

MIT Open Access Articles

Increasing Supply Chain Robustness through Process Flexibility and Inventory

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Simchi-Levi, David, Wang, He and Wei, Yehua. 2018. "Increasing Supply Chain Robustness through Process Flexibility and Inventory." *Production and Operations Management*, 27 (8).

As Published: <http://dx.doi.org/10.1111/poms.12887>

Publisher: Wiley

Persistent URL: <https://hdl.handle.net/1721.1/140771>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Author Manuscript

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the [Version of Record](#). Please cite this article as [doi: 10.1111/poms.12887](https://doi.org/10.1111/poms.12887)

This article is protected by copyright. All rights reserved

Increasing Supply Chain Robustness through Process Flexibility and Inventory

David Simchi-Levi *

Department of Civil and Environmental Engineering, Institute for Data, Systems, and Society, Operations Research Center,
Massachusetts Institute of Technology, 77 Massachusetts Avenue, dslevi@mit.edu

*Corresponding Author

He Wang

School of Industrial and Systems Engineering, Georgia Institute of Technology,
755 Ferst Drive, NW, Atlanta, GA 30332, he.wang@isye.gatech.edu

Yehua Wei

Carroll School of Management, Boston College,
140 Commonwealth Avenue, Chestnut Hill, MA 02467, yehua.wei@bc.edu

We study a hybrid strategy that uses both process flexibility and finished goods inventory for supply chain risk mitigation. The interplay between process flexibility and inventory is modeled as a two-stage robust optimization problem. In the first stage, the firm allocates inventory before disruption happens; in the second stage, after a disruption happens, the firm determines production quantities at each plant to minimize demand loss. Our robust optimization model can be solved efficiently using constraint generation, and under some stylized assumptions, can be solved in closed form. For a canonical family of flexibility designs known as the K -chains, we provide an analytical expression for the optimal inventory solution, which allows us to study the effectiveness of different degrees of flexibilities. Moreover, we find that firms should allocate more inventory to high variability products when its level of flexibility is low, but as flexibility increases, the inventory allocation pattern “flips” and firms should allocate more inventory to low variability products. These observations are further verified through a numerical case study of an automobile supply chain. Finally, we extend our robust optimization model to the *time-to-survive* metric, a metric that computes the longest time a supply chain can guarantee a predetermined service level under disruption.

Key words: supply chain risk, process flexibility, inventory, robust optimization

History: received: March 2017; accepted: April 2018 by Qi Annabelle Feng, after 2 revisions

1. Introduction

Over the past decade, managing supply chain disruption risks has emerged as one of the top business challenges. As reported by Forbes, 80 percent of companies worldwide see better protection of

supply chains as a priority (Culp 2013). Supply chain disruption is difficult to manage because it can come from a wide range of sources, whether from natural disaster, epidemics, factory fire or political upheavals (Simchi-Levi et al. 2014). The problem is further exacerbated as new sources of disruptions are constantly emerging, exemplified by the recent Equifax data breach (O'Marah 2017) and the computer system outage at Delta Air Lines (Jansen 2017).

To effectively manage supply chain risks, a firm needs to coordinate different risk mitigation strategies under supply and demand uncertainty. In this study, we focus on two popular risk mitigation strategies: holding finished goods inventory and employing process flexibility. While it is clear that both strategies improve supply chain resilience, it is much more challenging to understand the hybrid approach that combines both strategies. Therefore, for a firm to successfully implement the hybrid approach, one needs to answer the following questions: how does the firm's flexibility design affect its inventory decisions? And how should the firm allocate inventory among different products to achieve a required service level while minimizing cost?

To study the hybrid strategy, we introduce a two-stage robust optimization model in which finished goods inventory is stored in the first stage, while in the second stage, after disruption happens, the firm decides what and how much to produce at each plant using process flexibility to minimize the impact of disruption. Our two-stage risk mitigation model is intended for relatively infrequent disruption events, in contrast to stochastic dynamic models that are used to model frequent disruptions, e.g., (Tomlin 2006). In our model, the second stage decisions represent the firm's actions during a disruption, and the first stage inventory decisions represent the inventory levels the firm wants to build up prior to disruptions. Moreover, the inventory flexibility strategy we consider is most practical when a firm is facing disruptions that last for a short term, e.g., a few weeks, as holding months of inventory is a very expensive risk mitigation strategy. For longer term disruptions, the firms may consider our inventory flexibility strategy as an intermediate solution, while remodeling their supply chains to resume production quickly. For example, after the 2011 Japanese Earthquake, Toyota started to create a quake-proof supply chain that is guaranteed to recover in less than two weeks when another disruption occurs (Kim 2011).

1.1. Risk Mitigation Strategies

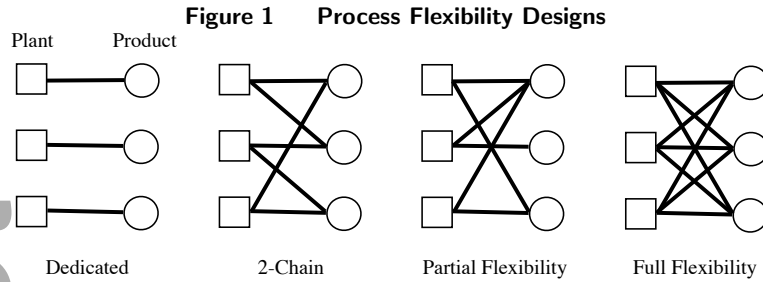
Next, we provide basic intuitions behind the two risk mitigation strategies, process flexibility and inventory, and then explain why it is important to coordinate the two strategies.

Risk mitigation inventory, also known as protective inventory, has been identified in a number of papers as an important tool for dealing with supply chain risk (see literature review in section 1.4). However, holding a large amount of inventory would incur a high loss in cost efficiency (Chopra and Sodhi 2014), as these inventories are only used when disruptions occur. As a result, firms would typically hold the minimum amount of inventory subject to a given performance requirement.

Process flexibility is defined as the “ability to build different types of products in the same manufacturing plant or on the same production line at the same time” (Jordan and Graves 1995). For example, under *full flexibility* design, each plant is capable of producing all products; while under *dedicated* design (i.e., no process flexibility), each plant is capable of producing just a single product (see Figure 1).

With process flexibility, the firm is capable of adjusting its production in the existence of uncertainties, and thus, is in a better position to match its capacity and demand. Unfortunately, implementing full flexibility is often too expensive since each plant needs to be capable of producing all products (Simchi-Levi 2010). As a result, *partial* flexibility designs are considered. A popular class of partial flexibility designs studied in the literature are known as *K-chains* (Hopp et al. 2004, Chou et al. 2014). In a *K-chain*, there are N plants and N products, and for each $1 \leq i \leq N$, plant i has the flexibility to produce products $i, i + 1, \dots, i + K - 1$ modulo N (see Figure 1 for an example of 2-chain). Note that if $K = 1$, we get dedicated design in the N plants and N products system, and if K is equal to the number of products, we get full flexibility design.

It is important to note that process flexibility can greatly impact inventory decisions. Consider dedicated and full flexibility designs in Figure 1 as an example. Suppose that the firm holds inventory for Product 2, and a disruption occurs at Plant 1. Then, under dedicated design, the inventory of Product 2 would not be helpful, since the inventory cannot be used to satisfy any



of the demand for Product 1. However, under full flexibility design, inventory of Product 2 can be used to save the flexible production capacities at Plant 2, which allows the firm to reallocate the capacities at Plant 2 to produce Product 1. In other words, when a firm has both inventory and process flexibility, inventory helps to free up excess flexible capacities during an unforeseen event, and thus improves the firm's ability to mitigate risk. As a result, the existence of process flexibility can reduce the amount of inventory needed, and the synergy between inventory and process flexibility makes the hybrid strategy very compelling.

Although it is well documented that both flexibility and inventory can improve supply chain resilience (see literature review in section 1.4), the synergy between inventory and process flexibility is not well understood. To the best of our knowledge, no existing literature has considered a hybrid strategy combining (partial) process flexibility designs and inventory as a way to mitigate against risks in a multi-product supply chain. We note that because holding extra amounts of inventories can be inefficient and expensive, our hybrid strategy is mainly intended for infrequent disruptions when the firm can restore its original production capacity in a few weeks. In the latter part of the paper, we will also discuss how *process flexibility* may greatly increase the effectiveness of inventories, and potentially allow firms to use this strategy for longer disruptions.

1.2. Robust Optimization Model

In both process flexibility and risk mitigation literature, uncertainties in demands and capacities are typically specified through probability distributions. Given probability distributions, the firm's objective is modeled using two different approaches. In the first approach, the firm's objective is to maintain a high level of profit/production with high probability (aka Type 1 service level). In

this case, the inventory and process flexibility problem would be modeled as a two-stage stochastic program with chance constraints. Unfortunately, these service constraints are non-convex (see Online Appendix for more details), which makes the stochastic problem intractable. A different approach is to minimize the firm's expected cost. This would lead to a two-stage linear stochastic program, which is a convex optimization problem. However, even for discrete distributions, solving the two-stage linear stochastic program exactly is $\#P$ -hard (Dyer and Stougie 2006), and practitioners would often resort to approximating the true problem using sample average approximation (SAA). The quality of the SAA solutions depends on the number of samples, and the number of samples required for SAA solutions to achieve strong theoretical guarantees is often very large (Shapiro and Nemirovski 2005). Therefore, SAA is only capable of producing good solutions for problems with a small number of variables.

In the interest of computational tractability, we propose a two-stage robust optimization model that finds the optimal inventory allocation under general process flexibility designs. In the robust optimization model, product demands and plant capacities lie in an uncertainty set. In the first stage, given a process flexibility design, the firm allocates its inventory levels across the products; and in the second stage, after the uncertainties are realized, the firm uses both inventory and process flexibility to either minimize its cost or satisfy a certain service level requirement. We find this robust model to be more tractable than its stochastic counterparts. In our numerical example (section 4), we find that two-stage robust problems for supply chains can be solved much faster than its stochastic counterparts. In some stylized settings, our robust model can even be solved in closed form (section 3), allowing for detailed analysis in order to identify useful insights. Finally, further computational experiments demonstrate that our robust model containing hundreds of nodes can be solved to optimality in less than 10 minutes (see chapter 4 of Wang 2016).

Apart from tractability, the worst-case model provides several additional advantages compared with its stochastic counterparts in the context of risk mitigation. First, it is difficult and often impossible to accurately assess the uncertainties in plant disruptions, because production disruption

can occur from so many different sources, whether from natural disaster, epidemics, or factory fire (Simchi-Levi et al. 2014). Because the optimal inventory level can be very sensitive to the probability of disruption, considering the worst-case scenario may offer a more “robust” approach. Second, for small probability events such as plant disruptions, managers might find it useful to understand the maximum possible shortage of demand under a wide range of scenarios. This understanding may lead them to better identify scenarios where the supply chain is most vulnerable. Finally, it has been suggested in Graves and Willems (2000) that it is easier for managers to communicate with customers by committing to a certain service level under a range of scenarios, rather than providing a probabilistic guarantee, which is difficult to understand or verify.

1.3. Overview and Summary of Results

Robust Optimization Model. In section 2, we propose a robust optimization formulation for supply chain risk mitigation. The decision variables are inventory levels in the first stage and production quantities in the second stage. The formulation ensures that the total demand lost is bounded by some target value for all scenarios in the uncertainty set, with the objective of minimizing the inventory costs. It turns out that our formulation is highly tractable numerically, and under some stylized assumptions, can be even solved in closed form. We also demonstrate that our model can be extended in several ways in section 5. In particular, our model can be extended to find the optimal inventory allocation that maximizes a supply chain’s *time-to-survive* (TTS), a metric that measures the longest time the supply chain can maintain customer service levels.

Analysis and Insights. In section 3, we analyze a special family of flexibility designs called the K -chains to provide insights into inventory allocations under different degrees of flexibility. First, we provide a closed-form characterization of the optimal inventory decision when the uncertainty set is symmetric. Using this characterization, we find that while changing from a dedicated network to the well-studied 2-chain design provides a large portion of benefit, when plants are subject to disruptions, 2-chain does not completely capture the benefit of full flexibility even under moderate demand variabilities, and there are significant benefits achieved by increasing the degree of

flexibility beyond 2-chain. The characterization also allows us to understand how the gap between inventory costs of 2-chain and K -chain ($K > 2$) changes when we vary the magnitude of potential disruptions.

Second, we show that the optimal inventory allocation can be drastically different under different flexibility designs. In particular, there tends to be a shift of inventories from high variability products to low variability products as the degree of process flexibility increases. We refer to such phenomenon as the *flipping effect*, since the inventory levels flip as the degree of flexibility increases. Therefore, it is unlikely that a single inventory strategy can perform well for all kinds of flexibility designs. Optimizing over inventory decisions that are tailored to the firm's particular *supply chain network structure* using our robust optimization model may significantly improve the firm's risk mitigation strategy, compared with other simple inventory heuristics that ignore the firm's supply chain structure.

Numerical Case Study. In section 4, we apply our robust optimization approach to a risk mitigation example of General Motors. The example is based on a real-world case introduced by Jordan and Graves (1995). The numerical experiment supports our finding that 2-chain is not as effective as K -chain for $K > 2$ when there is supply uncertainty, even when the assumptions in section 3 do not hold. We also demonstrate the flipping effect in this numerical case study.

1.4. Related Literature

There is rich literature on inventory for risk mitigation. Many of the earlier papers (e.g., Meyer et al. 1979, Song and Zipkin 1996, Arreola-Risa and DeCroix 1998) studied inventory risk mitigation in a single product setting. More recently, inventory mitigation strategies under multi-period, multi-echelon settings have also been explored (e.g., Bollapragada et al. 2004, DeCroix 2013). Moreover, researchers have also considered hybrid strategies in which firms use both dual sourcing and inventory to mitigate risks (Gürler and Parlar 1997, Tomlin 2006).

Process flexibility, also referred to as “mix flexibility” or “product flexibility”, has also been observed as a potential risk mitigation tool. Tomlin and Wang (2005) consider a risk mitigation

strategy that uses a combination of mix-flexibility and dual sourcing. Tang and Tomlin (2008) suggest process flexibility as one of the five types of flexibility strategies that can be used to mitigate supply chain disruptions. And finally, Sodhi and Tang (2012) list flexible manufacturing processes as one of the eleven robust supply chain strategies.

While many researchers investigated the effectiveness of fully flexible resources (e.g., Fine and Freund 1990, Bish and Wang 2004, Tomlin and Wang 2005), our model allows one to study arbitrary partial process flexibility designs in a multi-plant, multi-product system. This feature is motivated by the seminal work of Jordan and Graves (1995), who argue that fully flexible resources are often too expensive or impossible to implement, while a little bit of flexibility in the system often provides the same benefit as full flexibility. More recently, the effectiveness of sparse flexibilities designs have been verified by a number of theoretical developments (e.g., Chou et al. 2010, 2011, Simchi-Levi and Wei 2012, Wang and Zhang 2015). Moreover, Bassamboo et al. (2010, 2012) observed that when the firm is free to purchase any flexible resource, sparse configurations such as chaining and tailored pairing are often near optimal.

Production postponement and delayed product differentiation is another stream of literature closely related to the two-stage flexibility and inventory model we study. The classical research of Fisher and Raman (1996) and Swaminathan and Tayur (1998) studied two-stage production models, but neither paper proposed a model that incorporates partial second-stage production flexibility. To the best of our knowledge, the only paper in the postponement literature that considers partial flexibility is the work of Chou et al. (2014). However, that work differs from our paper significantly as they focused on a stochastic model with no capacity uncertainties, whereas our paper studies a robust model with capacity disruptions.

In this paper, we propose a robust optimization approach to study the risk mitigation strategy for combining partial process flexibility and inventory. The model is a two-stage robust optimization problem, where the firm allocates inventory before the uncertainties are realized (*ex-ante*), and schedules its flexible production after uncertainties are realized (*ex-post*). Our approach to

model uncertainties follows the recent proposals in the robust optimization literature (e.g., Ben-Tal et al. 2009, Bandi and Bertsimas 2012), which propose to replace probabilistic distributions with uncertainty sets. This approach is similar to the model studied by Graves and Willems (2000), who argued that in supply chain management, it is often easier for the firm to commit to service level guarantees under a range of scenarios, rather than probabilistic service guarantees.

In the two-stage robust optimization literature, Ben-Tal et al. (2004) studied the general two-stage robust linear programs (which they called “Adjustable Robust Counterpart”). They proved that the two-stage robust linear programs are NP-hard in general, and proposed the Affine Adjustable Robust Counterpart, which is shown to be computationally tractable in many settings. Atamtürk and Zhang (2007) studied a general two-stage robust network flow problem, showing that while the problem is NP-hard in general, it can yield polynomial-timed solutions when the networks have certain structures. Computational methods for solving general two-stage robust linear programs include, for example, Thiele et al. (2010), who proposed a cutting-plane method, and Zeng and Zhao (2013), who proposed a column-and-constraint generation algorithm. For a more detailed overview of the recent advances and applications of two-stage robust optimization, we refer the readers to the recent surveys of Bertsimas et al. (2011) and Gabrel et al. (2014).

2. The Model

We consider a firm that has M plants (denoted by \mathcal{S}_i for $i = 1, \dots, M$) and produces N products (denoted by \mathcal{T}_j for $j = 1, \dots, N$). A plant may produce either one or multiple products, and the *flexibility design* specifies which product each plant can produce. More formally, the firm’s flexibility design can be represented as a bipartite network, where a link $(\mathcal{S}_i, \mathcal{T}_j)$ exists if and only if plant i is able to produce product j . We refer to the set of such links as \mathcal{F} . For a given flexibility design \mathcal{F} and a subset of products $A \subset \{\mathcal{T}_1, \dots, \mathcal{T}_N\}$, we define $P_{\mathcal{F}}(A) = \{\mathcal{S}_i : (\mathcal{S}_i, \mathcal{T}_j) \in \mathcal{F}, \mathcal{T}_j \in A\}$ as the plants that can produce at least one of the products in A .

We assume that both plant capacity and product demand are uncertain quantities due to unknown risks such as plant disruptions or demand fluctuations. We denote the realized capacity

by $\mathbf{c} = (c_1, c_2, \dots, c_M)$, and the realized demand by $\mathbf{d} = (d_1, d_2, \dots, d_N)$. We use $\mathcal{U} \subset \mathbb{R}_+^M \times \mathbb{R}_+^N$ to denote the uncertainty set for both capacity and demand, and assume $(\mathbf{c}, \mathbf{d}) \in \mathcal{U}$. Throughout the paper, we will consider polyhedral uncertainty sets; that is, \mathcal{U} will be defined by a system of linear inequalities.

2.1. Service Guarantee Model

Next, we present the *service guarantee* (SG) model, where the firm ensures that total demand loss for any scenario in uncertainty set \mathcal{U} is bounded by a fixed shortage allowance, δ , while minimizing inventory cost. This can be modeled using a two-stage decision problem. In the first stage, the firm determines inventory level for each product, where the inventory level of product j is denoted by variable s_j , with unit holding cost h_j . In the second stage, the firm observes realized product demand and available plant capacities, and determines its production schedule. We use $x_{ij}(\mathbf{c}, \mathbf{d})$ to denote the quantity of product j produced by plant i under demand \mathbf{d} and capacity \mathbf{c} . Similarly, we use $l_j(\mathbf{c}, \mathbf{d})$ to denote the lost sales of product j . The inventory decision problem can then be formulated as follows.

$$\text{(SG)} \quad \min_{\mathbf{s}, \mathbf{x}, \mathbf{l}} \quad \sum_{j=1}^N h_j s_j \quad (1)$$

$$\text{s.t.} \quad \sum_{i: (S_i, T_j) \in \mathcal{F}} x_{ij}(\mathbf{c}, \mathbf{d}) + l_j(\mathbf{c}, \mathbf{d}) \geq d_j - s_j, \quad \forall 1 \leq j \leq N, (\mathbf{c}, \mathbf{d}) \in \mathcal{U},$$

$$\sum_{j: (S_i, T_j) \in \mathcal{F}} x_{ij}(\mathbf{c}, \mathbf{d}) \leq c_i, \quad \forall 1 \leq i \leq M, (\mathbf{c}, \mathbf{d}) \in \mathcal{U},$$

$$\sum_{j=1}^N l_j(\mathbf{c}, \mathbf{d}) \leq \delta, \quad \forall (\mathbf{c}, \mathbf{d}) \in \mathcal{U},$$

$$\mathbf{x}(\mathbf{c}, \mathbf{d}), \mathbf{l}(\mathbf{c}, \mathbf{d}), \mathbf{s} \geq 0. \quad (2)$$

The first constraint defines lost sales where demand is satisfied by either production or inventory, and the second constraint enforces plant capacity limits. The third constraint ensures that the total lost sales is bounded by a fixed budget δ .

While the SG model defines service guarantee aggregated over all products, the model can also be used if the firm requires service guarantees on an individual product level. To do this, we need

to make two changes in the model: first, the lost sales allowance δ is set to zero; second, the uncertainty set \mathcal{U} is modified to reflect service guarantee for individual products. For example, if the firm wants to commit to 95% of the demand for each product under all possible scenarios in uncertainty set \mathcal{U} , then we can define the uncertainty set to \mathcal{U}' , where $\mathcal{U}' = \{(\mathbf{c}, 0.95\mathbf{d}) \mid (\mathbf{c}, \mathbf{d}) \in \mathcal{U}\}$. Note that the uncertainty set \mathcal{U}' can easily be modified to allow different products having different service guarantee levels (by making an affine transformation on \mathcal{U}).

2.2. Optimization Algorithm

We claim that the service guarantee model (SG) can be solved efficiently using an optimization method called constraint generation. Because this paper is focused on deriving operational insights on the interplay between flexibility and inventory, we omit the technical details of the algorithm to solve the robust optimization model; instead, we only give an outline of the algorithm. Full details of the constraint generation algorithm can be found in chapter 4 of Wang (2016).

The algorithm for solving the SG model can be divided into two main steps. First, we can reformulate the two-stage robust optimization problem associated with the SG model as the following linear program:

$$\begin{aligned}
 \min \quad & \sum_{j=1}^N h_j s_j \\
 \text{s.t.} \quad & \sum_{\mathcal{T}_j \in A} (d_j^A - s_j) - \sum_{\mathcal{S}_i \in P_{\mathcal{F}}(A)} c_i^A \leq \delta, \quad \forall A \subset \{\mathcal{T}_1, \dots, \mathcal{T}_N\}, \\
 & \mathbf{s} \geq 0.
 \end{aligned} \tag{3}$$

The derivation of this reformulation can be found in chapter 4 of Wang (2016). Note that this linear program has N variables and 2^N linear constraints (with one constraint associated with each subset of the N products; d_j^A and c_i^A are constants defined for each subset A), so the number of constraints grows exponentially as the number of products increases. When N is not a large number, we can directly solve the linear programming formulation (3). However, in general when $N > 15$, solving (3) with all the 2^N constraints becomes extremely slow.

To solve the SG model faster, in the second step, we apply a constraint generation algorithm to the linear programming problem (3). We provide the basic ideas of our algorithm here. The constraint generation algorithm for solving the SG model runs iteratively, starting from a linear program with the same objective function as in (3) but with an empty set of constraints. At each iteration, the algorithm first solves the linear program with objective (3) and the current set of constraints, then runs a subroutine that either proves that the current inventory vector \mathbf{s} is optimal, or returns a violating linear constraint in (3). Depending on the outcome of the subroutine, the algorithm either terminates with an optimal solution, or adds another inequality to the current set of constraints. When the uncertainty set is polyhedral, the subroutine can be reformulated as a compact mixed integer linear program (MILP) (see chapter 4 of Wang (2016) for details), which can be solved quickly using off-the-shelf optimization solvers. In general, the algorithm stops when only a tiny fraction of the 2^N constraints are added to the set, so the running time of the constraint generation algorithm (and the running time of the subroutines) is much shorter than the time to solve the full linear programming formulation with exponentially many constraints. For example, when the number of products is equal to 50 and the uncertainty set is polyhedral, the full linear programming formulation has $2^{50} \approx 1.12 \times 10^{15}$ constraints, and solving a linear program with this size is beyond the capability of any existing LP solvers. In contrast, we observe in numerical experiment that our constraint generation algorithm requires only (on average) a few hundreds constraints being generated, even if the LP contains more than 10^{15} constraints; furthermore, each constraint generation step takes a fraction of a second. For a detailed discussion of the computational complexity of the algorithm, we refer readers to chapter 4 of Wang (2016).

We note that while constraint generation is a common standard technique for solving linear programs with a large number of constraints (Bertsimas and Tsitsiklis 1997), different tricks are often needed to come up with the subroutine that efficiently generates constraints depending on specific problem structures. In our constraint generation algorithm, we take advantage of the network flow structure to come up with a subroutine that solves a mixed-linear integer program with moderate

size. Finally, while we focus on the SG model in the paper, our computational algorithm can be applied to more general models with different lost sales costs for different products, or multiple stages of productions. These extensions are discussed in section 5.

2.3. Constructing Uncertainty Sets for Stochastic Fluctuations

So far, we have introduced a robust optimization formulation of the supply risk mitigation problem, and our model framework allows using any polyhedral uncertainty set (i.e., a finite set of linear inequalities) to model demand and supply uncertainty. Typically, in robust optimization, the uncertainty set is chosen based on some *probabilistic interpretation* (Bertsimas and Thiele 2006). In this section, we introduce a specific class of parametric uncertainty sets and discuss how to choose the values of the uncertainty set parameters in applications. For the rest of the paper, we will mainly focus on this class of uncertainty sets.

We start with the demand uncertainty. In the literature, product demands are often described as stochastic quantities. Suppose that the demand for each product j is denoted by some random variable D_j with mean \bar{d}_j and standard deviation σ_j . In this case, applying the method described by Ben-Tal et al. (2009), we use the following set of linear inequalities to restrict the variations of demand, \mathbf{d} .

$$\sum_{j=1}^N (d_j - \bar{d}_j) \leq \gamma, \sum_{j=1}^N \frac{|d_j - \bar{d}_j|}{\sigma_j} \leq \beta, |d_j - \bar{d}_j| \leq \alpha \sigma_j, \forall 1 \leq j \leq N. \quad (4)$$

The constraint defined by parameter γ allows us to upper-bound the maximum total demand increase compared with the mean. The constraint defined by parameter β is known as the budgeted uncertainty constraint, and allows us to control the total normalized deviation between the actual demand and its mean. Finally, the constraints defined by parameter α is known as box uncertainty constraints and allow us to control demand deviations for each individual product. The values of α , β and γ can be chosen based on empirical data or probability theories such as the central limit theorem (Ben-Tal et al. 2009). Alternatively, the parameters can be chosen in order to optimize or estimate stochastic quantities (Bertsimas and Thiele 2006, Bandi and Bertsimas 2012). To avoid negative product demand, we assume throughout the paper that $\alpha \sigma_j \leq \bar{d}_j$ for any $1 \leq j \leq N$.

The uncertainty set (4) includes several types of commonly used uncertainty sets as special cases. For instance, one may consider a *box uncertainty set* where each demand variable d_j is bounded individually, which is a special case of (4) with $\beta = \gamma = +\infty$. The *budget uncertainty set* is also a special case of (4) with $\beta = +\infty$.

Although we will focus on demand uncertainty sets defined by (4) in the rest of the paper, we note that other forms for demand uncertainty set can be used in our model. A possibility is bounding the standard deviation of demands instead of the absolute deviation of demands as in (4). However, this modification involves quadratic constraints in the uncertainty set. When the uncertainty set contains only linear inequalities as in (4), the optimization algorithm (described in §2.2) can be implemented efficiently by solving a sequence of mixed integer linear programs (MILPs). If the uncertainty set is a quadratic conic set, our algorithm requires solving a sequence of mixed integer conic quadratic programs (MICQPs). Although there are several open-source solvers for MICQPs, the dimensions of solvable MICQP by current solvers are much smaller compared to MILPs (Bonami et al. 2012).

In the context of risk mitigation, most of the capacity uncertainty comes from plant disruptions due to low probability, unforeseeable events. Because standard deviations of the disruptions are difficult to estimate and plant disruptions may cause the entire plant to go down, we do not consider the budgeted uncertainty constraint to the capacity vector. Moreover, because plant capacities can only decrease due to disruption, we use the following set of inequalities to restrict the capacity vector, \mathbf{c} ,

$$\sum_{i=1}^M (\bar{c}_i - c_i) \leq \zeta, \quad 0 \leq c_i \leq \bar{c}_i, \quad \forall 1 \leq i \leq M, \quad (5)$$

where parameter ζ bounds the total loss of capacity for all of the plants. Alternatively, we can view ζ as the firm's judgment on the potential magnitude of a disruption.

In the rest of this paper, we will focus on the uncertainty set \mathcal{U} defined by Equations (4) and (5). For more techniques to construct uncertainty sets, we refer interested readers to Bandi and Bertsimas (2012) and the references therein.

3. Analysis and Insights

In the previous section, we proposed a robust optimization framework that coordinates inventory and production decisions to mitigate supply disruptions. To further understand this robust optimization approach, we study two stylized settings based on the service guarantee (SG) model. In both settings, we assume that the number of plants and products are equal ($M = N$), all products have a unit holding cost ($h_j = 1$), and the demand shortage parameter is zero ($\delta = 0$). Note that δ bounds the aggregate demand loss. As we commented in section 2.1, even when $\delta = 0$, demand loss for individual products can still be included in the model by making an affine transformation to the uncertainty set.

3.1. Inventory Costs Under Different Flexibility Levels

To analyze different flexibility levels, we consider the K -chain designs. The K -chains are commonly studied in the literature (e.g. Hopp et al. 2004, Chou et al. 2014, Wang and Zhang 2015) as a class of designs with different levels of flexibility. Recall that in a K -chain, for each $1 \leq i \leq N$, plant i has the flexibility to produce products $i, i + 1, \dots, i + K - 1$ modulo N . If $K = 1$, we get dedicated design in the N by N system, and if K is equal to the number of products, we get full flexibility design.

We use the class of uncertainty sets defined by Equations (4) and (5) in section 2.3, and assume that the uncertainty is symmetric, i.e., $\bar{d}_j = \bar{d}$, $\sigma_j = \sigma$ and $c_i = \bar{c}$, for all $1 \leq j \leq N$. Under our assumption, any \mathcal{U} can be expressed as $\mathcal{U} = \mathcal{U}_c \times \mathcal{U}_d$, where

$$\mathcal{U}_d = \left\{ \mathbf{d} \mid \sum_{j=1}^N (d_j - \bar{d}) \leq \gamma, \sum_{j=1}^N |d_j - \bar{d}| / \sigma \leq \beta, |d_j - \bar{d}| / \sigma \leq \alpha, \forall 1 \leq j \leq N \right\}, \quad (6)$$

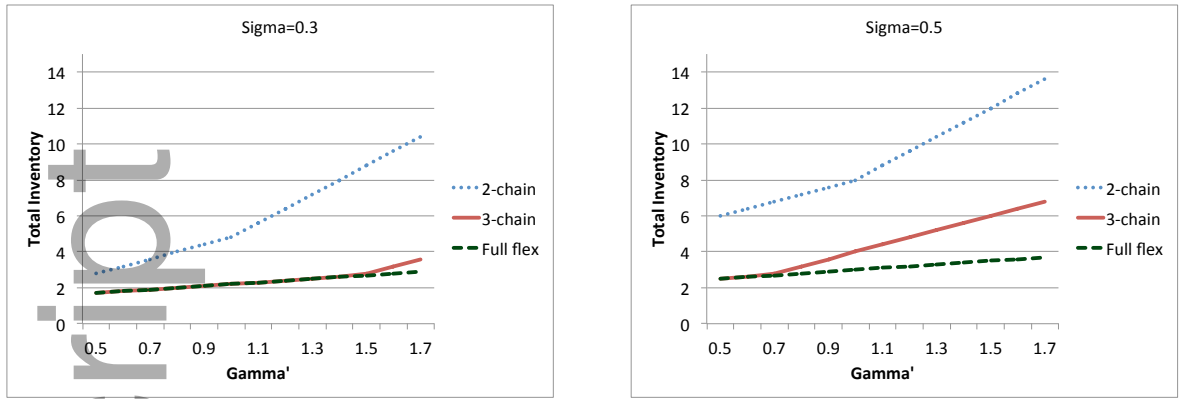
$$\mathcal{U}_c = \left\{ \mathbf{c} \mid \sum_{i=1}^N (\bar{c} - c_i) \leq \zeta, 0 \leq c_i \leq \bar{c}, \forall 1 \leq i \leq N \right\}, \quad (7)$$

for some parameters γ, ζ, β and α .

Under these assumptions, we can obtain a closed-form analytical expression for the optimal inventory of the SG model for K -chains. Let \mathbf{e} be a vector with all entries equal to 1. For any integer $1 \leq t \leq N$, we define

$$C^{\min}(t) = \min_x \{ \mathbf{c}^T \mathbf{x} \mid \mathbf{e}^T \mathbf{x} \geq t, 0 \leq \mathbf{x} \leq 1, \mathbf{c} \in \mathcal{U}_c \},$$

Figure 2 Inventory Cost for K-chain Designs



$$(N = 8, \bar{d} = \bar{c} = 1, \alpha = 2, \beta = \gamma = 4)$$

$$D^{\max}(t) = \max_x \{ \mathbf{d}^T \mathbf{x} \mid \mathbf{e}^T \mathbf{x} \leq t, 0 \leq \mathbf{x} \leq 1, \mathbf{d} \in \mathcal{U}_d \}.$$

PROPOSITION 1. Let \mathcal{F} be a K -chain, for some integer $1 \leq K \leq N$. Then the optimal inventory cost of the SG model is

$$N \max \left\{ \max_{1 \leq t \leq N-K} \frac{D^{\max}(t) - C^{\min}(t + K - 1)}{t}, \max_{N-K < t \leq N} \frac{D^{\max}(t) - C^{\min}(N)}{t}, 0 \right\}. \quad (8)$$

The characterization for the optimal inventory cost of K -chain in Proposition 1 provides us with a useful tool to think about the effectiveness of different degrees of flexibility. For example, it allows us to solve the optimal inventory cost of K -chains by solving N linear programs, which is much faster than the constraint generation algorithm for the general SG model. In addition, the characterization allows us to further analyze the optimal cost for a K -chain as we change different parameters. In Figure 2, we plot an example of the inventory cost as magnitude of a disruption (measured by ζ) changes, under different levels of demand variability σ and different flexibility designs including 2-chain, 3-chain and full flexibility. The reason we consider the impact of change in ζ is because it is typically difficult to estimate, as disruption may come from many different sources. From another perspective, one can view ζ as a firm's conservativeness thinking about disruption - larger ζ would imply that the firm is planning more conservatively as it prepares for disruptions impacting a larger portion of its production capacity.

From Figure 2, we observe that there is a significant gap between the performance of 2-chain and K -chain for $K > 2$ with both capacity and demand uncertainty. This contrasts to the classical observations that 2-chain performs almost as well as full flexibility (and hence K -chain for any $K > 2$) when there is no capacity uncertainty (Jordan and Graves 1995). For illustration, in the 10 plants 10 products example studied by Jordan and Graves (1995), if demand is normally distributed with coefficient of variation of 0.4, then the expected utilization rate of full flexibility is 95.4%, while that of 2-chain is already 95%. In comparison, the optimal inventory of 2-chain in Figure 2 is more than twice that of full flexibility, under most of the parameter ranges.

Interestingly, we also observe that most of the gap between 2-chain and full flexibility can be closed by 3-chain, which is just one additional degree of flexibility higher than 2-chain. In hindsight, this is not entirely surprising. Intuitively, while 2-chain is very effective in satisfying uncertain demand without capacity uncertainty, even the disruption of a single production plant can break up the chain; in comparison, 3-chain would still form a chain even if one plant is down. We note that recent findings suggest that 2-chain may not be effective when the demand variability is large compared with capacity (Chou et al. 2014, Wang and Zhang 2015). A crucial difference between our observation to the existing literature is that in our setting, a relatively small amount of uncertain supply disruption, e.g. shutdown of a single plant, creates a significant gap between 2-chain and full flexibility, as illustrated in Figure 2 under moderate demand variability ($\sigma = 0.3$).

Given that there is a significant gap between 2-chain and K -chain for $K > 2$ with both capacity and demand uncertainty, another interesting question is the comparative static of this gap and potential capacity loss. From Figure 2, we observe that when capacity loss (measured by ζ) increases, the gap between the inventory costs for 2-chain and K -chain ($K > 2$) widens. Next, we formally prove this observation under a wide range of ζ .

PROPOSITION 2. *When $0 \leq \zeta \leq 2\bar{c}$, the difference between the optimal inventory costs of 2-chain and K -chain ($K > 2$) is increasing as ζ increases.*

Proposition 2 is useful for a decision maker (DM) thinking about adding flexibility in the following way: suppose that the DM is considering between 2-chain and 3-chain but is unsure about the potential capacity loss caused by disruptions, which is described by ζ . Then, the DM can always calculate the difference between the inventory costs of 2-chain and 3-chain under the minimum (maximum) reasonable ζ , and this quantity would indeed be a lower-bound (upper-bound) on the difference under other values of ζ , as long as $0 \leq \zeta \leq 2\bar{c}$. Finally, in the following corollary, we show that if $\zeta > 2\bar{c}$, then the optimal inventory costs of 2-chain is the same as dedicated design. Thus, when $\zeta > 2\bar{c}$, the firm should ignore 2-chain completely, and consider either implementing flexibility with more than 2 degrees, or not implementing flexibility at all.

COROLLARY 1. *When $\zeta > 2\bar{c}$, the optimal inventory costs of 2-chain is equal to $ND^{\max}(1)$, which is equal to the optimal inventory costs of a dedicated design.*

The intuition behind Corollary 1 is reasonably straightforward. Consider product 1. Intuitively, in the worst-case scenario for product 1, both plants capable of producing product 1 are disrupted, while product 1 has demand $D^{\max}(1)$. This forces us to allocate $D^{\max}(1)$ amount of inventory at product 1. By symmetry of 2-chain, we also need to allocate $D^{\max}(1)$ amount of inventory at product 2, 3, ..., N , which is exactly the optimal inventory allocation for the dedicated design.

3.2. Inventory Allocation Strategy

In this section, we study a setting where product demand variability is different. Like section 3.1, we consider a special case of the uncertainty set defined in section 2.3 given by

$$\mathcal{U}_d = \left\{ \sum_{j=1}^N (d_j - \bar{d}) \leq \gamma, \sum_{j=1}^N |d_j - \bar{d}|/\sigma_j \leq \beta, |d_j - \bar{d}|/\sigma_j \leq \alpha, \forall 1 \leq j \leq N \right\}, \quad (9)$$

$$\mathcal{U}_c = \left\{ \sum_{i=1}^N (\bar{c} - c_i) \leq \zeta, 0 \leq c_i \leq \bar{c}, \forall 1 \leq i \leq N \right\}, \quad (10)$$

with $\mathcal{U} = \mathcal{U}_c \times \mathcal{U}_d$. The parameter σ_j measures the demand variability of product j .

We want to understand how inventory allocations change as we add more flexibility to the system, starting from a dedicated flexibility design. First, for dedicated design ($K = 1$), we show that under the optimal decision, the firm places more inventories to the high variability products.

PROPOSITION 3. When \mathcal{F} is a dedicated design (i.e., 1-chain), there always exists an optimal inventory solution \mathbf{s} , such that for any two products j_1, j_2 with $\sigma_{j_1} \geq \sigma_{j_2}$, we have $s_{j_1} \geq s_{j_2}$.

In contrast, for full flexibility design ($K = N$), it can be shown that the optimal inventory allocation places more inventories to the low variability products, under a large range of parameters.

PROPOSITION 4. When \mathcal{F} is a full flexibility design (i.e., N -chain), if $\bar{c}N \geq \alpha \sum_{j=1}^N \sigma_j + \gamma + \zeta$, there always exists an optimal inventory solution \mathbf{s} , such that for any two products j_1, j_2 with $\sigma_{j_1} \geq \sigma_{j_2}$, we have $s_{j_1} \leq s_{j_2}$.

We note that condition $\bar{c}N \geq \alpha \sum_{j=1}^N \sigma_j + \gamma + \zeta$ is usually satisfied in practice: ζ (the loss of capacity due to disruption), $\alpha \sum_{j=1}^N \sigma_j$ (total absolute demand deviation), and γ (total demand variation) all measure the *variability* of the system; due to the central limit theorem, all three quantities are usually smaller than $\bar{c}N$, the total capacity of the system, because the total capacity scales linearly in N and variability scales in the order of \sqrt{N} . Moreover, a condition of this form is necessary—for example, if $\bar{c}N < \zeta$, then the uncertainty set would contain scenarios where the total available capacity is zero, rendering any flexibility design useless.

Propositions 3 and 4 demonstrate that if we start from dedicated design and add flexibilities until we arrive at full flexibility, the inventory allocations to high and low variability products will be eventually *flipped*. We call this the *flipping effect*, which is the phenomenon that the optimal solution holds more inventory for products with high demand variability when the degree of flexibility is low, but holds more inventory for products with lower variability when the firm has a high degree of flexibility, and thus the pattern of inventory decision “flips” as the degree of flexibility grows. The intuition behind the flipping effect is that when there is full flexibility, the firm is more concerned about being able to use inventory to free up capacities anywhere in the system because any capacity is fully flexible, therefore placing more inventories to low variability products so that the inventories can be consumed when needed. On the contrary, when there is dedicated design, the firm should place inventories to high variability products, in order to ensure its service level.

We also note that Propositions 3 and 4 do not require the assumption that the capacity loss is positive ($\zeta > 0$), suggesting that demand variability is the main driver for the flipping effect. In section 4, we further confirm this insight as our numerical studies show that the flipping effect also holds when we have deterministic capacity loss, as well as when demand variation and capacity loss are stochastic.

The flipping effect has also been partially observed in production postponement under stochastic demand, as Fisher and Raman (1996) showed that when the second stage production is fully flexible, the firm should produce more low variability products in the first stage. However, Fisher and Raman (1996) do not consider the change in inventory allocations as flexibilities are added into dedicated design. Our finding augments Fisher and Raman (1996), and demonstrates that when the firm adds flexibility to its manufacturing system, its optimal inventory allocation strategy can change drastically, with the trend of allocating more inventories to low variability products. The flipping effect also demonstrates that when a manufacturing system has partial process flexibility, there does not exist a universal rule-of-thumb method for inventory allocation. In these cases, our model and computational algorithm provide a tool for effectively making optimal inventory decisions.

The placement of inventories to low variability products in full flexibility is also related to the *push and pull* strategy in the operations management literature (see, e.g., Simchi-Levi 2010). In a pull strategy, the firm produces to order after demand is realized. In a push strategy, the firm produces to stock and has less ability to adjust. With full flexibility, the firm should apply pull strategy to products with high demand uncertainty, and apply push strategy to products with low demand uncertainty, which means holding more inventory for products with lower variability in the first stage.

4. Application: Risk Mitigation for an Automotive Supply Chain

In this section, we apply our robust optimization method to a case study for the automobile supply chain. We consider an automotive manufacturing system based on the General Motors data

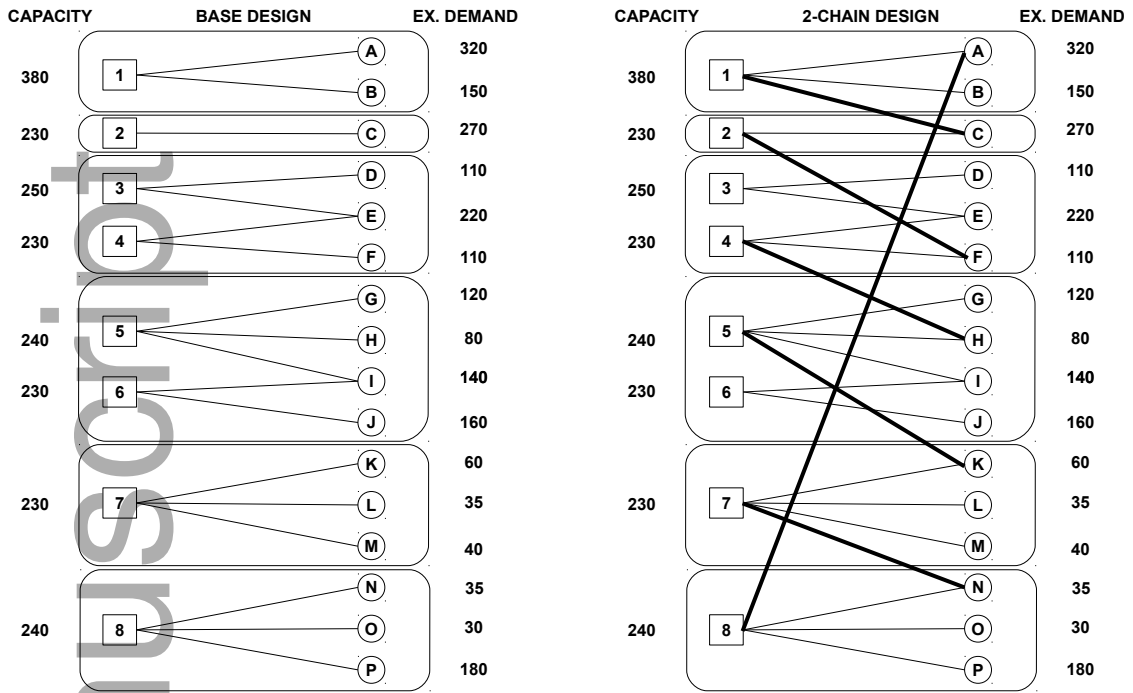
provided by Jordan and Graves (1995). The dataset consists of 16 vehicle models and 8 assembly plants. We use the same plant capacities and mean product demands as in their paper, shown in the left hand side of Figure 3. Since the dataset does not include cost information, we apply the service guarantee (SG) model, and assume all products have a unit holding cost.

The uncertainty set is created according to the section 2.3, with parameters α , β , γ , and ζ . We present our findings with parameters set to $\alpha = 2$, $\beta = 8$ and $\gamma = 0.5 \sum_{j=1}^N \sigma_j$. Furthermore, we assume at most 10% of total capacity may be lost under plant disruption, namely, $\zeta = 0.1 \sum_{i=1}^M \bar{c}_i$. The demand shortage allowance δ is set to 5% of the total mean demand. While we only present our numerical results under this set of parameter values, we have also tested a wide range of different parameter values within 50% of the presented values. For all of the different sets of values tested, we obtained similar qualitative insights.

To simulate the scenarios where products have different demand variability, we assume that half of the products have high variability demand with $\sigma_j = v_h \bar{d}_j$, and the rest of products have low variability demand with $\sigma_j = v_l \bar{d}_j$, where $v_h > v_l > 0$. We tested three groups of variabilities: $(v_h, v_l) = (0.2, 0.1)$, $(0.3, 0.15)$, and $(0.4, 0.2)$. In each case, we test 200 scenarios where high variability products are randomly selected.

To study systems with different flexibility levels, we construct flexibility designs based on the K -chain design in balanced ($M = N$) systems. Recall that in a K -chain design, plant i is capable of producing products i , $i + 1$, to $i + K - 1$. Since the number of plants and products are unequal in this example, we construct “generalized K -chains” by following the clustering idea of Jordan and Graves (1995). More specifically, Jordan and Graves (1995) proposed clustering plants and products into separate groups (six groups in this example), and consider each group as a plant dedicated to producing one product family. Once clustered, the base design can then be viewed as a dedicated or 1-chain design. Finally, to create a generalized K -chain, for each group i , one adds flexibility arcs to connect group i to groups $i + 1$, \dots , $i + K - 1$. An example of the generalized 2-chain is given in the right side of Figure 3.

Figure 3 Flexibility Design in the GM Example



There are many ways to add flexibility links to create a generalized K -chain design: the resulting design depends on the order that the groups appear in the chain, as well as plants and products being connected in each group. Because not all flexible designs are feasible due to physical and technological limitations, and these limitations may differ across different firms, in our simulation, we do not use the optimal K -chain. Instead, for each of the 200 scenarios, we generate a generalized K -chain by randomly selecting a group permutation, and then connect a random plant in each group i to a random product in each of the groups $i + 1, \dots, i + K - 1$. Finally, we also consider full flexibility design, where any plant can produce all products. Note that full flexibility design contains $M \times N = 128$ flexibility arcs, while the base design, 2-chain, 3-chain, and 4-chain contains 16, 24, 32, and 40 flexibility arcs, respectively.

Table 1 and Table 3 list the optimal total inventory cost and the percentage of inventory allocated to high variability products (averaged over 200 random instances) respectively. We start with the observation from Table 1. From Table 1, we see that while changing from the base design to a 2-chain design reduces inventory cost, there is also significant inventory cost reduction from a 2-chain

Table 1 Total Inventory Cost

(v_h, v_l)	Base	2-chain	3-chain	4-chain	Full
(0.2, 0.1)	1219.5 (439%)	573.1 (206%)	343.8 (124%)	292.9 (105%)	277.9
(0.3, 0.15)	1514.3 (430%)	786.4 (224%)	497.4 (141%)	405.2 (115%)	351.8
(0.4, 0.2)	1809.0 (425%)	1014.1 (238%)	670.7 (158%)	535.1 (126%)	425.7

Note: The percentages in round brackets display the cost ratios of the designs to full flexibility.

Table 2 Total Inventory Cost (Deterministic Capacity)

(v_h, v_l)	Base	2-chain	3-chain	4-chain	Full
(0.2, 0.1)	526.2 (189%)	331.6 (119%)	277.9 (100%)	277.9 (100%)	277.9
(0.3, 0.15)	784.3 (223%)	510.4 (145%)	360.6 (103%)	355.1 (101%)	351.8
(0.4, 0.2)	1047.5 (246%)	706.6 (166%)	474.9 (112%)	445.3 (105%)	425.7

Note: The percentages in round brackets display the cost ratios of the designs to full flexibility.

design to designs with higher degrees of flexibility, such as 3-chain or 4-chain. This observation echoes the insight from our analysis in section 3.1, which shows that in the presence of supply disruption, 2-chain is no longer as effective as full flexibility.

We also want to illustrate that the uncertainty in supply significantly increases the ineffectiveness of a 2-chain, compared with systems that suffer from deterministic loss in capacity. To illustrate this, we test another numerical example with the same setting, except that all plant capacities are deterministic and set to 90% of their original capacity levels. The result is shown in Table 2. Note that full flexibility design has the same inventory cost in Table 1 and Table 2, because the two examples assume the same worst-case total capacity. However, the performance of 2-chain is much better in Table 2. Therefore, we conclude that as capacity uncertainty increases, 2-chain becomes less effective.

Next, we discuss the percentages of inventories allocated to high variability products (Table 3). Recall that in every example, 8 products are randomly selected with high demand variability $\sigma_j = v_h \bar{d}_j$, and the remaining 8 products have low variability demand $\sigma_j = v_l \bar{d}_j$. For different variability

Table 3 Percentage of Inventory for High Variability Products

(v_h, v_l)	Base	2-chain	3-chain	4-chain	Full
(0.2, 0.1)	50.1%	49.8%	48.4%	45.6%	41.3%
(0.3, 0.15)	51.4%	52.0%	50.1%	46.2%	38.8%
(0.4, 0.2)	53.9%	54.7%	52.0%	47.6%	33.0%

Table 4 Percentage of Inventory for High Variability Products (Deterministic Capacity)

(v_h, v_l)	Base	2-chain	3-chain	4-chain	Full
(0.2, 0.1)	50.8%	43.6%	42.2%	41.5%	41.3%
(0.3, 0.15)	53.9%	47.2%	41.6%	39.8%	38.8%
(0.4, 0.2)	57.3%	51.0%	43.2%	37.2%	33.0%

parameters (v_h, v_l) , the result shows the same trend: the optimal solutions hold more inventory for products with high demand variability when the degree of flexibility is low (e.g., base design). But as the degree of flexibility increases (e.g., 4-chain, full flexibility), it is optimal to hold more inventory for products with lower variability, and to hold less inventory for products with higher variability. The simulation result verifies the *flipping effect* observed in section 3.2, furthermore, the percentages of inventories allocated to high variability products are monotonically decreasing as we add flexibility to the system. In general, the flipping effect occurs gradually as we add flexibility to the system, and the magnitude of the flipping depends on the variabilities of the products. Interestingly, we find that the flipping effect is not unique to systems with capacity uncertainty. In Table 4, we list the percentages of inventories allocated to high variability products when capacities are deterministically reduced by 10%. As Table 4 shows, the flipping effect still exists when there is no capacity uncertainty.

Furthermore, we note that the findings above are not unique to the robust optimization model, as similar results also hold for the stochastic model (see Online Appendix for a description of the model). In the stochastic model, we consider a case where product demands are independent and normally distributed truncated at two standard deviations. We assume that 8 randomly selected

products have high variability demands with a coefficient of variation equal to 0.5; the remaining 8 products have a coefficient of variation equal to 0.25. We assume that each plant is disrupted independently with probability 0.1. The objective is to find the optimal inventory allocation such that the expected fill rate is at least 95%. The model is solved by the SAA method (described Online Appendix) with 1,000 samples.

Comparing with the robust optimization model, the stochastic model takes *much longer* to solve. For this example with 16 vehicle models, the robust optimization model is solved by constraint generation in less than a second, while the SAA approach takes about 10 minutes. For a large scale network with hundreds of nodes, the robust optimization model can still be solved in a few minutes on a standard laptop, while the SAA model becomes prohibitively large to solve.

Table 5 shows the total inventory level and the percent of the inventory allocated to high variability products, averaged over 200 random examples. The result provides similar qualitative insights obtained from the robust optimization model: 1) there is a significant improvement in inventory cost when more flexibility is added to the 2-chain; 2) less and less inventory is allocated to high demand variability products as the degree of flexibility increases.

Table 5 Inventory Allocation for the Stochastic Model

	Base	2-chain	3-chain	4-chain	Full
Total Inventory	902.4	485.3	335.8	299.6	269.0
High Var Products	44.4%	42.5%	40.2%	34.9%	13.1%

5. Extensions to the Service Guarantee Model

5.1. Time-to-Survive Model

In this subsection, we introduce a metric called Time-to-Survive (TTS) to measure the resilience of supply chain when the duration of the disruption is unknown. Given a facility in the supply chain (e.g., a plant or a distribution center), TTS is defined as the longest time that customer service level is guaranteed if this facility is disrupted. The TTS metric has already been adopted by the

Ford Motor Company to assess risk exposure in its complex supply chain and evaluate Ford's risk mitigation strategies (Simchi-Levi et al. 2014, 2015).

The definition of TTS is motivated by the concept of Time-to-Recover (TTR), which is the time for a facility to return to full capacity after a disruption. TTR is widely used to evaluate supply chain risk (see e.g. Hopp et al. 2012). If TTS is greater than TTR, then a disruption in that facility is not going to affect the firm's service level. On the other hand, when TTS for a specific facility is shorter than TTR, then a disruption at that facility will have an impact on service. Thus, an important challenge in supply chain risk management is to allocate inventory to different products to increase the supply chain's time to survive.

For this purpose, we define the *supply chain TTS* as the minimum (worst) TTS among all of the potential scenarios. The longer the supply chain TTS, the more robust the supply chain is. Below we show that the problem of allocating inventory to maximize supply chain TTS can be reduced to a special case of the model in section 2.

Suppose that plant i has capacity c_i per unit time ($i = 1, \dots, M$), product j has a demand rate d_j per unit time ($j = 1, \dots, N$). Let r_j be the amount of inventory allocated to product j , and assume that the sum of inventory among all products cannot exceed a given budget R . We use $\mathcal{U} \subset \mathbb{R}_+^M \times \mathbb{R}_+^N$ to denote the uncertainty set that models the uncertainty of plant capacity and product demand. We then formulate the problem of maximizing supply chain TTS as the following nonlinear program:

$$\begin{aligned}
 \text{(TTS)} \quad & \max_{r_j, x_{ij}^{(\mathbf{c}, \mathbf{d})}} t \\
 \text{s.t.} \quad & t \leq \frac{r_j}{d_j - \sum_{i: (i,j) \in \mathcal{F}} x_{ij}^{(\mathbf{c}, \mathbf{d})}}, \quad \forall 1 \leq j \leq N, (\mathbf{c}, \mathbf{d}) \in \mathcal{U}, \\
 & \sum_{j: (i,j) \in \mathcal{F}} x_{ij}^{(\mathbf{c}, \mathbf{d})} \leq c_i, \quad \forall 1 \leq i \leq M, (\mathbf{c}, \mathbf{d}) \in \mathcal{U}, \\
 & \sum_{j=1}^N r_j \leq R, \\
 & r_j \geq 0, x_{ij}^{(\mathbf{c}, \mathbf{d})} \geq 0, \quad \forall 1 \leq i \leq M, 1 \leq j \leq N, (\mathbf{c}, \mathbf{d}) \in \mathcal{U}.
 \end{aligned}$$

In this formulation, $x_{ij}^{(\mathbf{c}, \mathbf{d})}$ denotes the production rate of product j by plant i if realized capacity rate and demand rate is (\mathbf{c}, \mathbf{d}) . The first constraint follows the definition of TTS: In any scenario, the supply chain must survive the disruption by continue supplying all products for at least t time units. The second constraint requires that production does not exceed plant capacity. The last constraint specifies that total inventory cannot exceed a given budget R .

Let $s_j = r_j/t$, we can reformulate the TTS model as

$$\begin{aligned}
 (\text{TTS}') \quad & \min_{s_j, x_{ij}^{(\mathbf{c}, \mathbf{d})} \geq 0} \sum_{j=1}^M s_j & (11) \\
 \text{s.t.} \quad & d_j - \sum_{i: (i,j) \in \mathcal{F}} x_{ij}^{(\mathbf{c}, \mathbf{d})} \leq s_j, \quad \forall 1 \leq j \leq N, (\mathbf{c}, \mathbf{d}) \in \mathcal{U}, \\
 & \sum_{j: (i,j) \in \mathcal{F}} x_{ij}^{(\mathbf{c}, \mathbf{d})} \leq c_i, \quad \forall 1 \leq i \leq M, (\mathbf{c}, \mathbf{d}) \in \mathcal{U}, \\
 & s_j \geq 0, x_{ij}^{(\mathbf{c}, \mathbf{d})} \geq 0, \quad \forall 1 \leq i \leq M, 1 \leq j \leq N, (\mathbf{c}, \mathbf{d}) \in \mathcal{U}.
 \end{aligned}$$

This is a special case of the SG model in section 2.1 where the lost sales allowance is equal to $\delta = 0$. Therefore, the analysis and algorithms developed for the SG model also applies to the TTS model.

Based on the optimization model, we note that the optimal objective value of (11) gives the ratio between R , the amount of inventory that the firm holds for all products, and the TTS. As a result, under a fixed flexibility design, its TTS scales *linearly* with inventory budget R . Therefore, under a fixed flexibility design, the firm only needs to compute the TTS model with $R = 1$, in order to consider the full trade-off between keeping inventory and increasing TTS.

To illustrate how the TTS model can be used to measure supply chain resilience, let us consider the GM example introduced in Section 4. Figure 3 in that section shows the base and 2-chain designs for the manufacturing system, as well as the demand rates for all products and the capacities for all production plants. We assume that *at most one* plant is disrupted at any given time, which is the same assumption made for the implementation of the TTS model in Simchi-Levi et al. (2015). Therefore, the capacity uncertainty set can be expressed as:

$$\mathcal{U}_c = \{ \mathbf{c} \in \mathbb{R}^M \mid \sum_{i=1}^M (\bar{c}_i - c_i) / \bar{c}_i \leq 1, 0 \leq c_i \leq \bar{c}, \forall 1 \leq i \leq M \}.$$

For demand uncertainty, the demand uncertainty set follows the structure defined in section 2.3:

$$\mathcal{U}_d = \{\mathbf{d} \in \mathbb{R}^N \mid \sum_{j=1}^N (d_j - \bar{d}_j) \leq \gamma, \sum_{j=1}^N \frac{|d_j - \bar{d}_j|}{\sigma_j} \leq \beta, |d_j - \bar{d}_j| \leq \alpha \sigma_j, \forall 1 \leq j \leq N\}$$

with parameters $\sigma_j = 0.15\bar{d}_j$, $\alpha = 2$, $\beta = 8$, $\gamma = 0.5 \sum_{j=1}^N \sigma_j$. Under these assumptions, we solve the TTS model (11) for GM's network. Note that the TTS model requires that all of the demand in \mathcal{U}_d to be filled during TTS. This is more conservative than the SG model we studied numerically in §4, because in §4, we allow 5% of the total demand to be lost. Following the discussion from §2.1, we can also relax the conservativeness by shrinking the uncertainty set. For example, if we let $\mathcal{U}'_d = 0.95 \cdot \mathcal{U}_d$, the TTS under \mathcal{U}'_d corresponds to firm requiring that 95% of demand for each product to be filled during TTS.

Table 6 and Table 7 list the TTS of GM's network under different flexibility designs given inventory budget R (measured in weeks of demand) with demand uncertainty set \mathcal{U}_d (100% fill rate) and \mathcal{U}'_d (95% fill rate per product), respectively. To populate Table 6 and 7, the TTS under each flexibility design only needs to be computed for $R = 1$, as the TTS under the same flexibility design with $R = x$ can be obtained by multiplying the TTS with $R = 1$ by a factor of x for any positive x . Once the TTS under a flexibility structure is computed with $R = 1$ (one week of extra inventory), the firm can use it to consider how much extra inventory is needed to protect its supply chain against disruptions. For example, under the base design with no flexibility (see Figure 3), if GM wants to maintain 100% of its demand with one week of extra inventory, its TTS is about 0.91 week; and this number would increase threefold if GM is willing to hold three weeks of inventory.

From Table 6 and Table 7, we observe that adding partial flexibility to the system can greatly increase TTS of a supply chain. Under 3-chain, the TTS with 100% fill rate and $R = 1$ is greater than 3 weeks, so GM's supply chain is robust against any disruption less than three weeks with one week of inventory. If GM sets its service level to just 95% for each product, under 3-chain, the TTS is further increased to 3.76 weeks with $R = 1$. Under this circumstance, GM can hedge against disruption for over 11 weeks (almost three months) with just three weeks of inventory.

Table 6 TTS Achieved under Different Inventory Budget and \mathcal{U}_d

	Base	2-chain	3-chain	4-chain	Full
$R = 1$ (week)	0.91	2.00	3.10	3.43	3.77
$R = 2$ (weeks)	1.82	4.00	6.21	6.86	7.54
$R = 3$ (weeks)	2.73	6.00	9.31	10.29	11.30

Table 7 TTS Achieved under Different Inventory Budget and \mathcal{U}'_d

	Base	2-chain	3-chain	4-chain	Full
$R = 1$ (week)	0.96	2.19	3.76	4.22	4.54
$R = 2$ (weeks)	1.92	4.38	7.52	8.44	9.08
$R = 3$ (weeks)	2.89	6.57	11.27	12.66	13.62

5.2. Cost Minimization Model

Sometimes, a supply chain manager does not set a particular service guarantee level, but would rather include the service level as a part of overall cost. In this case, the problem of cost minimization facing demand and capacity uncertainties can also be formulated as a two-stage robust optimization problem, which we call the *cost minimization* (CM) model. Like the service guarantee model, the firm determines the inventory level for each product in the first stage; in the second stage, the firm minimizes the cost based on the realized product demand and available plant capacities.

The total cost in the CM model may include inventory, production, underage (lost sales) and overage costs. For product j , we denote the unit inventory holding cost to be h_j , the unit underage (lost sales) cost to be u_j , and the unit overage (excess inventory) cost to be o_j . The unit product cost of product j at plant i is denoted by f_{ij} . We define $\Pi(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d})$ to be the second stage cost, given flexibility design \mathcal{F} , inventory vector $\mathbf{s} = (s_1, s_2, \dots, s_N)$ (decision from the first stage), and realized capacities \mathbf{c} and demands \mathbf{d} . We introduce a new decision variable, η , to represent the second stage cost. The CM model can be then formulated as follows.

$$(\text{CM}) \quad \min_{\mathbf{s}, \eta} \quad \sum_{j=1}^N h_j s_j + \eta \quad (12)$$

$$\begin{aligned} \text{s.t. } \quad & \Pi(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d}) \leq \eta, \quad \forall (\mathbf{c}, \mathbf{d}) \in \mathcal{U}, \\ & \mathbf{s} \in X_s, \eta \in X_\eta, \end{aligned}$$

where

$$\begin{aligned} \Pi(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d}) = \min_{\mathbf{l}, \mathbf{x}} \quad & \sum_{(S_i, T_j) \in \mathcal{F}} f_{ij} x_{ij} + \sum_{j=1}^N u_j l_j + \sum_{j=1}^N o_j (s_j - d_j)^+ \\ \text{s.t. } \quad & \sum_{i: (S_i, T_j) \in \mathcal{F}} x_{ij} + l_j \geq d_j - s_j, \quad \forall 1 \leq j \leq N, \\ & \sum_{j: (S_i, T_j) \in \mathcal{F}} x_{ij} \leq c_i, \quad \forall 1 \leq i \leq M, \\ & l_j, x_{ij} \geq 0. \end{aligned} \tag{13}$$

In the last constraint of (12), X_s and X_η are polyhedral sets representing constraints on inventory decisions and second stage cost, respectively. Like the SG model, $x_{ij}(\mathbf{c}, \mathbf{d})$ denotes the production of product j by plant i , and $l_j(\mathbf{c}, \mathbf{d})$ denotes the lost sales of product j . To keep notations succinct, we use x_{ij} and l_j instead in formulation (13). As it turns out, the constraint generation algorithm we described in section 2.2 can be also applied to the cost minimization (CM) model.

We note that the cost minimization (CM) model is a generalization of the service guarantee (SG) model. More specifically, for the CM model, if we set $o_j = 0$, $f_{ij} = 0$, $u_j = 1$ for all i, j , and $X_\eta = \{\eta \mid \eta = \delta\}$, then we recover the SG model. Because of its generality, the CM model can be applied to many different problems. While the CM model is more general, the more specialized structure in the SG model makes it much easier to solve numerically, and allows us to incorporate additional features such as TTS (section 5.1) and multi-stage supply chains (section 5.3). Moreover, the simplicity of the SG model can be advantageous; in some situations, it is difficult to accurately estimate the costs of lost sales due to its long term effect (Latour 2001).

5.3. Multi-Stage Service Guarantee Model

In section 2.1, we showed that if the firm specifies a service guarantee for each individual product, the problem can be reduced to a special case of the SG model by setting $\delta = 0$ and modifying the

demand uncertainty set. In this subsection, we show that the SG model can be extended for a multi-stage supply chain when $\delta = 0$.

Suppose there are K stages in the supply chain. For all $k = 1, \dots, K$, stage k consists of M^k plants/suppliers, and N^k products/components. The stage indices are counted backward, so the final products are produced in stage 1, and the products/components in stage k require components from stage $k + 1$. Furthermore, we assume the bill of material information at stage k is represented by a matrix \mathbf{A}^k , such that one unit of product j at stage k requires $\mathbf{A}_{j'j}^k$ units of component j' at stage $k + 1$.

Let s_j^k be the inventory of product j at stage k , with unit holding cost h_j^k . Let $\tilde{\Pi}(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d})$ be the total lost sales for the final products given inventory \mathbf{s} , capacity \mathbf{c} and demand of final products \mathbf{d} . The multi-stage service guarantee (m-SG) model can be formulated as

$$\begin{aligned}
 (\mathbf{m}\text{-SG}) \quad & \min \sum_{k=1}^K \sum_{j=1}^{N^k} h_j^k s_j^k \\
 \text{s.t.} \quad & \tilde{\Pi}(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d}) \leq 0, \quad \forall (\mathbf{c}, \mathbf{d}) \in \mathcal{U} \\
 & s_j^k \geq 0,
 \end{aligned}$$

where

$$\begin{aligned}
 \tilde{\Pi}(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d}) = \min \quad & \sum_{j=1}^{N^1} l_j \\
 \text{s.t.} \quad & \sum_{i: (S_i^1, T_j^1) \in \mathcal{F}} x_{ij}^1 + l_j \geq d_j - s_j^1, \quad \forall 1 \leq j \leq N^1 \\
 & \sum_{i: (S_i^{k+1}, T_{j'}^{k+1}) \in \mathcal{F}} x_{ij'}^{k+1} + s_{j'}^{k+1} \geq \sum_{j=1}^{N^k} \mathbf{A}_{j'j}^k \sum_{i: (S_i^k, T_j^k) \in \mathcal{F}} x_{ij}^k, \quad \forall 1 \leq k \leq K-1 \\
 & \sum_{j: (S_i, T_j) \in \mathcal{F}} x_{ij}^k \leq c_i^k, \quad \forall 1 \leq k \leq K, 1 \leq i \leq M^k \\
 & l_j, x_{ij}^k \geq 0.
 \end{aligned}$$

In the subproblem $\tilde{\Pi}(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d})$, the first set of constraints defines the lost sales of final products. The second set of constraints ensures that there are enough components at stage $k + 1$ to satisfy production at stage k . The third set of constraints enforces capacity limits.

The m-SG model can also be solved efficiently using constraint generation algorithm (Wang 2016). To apply constraint generation, we can first decompose the subproblems for different stages, using the fact that the lost sales allowance δ is zero and making affine transformation on the uncertainty sets. To this end, we define uncertainty sets $\mathcal{U}^1 = \mathcal{U}$, and $\mathcal{U}^{k+1} = \{(\mathbf{c}, \mathbf{A}^k \mathbf{d}^k) \mid (\mathbf{c}, \mathbf{d}^k) \in \mathcal{U}^k\}$ for all $k = 1, \dots, K-1$. Let $\tilde{\Pi}^k(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d}^k)$ be the lost sales of stage k , we have

$$\begin{aligned}
 (\mathbf{m}\text{-SG}') \quad & \min \sum_{k=1}^K \sum_{j=1}^{N^k} h_j^k s_j^k \\
 \text{s.t.} \quad & \tilde{\Pi}^k(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d}^k) \leq 0, \quad \forall (\mathbf{c}, \mathbf{d}^k) \in \mathcal{U}^k, k = 1, \dots, K \\
 & s_j^k \geq 0,
 \end{aligned}$$

where

$$\begin{aligned}
 \tilde{\Pi}^k(\mathcal{F}, \mathbf{s}, \mathbf{c}, \mathbf{d}^k) = \min \quad & \sum_{j=1}^{N^k} l_j \\
 \text{s.t.} \quad & \sum_{i: (S_i^k, T_j^k) \in \mathcal{F}} x_{ij}^k + l_j \geq d_j^k - s_j^k, \quad \forall 1 \leq j \leq N^k \\
 & \sum_{j: (S_i, T_j) \in \mathcal{F}} x_{ij}^k \leq c_i^k, \quad \forall 1 \leq i \leq M^k \\
 & l_j, x_{ij}^k \geq 0.
 \end{aligned}$$

We note that for the general case of the SG model with $\delta > 0$, this reformulation technique does not work because the subproblem cannot be decomposed for different stages. Nevertheless, even when δ is fixed to zero, the firm can still adjust its target service guarantee by using different demand uncertainty sets (see section 2.1).

6. Conclusion

We consider a firm that coordinates process flexibility and inventory decisions to mitigate supply risk. This interplay of process flexibility and inventory is modeled as a two-stage robust optimization problem, which can be solved efficiently using constraint generation algorithm. Under some symmetric assumptions, the solution of the robust optimization model can be characterized, allowing us to better understand the change in optimal inventory as flexibility or capacity uncertainty

changes. Using the characterization, we find that the widely used 2-chain structure does not provide enough protection against supply disruption. Moreover, we show that the optimal inventory decision critically depends on the flexibility design of the supply network. In particular, there tend to be a shift of inventories from high variability products to low variability products as the degree of process flexibility increases. We refer to such phenomenon as the *flipping effect*, which is demonstrated using both theoretical analysis and numerical experiments.

Our robust optimization model is applicable to various types of risk mitigation problems, as the model can be extended to incorporate different underage, overage, and inventory costs. The model can be also extended to incorporate multi-stage supply chain and unknown disruption time. In the numerical experiment, we apply the robust optimization method to a risk mitigation example in an automotive supply chain. This application demonstrates that our model can be an effective tool for analyzing the hybrid risk mitigation strategy that combines process flexibility and inventory.

References

- Arreola-Risa, A. and DeCroix, G. (1998). Inventory management under random supply disruptions and partial backorders. *Naval Research Logistics*, 45(7):687–703.
- Atamtürk, A. and Zhang, M. (2007). Two-stage robust network flow and design under demand uncertainty. *Operations Research*, 55(4):662–673.
- Bandi, C. and Bertsimas, D. (2012). Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical Programming*, 134(1):23–70.
- Bassamboo, A., Randhawa, R. S., and Mieghem, J. A. V. (2012). A little flexibility is all you need: on the asymptotic value of flexible capacity in parallel queuing systems. *Operations Research*, 60(6):1423–1435.
- Bassamboo, A., Randhawa, R. S., and Van Mieghem, J. A. (2010). Optimal flexibility configurations in newsvendor networks: Going beyond chaining and pairing. *Management Science*, 56(8):1285–1303.
- Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. (2009). *Robust optimization*. Princeton University Press.
- Ben-Tal, A., Goryashko, A., Guslitzer, E., and Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376.

- Bertsimas, D., Brown, D., and Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM review*, 53(3):464–501.
- Bertsimas, D. and Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations Research*, 54(1):150–168.
- Bertsimas, D. and Tsitsiklis, J. (1997). *Introduction to linear optimization*. Athena Scientific Belmont, MA.
- Bish, E. and Wang, Q. (2004). Optimal investment strategies for flexible resources, considering pricing and correlated demands. *Operations Research*, 52(6):954–964.
- Bollapragada, R., Rao, U. S., and Zhang, J. (2004). Managing inventory and supply performance in assembly systems with random supply capacity and demand. *Management Science*, 50(12):1729–1743.
- Bonami, P., Kilinc, M., and Linderoth, J. (2012). Algorithms and software for convex mixed integer nonlinear programs. In *Mixed Integer Nonlinear Programming*, pages 1–39. Springer.
- Chopra, S. and Sodhi, M. S. (2014). Reducing the risk of supply chain disruptions. *MIT Sloan management review*, 55(3):73.
- Chou, M., Teo, C.-P., and Zheng, H. (2011). Process flexibility revisited: The graph expander and its applications. *Operations Research*, 59(5):1090–1105.
- Chou, M. C., Chua, G. A., Teo, C.-P., and Zheng, H. (2010). Design for process flexibility: Efficiency of the long chain and sparse structure. *Operations Research*, 58(1):43–58.
- Chou, M. C., Chua, G. A., and Zheng, H. (2014). On the performance of sparse process structures in partial postponement production systems. *Operations Research*, 62(2):348–365.
- Culp, S. (2013). Supply chain disruption a major threat to business. *Forbes*. Published on February 15, 2013.
- DeCroix, G. A. (2013). Inventory management for an assembly system subject to supply disruptions. *Management Science*, 59(9):2079–2092.
- Dyer, M. and Stougie, L. (2006). Computational complexity of stochastic programming problems. *Mathematical Programming*, 106(3):423–432.
- Fine, C. H. and Freund, R. M. (1990). Optimal investment in product-flexible manufacturing capacity. *Management Science*, 36(4):449–466.

- Fisher, M. and Raman, A. (1996). Reducing the cost of demand uncertainty through accurate response to early sales. *Operations research*, 44(1):87–99.
- Gabrel, V., Murat, C., and Thiele, A. (2014). Recent advances in robust optimization: An overview. *European Journal of Operational Research*, 235(3):471–483.
- Graves, S. C. and Willems, S. P. (2000). Optimizing strategic safety stock placement in supply chains. *Manufacturing & Service Operations Management*, 2(1):68–83.
- Gürler, Ü. and Parlar, M. (1997). An inventory problem with two randomly available suppliers. *Operations Research*, 45(6):904–918.
- Hopp, W., Iravani, S., and Liu, Z. (2012). Mitigating the impact of disruptions in supply chains. In Gurnani, H., Mehrotra, A., and Ray, S., editors, *Supply Chain Disruptions*, pages 21–49. Springer London.
- Hopp, W. J., Tekin, E., and Van Oyen, M. P. (2004). Benefits of skill chaining in serial production lines with cross-trained workers. *Management Science*, 50(1):83–98.
- Jansen, B. (2017). Delta outage a reminder of fragility of airline computers. *USA TODAY*. Published on January 30, 2017.
- Jordan, W. C. and Graves, S. C. (1995). Principles on the benefits of manufacturing process flexibility. *Management Science*, 41(4):577–594.
- Kim, C.-R. K. (2011). Toyota aims for quake-proof supply chain. *Reuters*. September 6, 2011.
- Latour, A. (2001). Trial by fire: A blaze in albuquerque sets off major crisis for cell-phone giants. *The Wall Street Journal*. Published on January 29, 2011.
- Meyer, R. R., Rothkopf, M. H., and Smith, S. A. (1979). Reliability and inventory in a production-storage system. *Management Science*, 25(8):799–807.
- O’Marah, K. (2017). Supply chain risk 2020: New worries. *Forbes*. Published on October 6, 2017.
- Shapiro, A. and Nemirovski, A. (2005). On complexity of stochastic programming problems. In *Continuous optimization*, pages 111–146. Springer.
- Simchi-Levi, D. (2010). *Operations Rules: Delivering Customer Value through Flexible Operations*. The MIT Press.

- Simchi-Levi, D., Schmidt, W., and Wei, Y. (2014). From superstorms to factory fires: Managing unpredictable supply chain disruptions. *Harvard Business Review*, 92(1–2):96–101.
- Simchi-Levi, D., Schmidt, W., Wei, Y., Zhang, P. Y., Combs, K., Ge, Y., Gusikhin, O., Sander, M., and Zhang, D. (2015). Identifying risks and mitigating disruptions in the automotive supply chain. *Interfaces*, 45(5):375–390.
- Simchi-Levi, D. and Wei, Y. (2012). Understanding the performance of the long chain and sparse designs in process flexibility. *Operations Research*, 60(5):1125–1141.
- Sodhi, M. S. and Tang, C. S. (2012). Strategic approaches for mitigating supply chain risks. In Hillier, F. S., editor, *Managing Supply Chain Risk*, volume 172 of *International Series in Operations Research and Management Science*, pages 95–108. Springer, New York.
- Song, J.-S. and Zipkin, P. H. (1996). Inventory control with information about supply conditions. *Management Science*, 42(10):1409–1419.
- Swaminathan, J. M. and Tayur, S. R. (1998). Managing broader product lines through delayed differentiation using vanilla boxes. *Management Science*, 44(12-part-2):S161–S172.
- Tang, C. and Tomlin, B. (2008). The power of flexibility for mitigating supply chain risks. *International Journal of Production Economics*, 116(1):12–27.
- Thiele, A., Terry, T., and Epelman, M. (2010). Robust linear optimization with recourse. University of Michigan, IOE Technical Report TR09-01.
- Tomlin, B. (2006). On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science*, 52(5):639–657.
- Tomlin, B. and Wang, Y. (2005). On the value of mix flexibility and dual sourcing in unreliable newsvendor networks. *Manufacturing & Service Operations Management*, 7(1):37–57.
- Wang, H. (2016). *Dynamic learning and optimization for operations management problems*. PhD thesis, Massachusetts Institute of Technology.
- Wang, X. and Zhang, J. (2015). Process flexibility: A distribution-free bound on the performance of k-chain. *Operations Research*.

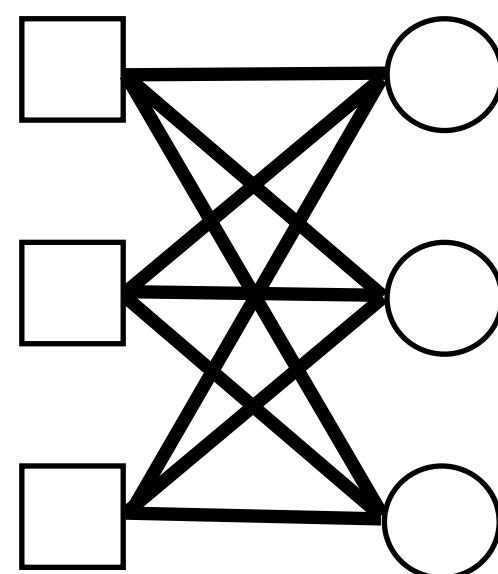


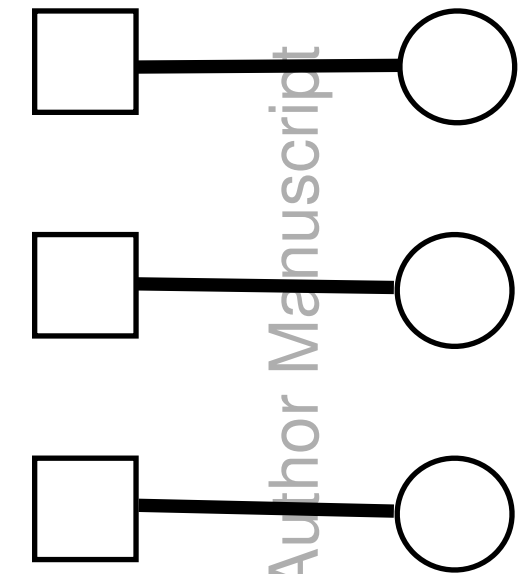
Zeng, B. and Zhao, L. (2013). Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–461.

Author Manuscript

Plant

Product

poms_12887_f1.pdf



Author Manuscript

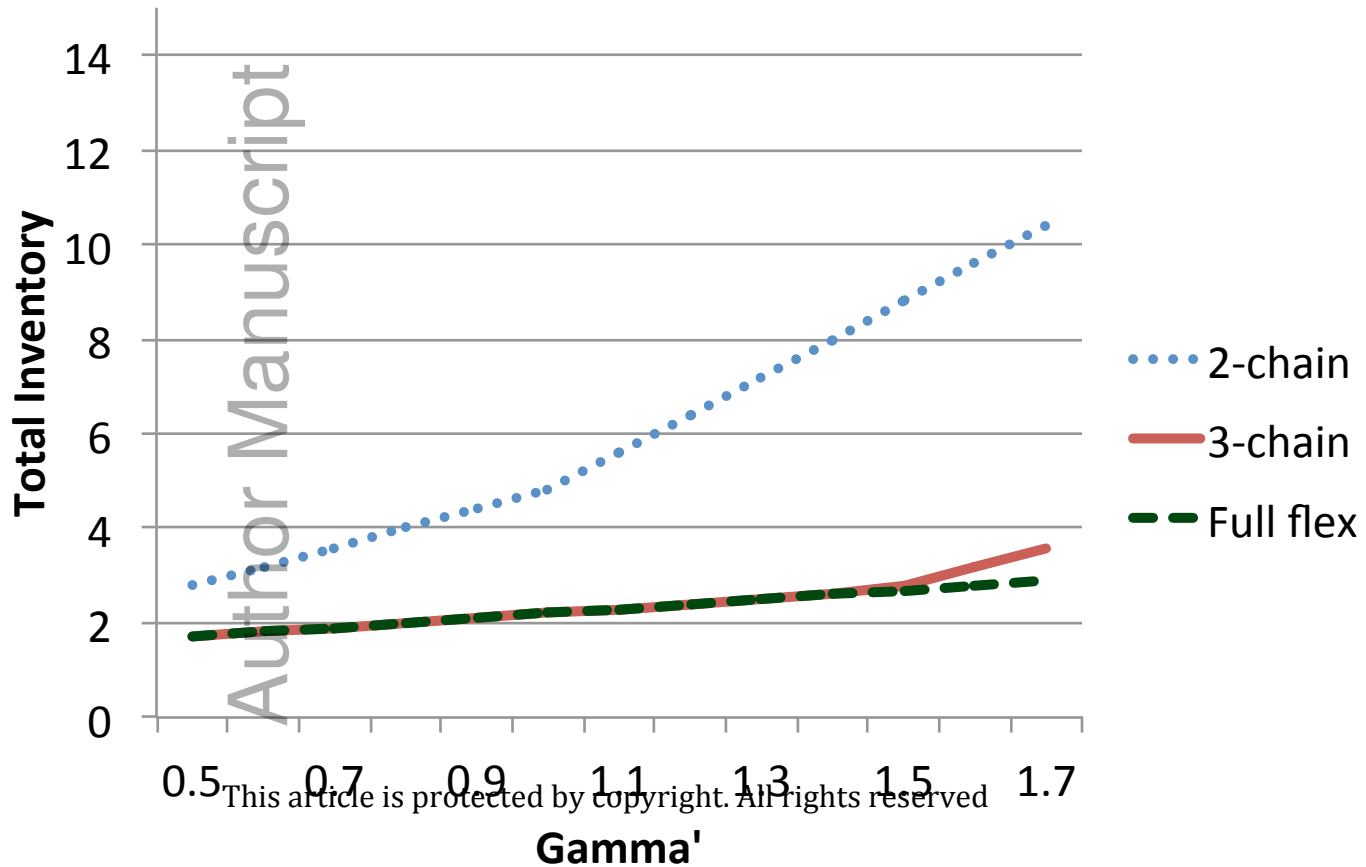
This article is protected by copyright. All rights reserved

Dedicated


2-Chain

Partial Flexibility

Full Flexibility




Author Manuscript



CAPACITY

BASE DESIGN

EX. DEMAND



CAPACITY

2-CHAIN DESIGN

EX. DEMAND

