a *wide variety* of opponents; in the case of maximum, we are emphasizing the discovery of solutions which specialize against *some* opponent. In our case, it would make sense to use mean or minimum for our mitigation strategies, to produce methods which are robust and general-purpose, while using maximum for our disruption approaches pushes them to hunt for unexpected ways to break existing mitigation strategies.

N may be any value ≥ 1 . With lower values (even as low as $N = 1$), assessment is noisier but we can assess more mitigation strategies and disruption methods. The particular optimal setting may vary from problem to problem. It will be of interest to explore the relative outcomes of using different values of N .

8 Closing Comments

In summary, this phase of the project affirmed the rigor and robustness of our general approach to analyzing supply chains under disruptions, in particular those due to criminal intent. The mapping of the supply chain keeping the focal company's viewpoint in mind, while judiciously aggregating external linkages, provides an effective pathway to analyzing the performance of supply chain under disruptive scenarios. This phase of the project uncovered the importance and significance of synchronization, and the potential for new kinds of vulnerabilities to criminal disruption, when two (or more) components have to be put together to make the product usable. In the case of MedTech both DS and EE components of the supply chain needed to be coordinated and any disruption to even one of these would have significant deterioration in performance. Our experiments leveraged "optimized" worst case scenarios to develop most effective mitigation strategies. Our work showed that multiple disruptions can cause significant disruption even if they may appear to be not as detrimental individually. Finally, this phase of the project underscored the importance of early detection through indicators and warnings and reactive speed capability that both play a crucial role in making backup plans effective. At the conclusion of this phase of our work, the project team feels confident of the generalizability of the developed methodology and its efficacy in leveraging knowledge of subject matter experts to produce tools that can lead to important insights about vulnerabilities and mitigations. The detailed but flexible parametrization of the model that drives the simulation experiments has demonstrated the real-world applicability of our approach which can serve as a template for analyses of other supply chains.

References

- [1] Talha Burki. Global shortage of personal protective equipment. *The Lancet Infectious Diseases*, 20(7):785–786, 2020.
- [2] Jennifer Cohen and Yana van der Meulen Rodgers. Contributing factors to personal protective equipment shortages during the covid-19 pandemic. *Preventive medicine*, 141:106263, 2020.
- [3] Jeremy R Beitler, Aaron M Mittel, Richard Kallet, Robert Kacmarek, Dean Hess, Richard Branson, Murray Olson, Ivan Garcia, Barbara Powell, David S Wang, et al. Ventilator sharing during an acute shortage caused by the covid-19 pandemic. *American journal of respiratory and critical care medicine*, 202(4):600–604, 2020.
- [4] Tom Mahler, Nir Nissim, Erez Shalom, Israel Goldenberg, Guy Hassman, Arnon Makori, Itzik Kochav, Yuval Elovici, and Yuval Shahar. Know your enemy: Characteristics of cyber-attacks on medical imaging devices. *arXiv preprint arXiv:1801.05583*, 2018.
- [5] Mark Goh, Robert De Souza, Allan N Zhang, Wei He, and PS Tan. Supply chain visibility: A decision making perspective. In *2009 4th IEEE Conference on industrial electronics and applications*, pages 2546–2551. IEEE, 2009.
- [6] Mohd Helmi Ali, Norhidayah Suleiman, Norlin Khalid, Kim Hua Tan, Ming-Lang Tseng, and Mukesh Kumar. Supply chain resilience reactive strategies for food smes in coping to covid-19 crisis. *Trends in food science & technology*, 109:94–102, 2021.
- [7] Sean Luke, Robert Simon, Andrew Crooks, Haoliang Wang, Ermo Wei, David Freelan, Carmine Spagnuolo, Vittorio Scarano, Gennaro Cordasco, and Claudio Cioffi-Revilla. The MASON simulation toolkit: Past, present, and future. In *International Workshop on Multi-Agent-Based Simulation (MABS)*, 2018.
- [8] Giuseppe D’Ambrosio and Sean Luke. Hybrid agent-based and discrete event simulation in mason. In *In Society for Simulation and Modeling Annual Modeling and Simulation Conference (ANNSIM)*, 2023.
- [9] Kathleen M Eisenhardt. Building theories from case study research. *Academy of management review*, 14(4):532–550, 1989.
- [10] Wallace J Hopp, Seyed MR Iravani, and Zigeng Liu. Mitigating the impact of disruptions in supply chains. *Supply chain disruptions: Theory and practice of managing risk*, pages 21–49, 2012.

- [11] Martin Christopher and Hau Lee. Mitigating supply chain risk through improved confidence. *International journal of physical distribution & logistics management*, 2004.
- [12] Patricia AH Williams and Andrew J Woodward. Cybersecurity vulnerabilities in medical devices: a complex environment and multifaceted problem. *Medical Devices: Evidence and Research*, pages 305–316, 2015.