

3D-Printing for Supply Chain Resilience

Ziyu He,^{a,b} Vishal Gupta,^a Nick Vyas^{a,b}

^aDepartment of Data Science and Operations, University of Southern California, Los Angeles, California, 90089; ^bRandall R. Kendrick Global Supply Chain Institute, University of Southern California, Los Angeles, California, 90089

Abstract. Problem definition: Most strategies for mitigating supply chain disruptions require upfront, dedicated investments for each supplier, making them impractical for large, disaggregated supply chains. We study using 3D printing (3DP) as a *flexible*, backup resource that can support any disrupted supplier. 3DP has traditionally been rejected as a viable resilience strategy due to high per-unit production costs and limited capacity. Inspired by the “a little flexibility is enough” literature, however, we explore when and why 3DP may be a cost-effective resilience strategy and how to deploy it. **Methodology/results:** We formulate an optimization model to decide which suppliers to backup with 3DP, how much 3DP capacity to procure, and how to allocate that capacity in real-time to unmet demand. The resulting mixed-binary, stochastic optimization problem is computationally challenging, even for moderate sized supply chains. Hence, we propose a novel algorithmic framework combining supermodular approximations based on Taylor series and first-order stochastic optimization to compute high-quality feasible solutions. Finally, we conduct an empirical case study based on bill-of-lading data from toy manufacturer Mattel. With cost estimates based on current technology, 3DP offer modest cost savings relative to traditional resilience strategies. However, its principal benefit is a reduction in demand shortfalls. This advantage is especially evident in larger systems with weakly correlated supplier disruptions. **Managerial implications:** Resonating with the well-known “a little flexibility is enough” principle, our findings provide both analytical and empirical evidence of 3DP’s transformative potential as a strategic resilience tool for large supply chains. We also provide concrete guidance on how to introduce 3DP into an existing portfolio of supply chain resilience strategies to complement existing capabilities.

Funding: The second author was supported by Randall R. Kendrick Global Supply Chain Institute. The first and second author were partially supported by a grant from the Institute for Outlier Research in Business at USC Marshall.

Key words: 3D-printing, Supply chain resilience, Stochastic mixed-integer programming, Supermodular minimization

1. Introduction

Managing supply chain disruptions has become part of normal operations of a firm (Simchi-Levi et al. 2015a). Some estimates suggest that poorly managed disruptions can incur up to a 45% loss in firm's annual earnings over a decade (McKinsey & Company 2021), prompting most firms to invest in supply chain resilience strategies. Indeed, in a recent survey, 97% of respondents had implemented at least one resilience measure following COVID-19 (McKinsey & Company 2023).

Common resilience strategies include inventory buffering (Simchi-Levi et al. 2015b, 2017), dual sourcing/siting (Tomlin 2006, Dada et al. 2007), and reserving capacity at backup suppliers (Yang et al. 2009, Demirel et al. 2017). Typically, these strategies are applied in a dedicated, product-specific manner, requiring separate upfront investments for each protected component before a disruption occurs. For example, an automotive company might pre-stock separate inventory buffers to guard against disruptions in its brake and engine parts.

Such dedicated strategies can be cost-prohibitive in large, disaggregated supply chains. For example, Simchi-Levi et al. (2015a) reports that Ford's supply chain involves over 35 billion parts, with 14,000 Tier 1 suppliers across 4,400 sites. Worse, a large fraction of suppliers provide high-volume, low-cost, low-margin components that are nonetheless critical to the final product. In such settings, firms are seemingly forced to choose which products merit an expensive dedicated back-up strategy, and which should be left unprotected.

In this paper, we explore a third, unconventional option: leveraging 3D Printing (3DP) as a "flexible" back-up resource. By building items layer by layer from scratch, a single printer can produce a wide range of items without incurring the product-specific fixed costs of traditional manufacturing techniques like injection molding. Thus, although it requires an upfront capital investment before disruption occurs, 3DP is flexible in the sense that we can strategically choose which components to print on-demand, *after the disruption*. Intuition suggests this flexibility might offer significant value precisely in large, disaggregated supply chains. Moreover, the well-known principle that "a little flexibility is all you need" from both manufacturing (Jordan and Graves 1995, Simchi-Levi and Wei 2012) and service systems (Bassamboo et al. 2012, Tsitsiklis and Xu 2013) suggests that even a small amount of well-used, 3DP capacity might yield substantive savings.

Despite this intuition, 3DP has traditionally *not* been seen as a viable resilience strategy. It often incurs high per-unit production costs, slower production speeds, and potentially expensive capital investments in printers. Research in operations management has largely focused on using 3DP's in low-volume manufacturing, such as spare parts inventory (Song and Zhang 2020, Zhang et al. 2022, Westerweel et al. 2021) and product customization (Chen et al. 2021, Sethuraman et al. 2023).

Outside academia, however, some firms have started exploring 3DP as a supply chain resilience measure. For example, during the COVID-19 Pandemic, CNH Industrial, a major farm equipment manufacturer, experienced a disruption of a small, but vital clamping fixture. Stopping assembly until the parts arrived would have incurred a loss of €189,000 (Materialise 2023). Instead, CNH chose to print missing clamp fixture. Even though printing costs were more than 7 times larger per part than injection molding, printing costs totaled a mere €806, the needed parts were ready in a week, and introduced a minimal delay (Materialise 2023). Implicit in this example is the idea that despite its criticality, CNH did not backup the clamp fixture, presumably because it was too expensive to merit doing so in a dedicated fashion. As a second example, in September of 2021, a change was made to the design of the close-out seal of the spoiler of the 2022 Chevy Tahoe. Approximately 60,000 new parts were needed in 6 weeks time to avoid delaying the release the truck. Manufacturing the part via injection molding was estimated to take 12 weeks (McEachern 2022). Instead, General Motors opted to print the requisite parts, meeting its desired time-frame (Stevens 2022). Again, note the dedicated back-up option was prohibitively expensive (in terms of set up time), whereas a the 3DP solution was more expensive per part, but faster to bring online. More broadly, GM has also invested in an additive manufacturing facility (Lopez 2020) to explore similar uses of 3DP. Anecdotal evidence like this strongly suggests rethinking the potential role of additive manufacturing in supply chain resilience.

To this end, we propose a novel model of 3DP and supply chain resilience in which a firm sources multiple products, each from its own primary supplier, and each supplier may suffer a random disruption (modeled as random yield). Prior to sourcing, the firm can select one of two distinct backup strategies for each product – dedicated backup or 3DP – that can be utilized to meet unmet demand due to disruptions. To capture the key trade-offs between the two strategies, we model dedicated backup as having unlimited capacity and a small per-unit cost of production, but incurring a fixed cost per product it protects. By contrast, 3DP requires only a single capacity investment to protect many products, but the per unit cost of production is higher and the total recoverable demand is limited by this capacity. To facilitate an apples-to-apples comparison, we assume both backup strategies are make-to-order, i.e., we neglect any potential benefits of 3DP arising from co-locating printers closer to demand and optimistically assume dedicated back-ups can be brought online instantaneously. Both assumptions favor dedicated back-ups over 3DP. Indeed, throughout, we make modeling assumptions which favor dedicated back-ups in order to conservatively assess the potential value of 3DP.

With this model, we prove certain structural features of the optimal policy after introducing 3DP: The sets of products backed up by 3DP and by dedicated strategies form a partition, i.e., no product is unprotected (i.e. not backed up by any resource) nor is any product doubly backed up. Moreover, the set of products backed up by 3DP is a (potentially strict) superset of those that are unprotected in the absence of 3DP. With respect to the first stage order quantities, we can upper and lower bound these quantities by simple, easily computed formulas. These formulas are immediately interpretable as the first stage order quantities in a proxy system with infinite 3DP printing capacity and a proxy system with no printing capacity. Finally, we can also bound the optimal printing capacity by an appropriate quantile of the demand shortfall in those two proxy systems. Overall, these structural features provide insights into which types of products might benefit from 3DP backup, how much capital investment might be needed and that systems where suppliers fail independently (or are uncorrelated) benefit most from 3DP backup.

Moving beyond bounds to precise values, however, requires solving for the firm's optimal strategy. We formulate this problem as a mixed-binary, stochastic optimization problem and describe an exact algorithm for computing an optimal solution based on mixed integer optimization and Bender's cuts. Unfortunately, for even moderately sized supply chains, this exact approach is computationally expensive. Part of the challenge is that the problem is neither submodular nor supermodular in the set of products backed up by 3DP.

Consequently, we develop an efficient algorithmic approach to find high-quality feasible solutions. Specifically, because it is so expensive, we might intuit that the optimal amount of 3DP capacity is small relative to the demand. Inspired by this intuition, we show that replacing the expected second stage costs by a suitable Taylor series expansion around a printing capacity of zero yields an approximate objective function that is supermodular, but non-monotonic. While non-monotonic supermodular minimization is NP-Hard, it admits highly efficient approximation algorithms Feige et al. (2011). Our procedure combines these approximation algorithms with a polishing step based on projected stochastic gradient descent and scales easily to very large supply chains.

While our above algorithm computes a near-optimal policy, in practice, a firm not currently engaging in 3DP may be reticent to change their entire resilience strategy. More likely, they might consider using 3DP to back up otherwise unprotected products, i.e., products not currently backed up by a dedicated resource. For such a firm, we provide a necessary and sufficient condition for 3DP to offer a savings, and show that when there exists a positive savings, any positive, sufficiently small capacity investment yields a benefit. Hence, a firm might truly benefit from introducing just "a little" flexibility.

We conclude with an extensive numerical case study using a combination of real-world bill-of-lading data, published 3D printing specifications, and web-scraped prices for Mattel, a global leader in toy manufacturing. Our analysis validates our key theoretical insights. First, our algorithmic approach provides extremely high-quality solutions, often within a few percent of optimal. Second, under current estimates of printing costs, 3DP typically covers 40–50% of Mattel’s product portfolio. This includes all previously unprotected products (approximately 20–40% of the total) and around 20% of the products previously backed-up by dedicated resources. Even marginal investments in 3DP capacity can yield a savings. For example, 3DP achieves a 4% cost savings and reduces demand shortfalls by 50% relative to relying solely on dedicated backups, even with a capacity investment of just 5% of total demand. Finally, the most significant benefits of 3DP are in meeting demand shortfalls. Even small investments can help reduce the probability of large shortfalls significantly.

We summarize our contributions as follows:

- In Section 2, we propose a model of a firm’s choice of resilience strategies that captures the essential trade-offs between 3DP and dedicated back-up strategies. We derive certain structural features of the optimal policy in terms of which products are backed up by 3DP, how primary order quantities change relative to a system without 3DP, and how large the optimal capacity investment might be.
- In Section 4 we formulate a mixed-binary stochastic optimization problem for firm’s optimal policy and describe a solution approach via Bender’s cuts. Since the approach does not scale effectively for large supply-chains, we propose a heuristic algorithm leveraging a series of approximations, supermodularity, and a polishing step in Section 5. We argue this heuristic approach provides high-quality solutions when the optimal capacity investment is small.
- In Section 6, we consider a firm that is not currently engaged in 3DP that is only willing to make a small investment. We provide a necessary and sufficient condition whereby any sufficiently small investment in 3DP yields a cost-savings.
- Finally, in Section 7, we assess the quality of our proposed heuristic approach and assess the value of 3DP as a resilience strategy through extensive numerical experiments using real data for Mattel. We find that our heuristic approach yields solutions within a few percent of optimal for moderate chains, and that for Mattel’s full chain, even a small investment in 3DP capacity could yield substantive cost-savings and reduce demand shortfalls.

1.1. Additional Related Literature

As mentioned, much of the existing literature focuses on low-volume manufacturing applications, such as spare-parts (Song and Zhang 2020, Zhang et al. 2022, Westerweel et al. 2021) and product

customization (Chen et al. 2021, Sethuraman et al. 2023), possibly because these applications require only small amounts of 3DP capital investment. By contrast, we study 3DP in supply chain resilience and quantify the required (possibly large) 3DP capital investment (relative to dedicated strategies).

In context of supply chain, Arbabian and Wagner (2020) consider co-locating printers closer to demand. Co-location avoids transportation costs and upstream holding costs, and makes the system more responsive to demand. They quantify these benefits (and how they depend on problem parameters) in a one manufacturer/one retailer setting with a single product. In our paper, we explicitly neglect the benefits of co-location and treat a chain with multiple suppliers to focus exclusively on the benefits of flexibility from 3DP.

Dong et al. (2022) also study 3DP flexibility, specifically how introducing 3DP affects the optimal product offering in design-intensive industries (home furnishing, apparel, jewelry). They focus on three aspects of 3DP printing: i) Design Freedom – 3DP can create products not possible with traditional techniques ii) Quality – 3DP may produce a higher or lower quality product than traditional techniques iii) Flexibility – 3DP can print multiple different types of products in a single run. Those authors develop structural insights on the optimal product offering under a particular choice model for demand and how the above features affect that assortment. Our work adds to this study of 3DP flexibility by considering a large, disaggregated supply chain. Motivated by the previous anecdotes, our primary focus is on low-cost, low-margin components – not design-intensive products. Hence, issues around customer perceptions of quality and design freedom are arguably second order relative to quantifying the needed capital investment in 3DP and identifying the “right” products to backup.

We also connect to the broader work studying flexible backup strategies in supply chain including Saghafian and Van Oyen (2016, 2012). Saghafian and Van Oyen (2016) study a multi-product, multi-supplier network with a rich, Markov Chain model of supply chain disruption that can capture heterogeneous rates of disruption and lengths of disruption. Their focus is on establishing that “a little flexibility” is sufficient to capture the benefits of a fully-flexible backup system. The model studied, however, is very general and technically challenging. While it can in principle be solved via an infinite dimensional Bellman equation, in practice, solving such problems is notoriously challenging. The authors do not offer a specialized algorithm for this task. Consequently, for the most part, the work does not study the question of *which* suppliers to back up with the flexible resource, which is a key question in our work. Indeed, the only results around choosing suppliers for flexible backup from Saghafian and Van Oyen (2016) are with respect to a simplified model with only 2 suppliers. Our work on the other hand adopts an admittedly coarser model of disruptions (random yield)

but seeks to develop a general purpose algorithm for identifying which suppliers to back up with 3DP. We demonstrate empirically that our algorithm scales gracefully to very large supply chains.

Finally, we contrast our work to Wang and Webster (2022). In that work, authors also study backup flexibility, but consider a flexible backup that might fail. By considering a 2 product system, they show that flexibility is not always valuable and one must consider whether to invest in flexibility amongst primary suppliers, or amongst backups. The notion of flexibility in this work is generic, meant to represent “supplier development efforts.” In our work, we focus explicitly on 3DP, and, consequently, model it as a *reliable* but expensive backup strategy. We again focus on a multiproduct system where the challenge is identifying which products to backup.

2. Model Setup

We consider a firm that sources n products (indexed by $j \in \mathcal{N} \equiv \{1, \dots, n\}$), where each product is sourced from a distinct supplier. In the absence of supply chain disruptions, the firm orders q_j units from the j^{th} primary supplier at a cost of c_j per unit and then sees a random demand D_j . Then, the firm pays holding costs of h_j per unit for excess inventory and v_j per unit of unmet demand. Thus, absent disruptions, the firm faces a simple newsvendor problem for the j^{th} product.

We, however, consider a setting with disruptions. Namely, let $s_j \in [0, 1]$ be a random variable representing the yield of supplier j . After ordering, the firm receives $q_j s_j$ units, and only pays for received units. To hedge against this yield uncertainty and the possibility of unmet demand, the firm can choose to invest in one of two resilience strategies for each product j :

- (1) DEDICATED BACKUPS (DB): The firm can order q_j^{DB} units from an expediting supplier at a per unit cost of c_j^{DB} . Engaging in this strategy also incurs a one-time, upfront fixed cost of C_j^{DB} , irrespective of the order quantity. This fixed cost implicitly models the cost of reserving production capacity with this expediting supplier.
- (2) 3D PRINTING (3DP): The firm can choose to use a 3D-printer to produce q_j^{3DP} units of product j at a per-unit cost of c_j^{3DP} . Unlike dedicated backups, 3DP has a finite capacity K across all products, i.e., we must have $\sum_{j \in \mathcal{N}} q_j^{\text{3DP}} \leq K$. We model the cost of investing in 3DP capacity by a non-decreasing, convex function $C^{\text{3DP}}(K)$, with $C^{\text{3DP}}(0) = 0$, paid before demand and yield are realized.

The above choice must be made for each product j . Let $\mathcal{A} \subseteq \mathcal{N}$ denote the set of products backed up using the 3DP strategy, let $\mathcal{T} \subseteq \mathcal{N}$ represent those backed up by DB, and note that a priori these sets need not be disjoint. We seek to choose $\mathcal{A}$ and $\mathcal{T}$ to minimize the total expected costs.

To avoid several trivial scenarios, we will assume throughout that:

ASSUMPTION 1 (Non-trivial Parameters). For each $j \in \mathcal{N}$, $c_j < c_j^{\text{DB}} < c_j^{\text{3DP}} < v_j$.

In particular, $v_j \geq \max(c_j^{\text{DB}}, c_j^{\text{3DP}})$ ensures that meeting unmet demand is always preferable to stocking out, while $c_j < \min(c_j^{\text{DB}}, c_j^{\text{3DP}})$ ensures that sourcing from primary suppliers is preferable to expediting or printing. Finally, $c_j^{\text{DB}} < c_j^{\text{3DP}}$ reflects the typical cost relationship between DB and 3DP given current technology.

In summary, the sequence of events are:

- i) The firm chooses sets $\mathcal{A}$ and $\mathcal{T}$ of products backed up by 3DP and DB, respectively, and a capacity $K \in [0, \infty)$ for the 3D printing resource. These choices induce a total fixed cost of $\sum_{j \in \mathcal{T}} C_j^{\text{DB}} + C^{\text{3DP}}(K)$.
- ii) The firm orders q_j units from the j^{th} supplier for each $j \in \mathcal{N}$.
- iii) The firm observes random yield s_j and demands D_j . It pays $\sum_{j \in \mathcal{N}} c_j q_j s_j$ for all successful deliveries.
- iv) The firm orders q_j^{DB} units from the expedited supplier for each $j \in \mathcal{T}$ and prints q_j^{3DP} units on the 3D printer, subject to $\sum_{j \in \mathcal{A}} q_j^{\text{3DP}} \leq K$, inducing a cost of $\sum_{j \in \mathcal{T}} c_j^{\text{DB}} q_j^{\text{DB}} + \sum_{j \in \mathcal{A}} c_j^{\text{3DP}} q_j^{\text{3DP}}$.
- v) The firm pays stock-out costs v_j per unit of any unfulfilled demand and holding costs h_j per unit of any excess production for each $j \in \mathcal{N}$.

Our goal is to minimize the total expected costs, i.e.,

$$\min_{\mathcal{T}, \mathcal{A} \subseteq \mathcal{N}, K \geq 0} C^{\text{3DP}}(K) + \sum_{j \in \mathcal{T}} C_j^{\text{DB}} + U^{\text{3DP}}(\mathcal{A}, K) + U^{\text{DB}}(\mathcal{T}) \quad (1)$$

where $U^{\text{DB}}(\mathcal{T})$ and $U^{\text{3DP}}(\mathcal{A}, K)$ represent the expected operational costs incurred from iii) to v). These terms are described by optimization problems over $q_j, q_j^{\text{3DP}}, q_j^{\text{DB}}$. For completeness, we state them now, but provide detailed derivations and commentary on these problems in Sections 2.1 and 2.2 below:

$$U^{\text{DB}}(\mathcal{T}) \equiv \sum_{j \in \mathcal{T}} U_j^{\text{DB}}, \quad U_j^{\text{DB}} \equiv \min_{q_j \geq 0} \mathbb{E}_{D_j, s_j} [c_j q_j s_j + c_j^{\text{DB}} [D_j - q_j s_j]^+ + h_j [D_j - q_j s_j]^-] \quad (2)$$

$$U^{\text{3DP}}(\mathcal{A}, K) \equiv \min_{\mathbf{q} \geq 0} \sum_{j \in \mathcal{A}} c_j q_j \mathbb{E}[s_j] + \mathbb{E}_{\mathbf{D}, \mathbf{s}} [V^{\text{3DP}}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, \mathcal{A}, K)] \quad (3)$$

where $\mathbf{q} \circ \mathbf{s}$ stands for the component-wise product between vectors $\mathbf{q}$ and $\mathbf{s}$, and

$$\begin{aligned} V^{\text{3DP}}(\mathbf{y}, \mathcal{A}, K) &\equiv \min_{\mathbf{q}^{\text{3DP}} \geq 0} \sum_{j \in \mathcal{A}} \left(c_j^{\text{3DP}} q_j^{\text{3DP}} + v_j [y_j - q_j^{\text{3DP}}]^+ + h_j [y_j - q_j^{\text{3DP}}]^- \right) \\ &\text{s.t.} \quad \sum_{j \in \mathcal{A}} q_j^{\text{3DP}} \leq K \end{aligned} \quad (4)$$

Before proceeding, we prove three simple properties of the optimal solution that we use in these derivations and throughout: First, because $c_j^{\text{DB}} < v_j$ and the expedited backup is unconstrained, it is always optimal to expedite all units of unmet demand for $j \in \mathcal{T}$, i.e., $q_j^{\text{DB}} = [D_j - q_j s_j]^+$. Second, since $c_j^{\text{DB}} < c_j^{\text{3DP}}$ and expedited backup is unconstrained, we always prefer to expedite unmet demand from a supplier $j \in \mathcal{T}$ over printing it. Hence, $\mathcal{A} \cap \mathcal{T} = \emptyset$. Finally, since we are always free to choose $q_j^{\text{3DP}} = 0$ for any $j \in \mathcal{A}$, we can without loss of generality take $\mathcal{A} = \mathcal{T}^c$. We summarize these last two observations in the following proposition:

LEMMA 1 (Backup Strategies Form a Partition). *Let $\mathcal{T}^*$ and $\mathcal{A}^*$ be optimal for (1). Then, $\mathcal{T}^* = \mathcal{A}^c$.*

Thus, in the remainder, we let $\mathcal{T} = \mathcal{A}^c$ and work with the set $\mathcal{A}$. Furthermore, when $\mathcal{A}$ and K are fixed and clear from context, we omit them from notation for brevity.

2.1. Dedicated Backups and $U^{\text{DB}}(\mathcal{A}^c)$

We next present the derivation of U^{DB} in Problem (2). By construction, dedicated backups decouple across products, i.e., $U^{\text{DB}}(\mathcal{A}^c) = \sum_{j \in \mathcal{A}^c} U_j^{\text{DB}}$, where U_j^{DB} includes primary ordering, expediting, and holding costs. Recall, under Assumption 1, any unmet demand $[D_j - q_j s_j]^+$ is fully covered by DB's unlimited capacity and there are no stock-outs. This gives rise to the newsvendor-type problem in Problem (2). We can solve this optimization explicitly:

LEMMA 2 (First-Stage Ordering for DB). *The optimal solution of Problem (2) with the smallest magnitude is given by:*

$$\bar{q}_j \equiv \inf \left\{ q \geq 0 : \mathbb{E}_{D_j, s_j} [s_j \mathbb{I}\{D_j \leq q s_j\}] \geq \left(\frac{c_j^{\text{DB}} - c_j}{c_j^{\text{DB}} + h_j} \right) \mathbb{E} s_j \right\}. \quad (5)$$

In special cases, $\bar{q}_j$ reduces to the usual newsvendor quantile, e.g., if D_j and s_j are independent. Otherwise, one must solve Eq. (5) numerically, e.g., by bisection search on q .

2.2. 3DP and $U^{\text{3DP}}(\mathcal{A}, K)$

We now present the derivation of U^{3DP} , c.f. Problem (3). Unlike dedicated strategies, the optimal 3DP backup quantities q^{3DP} lack a simple, closed-form due to the capacity constraint. Instead, U^{3DP} is defined by a two-stage stochastic optimization problem. The function V^{3DP} in Problem (4) represents the second-stage recourse cost that minimizes printing, holding, and stock-out costs with respect to the printing quantities q^{3DP} after the realization of q , s , and D . Using standard results from the stochastic programming literature, we can show that computing q_j for $j \in \mathcal{A}$ amounts to solving a convex optimization problem:

LEMMA 3 (Convexity of U^{3DP} and V^{3DP}). *The following holds for U^{3DP} and V^{3DP} :*

- (i) *The function $V^{3DP}(\mathbf{D} - \mathbf{q} \circ \mathbf{s})$ is convex in $\mathbf{q}$ for any given $(\mathbf{s}, \mathbf{D})$*
- (ii) *Problem (3) (which defines U^{3DP}) is convex.*
- (iii) *The function $K \mapsto U^{3DP}(\mathcal{A}, K)$ is convex in K for any $\mathcal{A}$.*

While many algorithms exist for solving convex problems, later in Section 5.4, we present a simple approach that leverages the specific structure of $V^{3DP}(\cdot)$ to efficiently solve Problem (3) and integrate it into a scalable heuristic framework for optimizing the first-stage decisions $\mathcal{A}$ and K .

2.3. Additional Model Discussion

We view the 3D printer as a flexible, make-to-order manufacturing resource capable of producing any product. In reality, some products, however, may not be printable (using current technology) because of their engineering specifications. We exclude such products from $\mathcal{N}$ because they do not affect the choice of $\mathcal{A}$, and computing their optimal order quantities and expedited shipping quantities can be done using simple newsvendor-like calculations, outside of the model, see Section 2.1.

We have made no assumptions on the dependence between D_j or s_j or across suppliers. Hence, we may without loss of generality assume a one-to-one correspondence between products and suppliers. Indeed, if a primary supplier supplies two products j and j' , we take s_j and $s_{j'}$ to be comonotonic.

Finally, while we have described dedicated backup in terms of expediting for ease of exposition, it can easily represent any dedicated strategy that entails a fixed cost and a cheaper per-unit cost, including employing dual sites or purchasing buffering inventory.

3. Properties of Optimal 3DP Strategies

In this section we present properties of the firm's optimal backup strategy by analyzing Problem (1). These properties provide concrete insights into how one should construct and operate a portfolio of resilience strategies that includes 3DP.

We first define the set $\mathcal{A}_0$ of unprotected products in the absence of 3DP, i.e., products that did not merit a dedicated backup. Unlike the optimal 3DP backup set $\mathcal{A}^*$ which may be difficult to compute, $\mathcal{A}_0$ is easily identified: Because dedicated backups decouple across products, a product $j \in \mathcal{N}$ is in $\mathcal{A}_0$ if and only if its costs when protected $U_j^{\text{DB}} + C_j^{\text{DB}}$ are larger than its costs when unprotected. These unprotected costs are given by solving Problem (2) after replacing c_j^{DB} by v_j , because there are no associated fixed costs for an unprotected product.

The next proposition relates $\mathcal{A}_0$ with the optimal 3DP backup set $\mathcal{A}^*$:

PROPOSITION 1 (Unprotected Products and Optimal Backup). *$\mathcal{A}_0 \subseteq \mathcal{A}^*$ and this containment can be strict.*

Thus, introducing 3DP backup impacts a firm's resilience practices in two ways: it covers all previously unprotected products *and* it *may* switch some products from dedicated to 3DP backup.

Proposition 1 has important implications for a firm piloting the use of 3DP as a backup strategy. Since $\mathcal{A}_0$ is easy to compute, such a firm might consider only backing up $\mathcal{A}_0$ instead of computing $\mathcal{A}^*$. While suboptimal, this strategy avoids the need to renegotiate terms with any existing expediting suppliers used by dedicated backup while the firm builds internal expertise around 3DP, and functions as a natural stepping stone towards backing up the entire set $\mathcal{A}^*$ later. In Section 6, we further explore the perspective of this firm piloting 3DP as a backup strategy and provide necessary and sufficient conditions for this suboptimal strategy to be profitable with even a small capacity investment. In the remainder of this section, we develop properties of the optimal solution for a generic $\mathcal{A}$, since we might be interested in $\mathcal{A}_0$ or $\mathcal{A}^*$, or some other set between these two.

For any choice of 3DP backups $\mathcal{A}$ and 3DP capacity K , we must identify corresponding first-stage order quantities $q_j^* = q_j^*(\mathcal{A}, K)$ for all $j \in \mathcal{A}$. Although computing these quantities exactly requires solving a two-stage convex optimization problem, Theorem 1 provides bounds that can easily be computed via simple bisection search.

THEOREM 1 (Bounds on Optimal First-Stage Orders). *For any set $\mathcal{A}$ and $K \geq 0$, let $q^*(\mathcal{A}, K)$ be the optimal solution to $U^{3DP}(\mathcal{A}, K)$ (c.f. Problem (3)) with minimal ℓ_1 -norm.¹ Then, for all $j \in \mathcal{A}$, we have $\bar{q}_j^\infty \leq q_j^*(\mathcal{A}, K) \leq \bar{q}_j^0$, where*

$$\bar{q}_j^0 = \arg \min_{q_j \geq 0} \mathbb{E} (c_j q_j s_j + v_j [D_j - q_j s_j]^+ + h_j [D_j - q_j s_j]^-). \quad (6)$$

$$\bar{q}_j^\infty = \arg \min_{q_j \geq 0} \mathbb{E} (c_j q_j s_j + c_j^{3DP} [D_j - q_j s_j]^+ + h_j [D_j - q_j s_j]^-). \quad (7)$$

When Problem (6) or Problem (7) has multiple optima, we tie break by taking the smallest solution.

Intuitively, $\bar{q}_j^0$ and $\bar{q}_j^\infty$ are the optimal first-stage orders for a system where $K = 0$ and a system where $K = \infty$, i.e., with no (resp. infinite) 3DP capacity. We stress, the bounds hold for each $j \in \mathcal{A}$. (We will use this property later when designing our algorithms.) Finally, we note that while computing $q_j^*(\mathcal{A}, K)$ requires knowledge of the full joint-distribution of (D_j, s_j) across products, computing the bounds above only requires the marginal distribution of each (D_j, s_j) pair. Calibrating such a distribution to data may be substantially easier in practice.

Our above bounds can also be used to estimate the optimal 3DP capacity investment. Specifically, let $K^*(\mathcal{A})$ be the minimizer of Problem (1) for a given $\mathcal{A}$ and $K^*(\mathcal{A}, \mathbf{q})$ be the optimal capacity

¹ Break ties arbitrarily.

for a given $\mathcal{A}$ and fixed first-stage order $q \geq 0$. In both cases, if the corresponding problem admits multiple optima, we define $K^*(\mathcal{A})$ or $K^*(\mathcal{A}, q)$ as the minimum optimal solution.

THEOREM 2 (Bounds on Optimal 3DP Capacity). *Let $C^{3DP}(K) = c^{\text{cap}} K$ for some $c^{\text{cap}} > 0$.*

- (i) *Suppose there exists an $r > 0$ such that $v_j - c_j^{3DP} = r$ for all $j \in \mathcal{A}$, then $K^*(\mathcal{A}, q)$ is the $(1 - \frac{c^{\text{cap}}}{r})$ -quantile of $\sum_{j \in \mathcal{A}} [D_j - q_j s_j]^+$, for any $q \geq 0$.*
- (ii) *Let $r^{\min} \equiv \min_{j \in \mathcal{A}} (v_j - c_j^{3DP})$ and $r^{\max} \equiv \max_{j \in \mathcal{A}} (v_j - c_j^{3DP})$. Then, $K^*(\mathcal{A})$ is bounded below by the $(1 - \frac{c^{\text{cap}}}{r^{\min}})$ -quantile of $\sum_{j \in \mathcal{A}} [D_j - \bar{q}_j^0 s_j]^+$ and bounded above by the $(1 - \frac{c^{\text{cap}}}{r^{\max}})$ -quantile of $\sum_{j \in \mathcal{A}} [D_j - \bar{q}_j^\infty s_j]^+$.*

When $v_j - c_j^{3DP}$ is equal across $j \in \mathcal{A}$, the first part of the theorem develops intuition around the size of the optimal $K^*(\mathcal{A}, q)$. Specifically, it is a quantile of the total demand shortfall and thus depends on the joint distribution of all (D_j, s_j) for $j \in \mathcal{A}$. The second part of the theorem builds on this intuition, leveraging our previous bounds on the first-stage order quantities and appropriately rounding $v_j - c_j^{3DP}$ to derive bounds on $K^*(\mathcal{A})$.

Finally, Theorem 2 also suggests settings where 3DP is likely to be beneficial. Indeed, if demand shortfalls are independent or anti-correlated, the quantile of the total shortfall tends to be small, and thus the optimal amount of 3DP capacity required will be small. By contrast, if shortfalls are highly correlated, a large amount of 3DP capacity may be needed, which is unlikely to be cost-effective. Leveraging results in the supermodular ordering of random variables, we extend this intuition in Theorem 3 which states that the cost of the 3DP strategy is highest when correlations among demand shortfalls are strongest. Recall, the random variables $X_1, \dots, X_n$ are comonotonic if and only there exists non-decreasing functions f_i for $i = 1, \dots, n$ and a random variable U such that $(X_1, \dots, X_n) \sim_d (f_1(U), \dots, f_n(U))$. Comonotonicity describes the strongest form of dependence between random variables.^d

THEOREM 3 (Comonotonic Shortfalls Are the Worst). *Suppose the marginal distributions of (D, s) are fixed. Then, the optimal cost of Problem (1) is maximal when the joint distribution of (D, s) is such that the demand shortfalls $[D_j - q_j s_j]^+$ are comonotonic across $j \in \mathcal{N}$ for all $q \geq 0$.*

Since dedicated backups decouple across backups by construction, the cost of a dedicated strategy depends only on the marginal distribution of each (D_j, s_j) pair, not their correlation structure. Hence, fixing marginal distributions, the cost of a dedicated backup is fixed. The theorem observes, however, that 3DP backups are sensitive to joint dependence, and describes the worst-case dependence. It provides an additional insight that well-chosen 3DP backups $\mathcal{A}$ will likely consist of suppliers that are anti-correlated or independent, not highly dependent.

4. Optimizing the 3DP Strategy via Mixed-Integer Optimization

In this section, we compute the optimal $\mathcal{A}^*$, $\mathbf{q}^*$ and K^* exactly, instead of bounding them, by reformulating (1) as a mixed-integer optimization (MIO) problem.

4.1. Mixed Integer Optimization Reformulation

For any $\mathcal{A} \subseteq \mathcal{N}$, we define a corresponding binary vector $\mathbf{x} \in \{0, 1\}^n$ such that $x_j = 1$ if $j \in \mathcal{A}$ and $x_j = 0$ otherwise. We can then rewrite U^{3DP} and V^{3DP} as functions in $\mathbf{x}$:

$$U^{3DP}(\mathbf{x}, K) \equiv \min_{\mathbf{q} \geq 0} \sum_{j \in \mathcal{N}} c_j q_j x_j \mathbb{E}[s_j] + \mathbb{E}_{\mathbf{D}, \mathbf{s}} [V^{3DP}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, \mathbf{x}, K)], \quad (8)$$

$$V^{3DP}(\mathbf{y}, \mathbf{x}, K) \equiv \min_{\mathbf{q}^{3DP} \geq 0} \sum_{j \in \mathcal{N}} \left(c_j^{3DP} q_j^{3DP} + v_j [y_j - q_j^{3DP}]^- + h_j [y_j - q_j^{3DP}]^+ \right) x_j \quad (9)$$

$$\text{s.t. } \sum_{j \in \mathcal{N}} q_j^{3DP} x_j \leq K.$$

This reformulation, although straightforward, introduces bilinear terms like $q_j x_j$ in the objective, which are notoriously difficult to handle. We prefer a “big M” type formulation in what follows, and develop bounds to ensure the “M” is not too large.

LEMMA 4 (“Big M” Formulation of V^{3DP}). Assume $0 \leq D_j \leq \bar{D}_j$ and $s_j^{\min} \leq s_j \leq s_j^{\max}$ almost surely for all $j \in \mathcal{N}$. Then for all $\mathbf{x} \in \{0, 1\}^n$, $K \geq 0$, and $\mathbf{y} = \mathbf{D} - \mathbf{q} \circ \mathbf{s}$:

$$V^{3DP}(\mathbf{y}, \mathbf{x}, K) = \min_{\mathbf{q}^{3DP}, \mathbf{z}^{3DP}} \sum_{j=1}^n \left(c_j^{3DP} q_j^{3DP} + v_j [y_j - z_j^{3DP}]^- + h_j [y_j - z_j^{3DP}]^+ \right) \quad (10a)$$

$$\text{s.t. } 0 \leq q_j^{3DP} \leq M_j^1 x_j, \quad \forall j = 1 \dots n, \quad (10b)$$

$$-M_j^2(1 - x_j) \leq z_j^{3DP} - q_j^{3DP} \leq M_j^2(1 - x_j), \quad \forall j = 1 \dots n, \quad (10c)$$

$$-M_j^3 x_j \leq y_j - z_j^{3DP} \leq M_j^3 x_j, \quad \forall j = 1 \dots n, \quad (10d)$$

$$\sum_{j=1}^n q_j^{3DP} \leq K, \quad (10e)$$

where for all $j \in \mathcal{N}$, we let

$$M_j^1 = \bar{D}_j - \bar{q}_j^\infty s_j^{\min}, \quad M_j^2 = \bar{D}_j, \quad M_j^3 = \max\{\bar{D}_j - \bar{q}_j^\infty s_j^{\min}, \bar{q}_j^0 s_j^{\max}\}. \quad (11)$$

Additionally, V^{3DP} is convex jointly in $(\mathbf{y}, \mathbf{x}, K)$.

It is well-known that the effectiveness of “big M” formulations like Problem (10) hinges on the choice of “M”, which in our case correspond to M_j^1 , M_j^2 , and M_j^3 . We have used the bounds from Theorem 1 to choose these values.

We can now use this reformulation of $V^{3DP}(\cdot)$ to reformulate Problem (1):

THEOREM 4 (Big “M” formulation for Optimizing $\mathcal{A}$). *Let the assumptions in Lemma 4 holds and let V^{3DP} be as reformulated in Lemma 4. Then, we can determine the optimal backup strategy $\mathcal{A}$, first-stage ordering quantities $\mathbf{q}$, and 3DP capacity K by solving the two-stage, stochastic MIO:*

$$\begin{aligned} \min_{\mathbf{x} \in \{0,1\}^n, K \geq 0, \mathbf{q}} \quad & C^{3DP}(K) + \sum_{j=1}^n (1 - x_j) U_j^{DB} + \sum_{j=1}^n c_j q_j \mathbb{E}[s_j] + \mathbb{E}_{\mathbf{D}, \mathbf{s}} [V^{3DP}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, \mathbf{x}, K)] \\ \text{s.t.} \quad & \bar{q}_j^\infty x_j \leq q_j \leq \bar{q}_j^0 x_j, \quad \forall j = 1 \dots n. \end{aligned} \quad (12)$$

Moreover, the objective function of Problem (12) is jointly convex in $\mathbf{x} \in \mathbb{R}^n$, $K \geq 0$ and $\mathbf{q} \geq 0$.

4.2. Solving Problem (12) at Scale

In principle, Problem (12) can be solved using an off-the-shelf solver (e.g., Gurobi) after approximating the expectation by S scenarios. However, this becomes impractical for large supply chains as the number of variables scales with nS , where n is the number of suppliers. While many variables are expected to be zero when $|\mathcal{A}|$ or K is small, generic branch-and-bound algorithms fail to exploit this sparsity or other structural properties of V^{3DP} .

An alternate approach frequently used in two-stage stochastic programs is constraint generation via Bender’s decomposition. The key idea is to replace the second-stage costs with a new auxiliary variable θ in the objective and introduce the convex, epigraphic constraint $\mathbb{E}[V^{3DP}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, \mathbf{x}, K)] \leq \theta$, where V^{3DP} is reformulated as in (10). The Benders approach iteratively refines the approximation of this constraint by introducing new valid cuts; see Birge and Louveaux (2011, Section 5.1) for an overview. Generating valid cuts in the Benders approach amounts to evaluating the subgradient of $V^{3DP}(\mathbf{y}, \mathbf{x}, K)$. This subgradient computation can be efficiently performed using a closed-form (c.f. Proposition 5 below). Details of Benders approach are deferred to Section B.1.

5. Approximations and Supermodularity Heuristics for Computing $\mathcal{A}$

Theorem 4 and Problem (12) describe an exact approach for computing the optimal policy $\mathcal{A}^*$. However, even with Benders decomposition, solving the MIP directly can be prohibitively expensive for a large supply chain. We next develop an approximation leveraging supermodularity and the intuition that, in an optimal solution, K is likely small relative to demand. We illustrate its computational effectiveness in Section 7 through our case-study with Mattel’s supply-chain (c.f. Fig. 3).

Recall, a set function $f : 2^{\mathcal{N}} \mapsto \mathbb{R}$ is supermodular if for all $S \subseteq T$ and $i \notin S$, it holds that:

$$f(T \cup \{i\}) - f(T) \geq f(S \cup \{i\}) - f(S) \quad (13)$$

Similarly, $f(\cdot)$ is submodular if $-f$ is supermodular. Intuitively, supermodularity captures increasing return to scale. There is a rich literature on optimizing submodular/supermodular functions. See, e.g., Krause and Golovin (2014). We would ideally leverage these techniques, but, as we next show, our problem is neither submodular nor supermodular.

5.1. Supermodularity of U^{3DP} : Positive and Negative Results

We start with a positive result, the objective function in the optimization defining $U^{3DP}(K, \mathcal{A})$ is supermodular:

PROPOSITION 2. *The function $\mathcal{A} \mapsto \sum_{j \in \mathcal{A}} c_j q_j \mathbb{E}[s_j] + \mathbb{E}[V^{3DP}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, \mathcal{A}, K)]$ is supermodular in $\mathcal{A} \subseteq \mathcal{N}$ for any fixed $\mathbf{q} \in \mathbb{R}^n$ and $K \geq 0$.*

Proposition 2 suggests that under an alternate simplified model, where first-stage order quantities are specified *exogenously* in a manner that does not depend on $\mathcal{A}$, one enjoys supermodularity. Unfortunately, in our model, first stage costs are chosen endogenously to optimize costs, which necessarily depends on the set $\mathcal{A}$. As a result, $U^{3DP}(K, \mathcal{A})$ is neither supermodular nor submodular.

PROPOSITION 3. *$U^{3DP}(K, \mathcal{A})$ is neither supermodular nor submodular in $\mathcal{A} \subseteq \mathcal{N}$.*

It seems then that optimizing the first stage costs partially drives the computational complexity of our model. This motivates seeking supermodular approximations.

5.2. A Supermodular Approximation of U^{3DP}

We next propose a supermodular approximation of U^{3DP} based on its Taylor expansion with respect to K at $K = 0$. The key intuition is that 3DP essentially serves as “flexibility” in the backup resource, and a host of operations literature suggests “a little flexibility is enough.” Moreover, given the costliness of printers, we intuit an optimal solution should have small K . Collectively, these observations suggest studying the Taylor series expansion of U^{3DP} for small K . The corresponding first order expansion turns out to be both supermodular *and* a lower bound:

THEOREM 5 (Supermodular Lower Bound for $U^{3DP}(K, \mathcal{A})$).

i) *Let*

$$L^{3DP}(\mathcal{A}, K) \equiv \sum_{j \in \mathcal{A}} U_j^0 - K \mathbb{E}_{\mathbf{D}, \mathbf{s}} \left[\max_{j \in \mathcal{A}} (v_j - c_j^{3DP}) \mathbb{I} \{D_j > \bar{q}_j^0 s_j\} \right]. \quad (14)$$

Then, for any $\mathcal{A}$, $L^{3DP}(\mathcal{A}, K) \leq U^{3DP}(\mathcal{A}, K)$.

ii) *For any $\mathcal{A}$, this bound becomes tight as K becomes small, i.e.,*

$$\lim_{K \downarrow 0} (U^{3DP}(\mathcal{A}, K) - L^{3DP}(\mathcal{A}, K)) = 0.$$

iii) For any fixed $K \geq 0$, the function $\mathcal{A} \mapsto L^{3DP}(\mathcal{A}, K)$ is supermodular.

Since $L^{3DP}(\mathcal{A}, K)$ is supermodular in $\mathcal{A}$, a natural heuristic might be to minimize this lower bound (in lieu of $U^{3DP}(\mathcal{A}, K)$), which should perform well if the optimal K is small. We pursue this heuristic in the next section, but first develop two weaker bounds that may be of interest in specialized settings:

PROPOSITION 4 (Simpler Supermodular Lower Bounds for U^{3DP}). *Let*

$$\widehat{L}^{3DP}(\mathcal{A}, K) \equiv \sum_{j \in \mathcal{A}} U_j^0 - K \left(\max_{j \in \mathcal{N}} v_j - c_j^{3DP} \right) \mathbb{P}(\exists j \in \mathcal{A} \text{ s.t. } D_j > \bar{q}_j^0 s_j). \quad (15)$$

Then,

- i) $L^{3DP}(\mathcal{A}, K) \geq \widehat{L}^{3DP}(\mathcal{A}, K)$.
- ii) For any fixed $K \geq 0$, the function $\mathcal{A} \mapsto \widehat{L}^{3DP}(\mathcal{A}, K)$ is supermodular.
- iii) $\widehat{L}^{3DP}(\mathcal{A}, K) = L^{3DP}(\mathcal{A}, K)$ whenever the per-unit margin $v_j - c_j^{3DP}$ is identical across $j \in \mathcal{N}$.

Similarly, for any fixed $\lambda > 0$, define

$$\widetilde{L}^{3DP}(\mathcal{A}, K, \lambda) \equiv \sum_{j \in \mathcal{A}} U_j^0 - \frac{K}{\lambda} \log \left(\sum_{j \in \mathcal{A}} \left[\mathbb{P}(D_j \leq \bar{q}_j^0 s_j) + \mathbb{P}(D_j > \bar{q}_j^0 s_j) e^{\lambda(v_j - c_j^{3DP})} \right] \right) \quad (16)$$

Then,

- iv) $L^{3DP}(\mathcal{A}, K) \geq \widetilde{L}^{3DP}(\mathcal{A}, K, \lambda)$ for all $\lambda > 0$.
- v) For any fixed $K \geq 0$, $\lambda \geq 0$, the function $\mathcal{A} \mapsto \widetilde{L}^{3DP}(\mathcal{A}, K, \lambda)$ is supermodular.

In Eq. (16), λ should be interpreted as a hyperparameter that is set exogenously.

In lieu of the expectation, $\widehat{L}^{3DP}(\mathcal{A}, K)$ depends on the probability that there is some unmet demand in $\mathcal{A}$ under the naive ordering strategy $\bar{q}^0$. Estimating this probability may be easier in practice than estimating the expectation in Eq. (14). On the other hand, evaluating this probability still requires the joint distribution of (D_j, s_j) across j which may be difficult to specify. The looser bound (16) avoids this and only depends on the distributions of the pair random variables (D_j, s_j) for each $j \in \mathcal{A}$. This might be substantially easier to estimate in practice. Thus, depending on data availability, either bound might be preferred to $\widehat{L}^{3DP}$.

As an aside, we note that replacing the maximum over $j \in \mathcal{N}$ in Eq. (15) with the maximum over $j \in \mathcal{A}$ also yields a tighter lower bound than $\widehat{L}^{3DP}(\mathcal{A}, K)$, but that bound is no longer supermodular.

5.3. Our Proposed Heuristic for Optimizing $\mathcal{A}$ and K

Our heuristic for optimizing $\mathcal{A}$ and K treats Problem (1) as a two-step optimization, and then replaces $U^{3DP}(\mathcal{A}, K)$ by $L^{3DP}(\mathcal{A}, K)$:

$$\min_{K \geq 0} C^{3DP}(K) + \min_{\mathcal{A} \subseteq \mathcal{N}} \left\{ \sum_{j \in \mathcal{A}^c} (C_j^{DB} + U_j^{DB}) + U^{3DP}(\mathcal{A}, K) \right\}, \quad (17)$$

$$\implies \min_{K \geq 0} C^{3DP}(K) + \min_{\mathcal{A} \subseteq \mathcal{N}} \left\{ \sum_{j \in \mathcal{A}^c} (C_j^{DB} + U_j^{DB}) + L^{3DP}(\mathcal{A}, K) \right\} \quad (18)$$

We solve the outer minimization in Problem (18) via grid search over K . For a fixed K , Theorem 5 shows that the inner problem entails minimizing a supermodular function, or, equivalently, maximizing a positive-valued, non monotone submodular function. (The non-monotonicity arises because adding products to the 3DP set may not reduce costs when K is fully utilized.) Feige et al. (2011) proves maximizing such functions is NP-Hard, but also shows that a simple local-search heuristic attains a $\frac{1}{3}$ approximation. Inspired by this result, we also use a local search to optimize the inner problem by iteratively adding or removing a product from $\mathcal{A}$. With clever bookkeeping, this can be done efficiently (c.f. Section B.2).

Once the local search terminates with a candidate $\hat{\mathcal{A}}(K)$, we seek to identify the best K from the grid. Lemma 3 shows that fixing K and $\hat{\mathcal{A}}(K)$, the optimization problem defining $U^{3DP}(\hat{\mathcal{A}}(K), K)$ (c.f. Problem (3)) is a convex stochastic optimization problem over $\mathbf{q}$, and hence, well-suited to a first-order optimization method. Thus, instead of comparing the objective values of the inner problem of Problem (18) to identify the best K , we run a first order method for each $(K, \hat{\mathcal{A}}(K))$ pair to evaluate the inner objective function of Problem (17) and compare these when choosing the best K . We summarize this heuristic in Algorithm 1. Experiments in Section 7 suggest Algorithm 1 already achieves near-optimal performance at a fraction of the MIO method's computational cost.

We next provide some details on how we implement projected stochastic gradient descent.

5.4. Evaluating $U^{3DP}(\mathcal{A}, K)$ via a Projected Stochastic Gradient Descent

Evaluating the inner objective function of Problem (17) amounts to computing $U^{3DP}(\mathcal{A}, K)$ by solving Problem (3) with $\mathcal{A} = \hat{\mathcal{A}}(K)$. Projected stochastic gradient descent is a natural approach to Problem (3), especially when $(\mathbf{D}, \mathbf{s})$ distributions are only accessible through historical data or simulations (see Bottou et al. (2018) for an overview). Additionally, as we will discuss, projected SGD allows us to leverage the specific structure of V^{3DP} to improve computational efficiency.

Specifically, the two key computational steps for any first-order method are:

Algorithm 1 SuperMod Approx: A Heuristic for Optimizing $\mathcal{A}$ and K **Require:** A grid $\mathcal{K}$ of K values

- 1: **for** each fixed $K \in \mathcal{K}$ **do**
- 2: Obtain $\hat{\mathcal{A}}(K)$ by solving the inner minimization of Problem (18) using local search (c.f. Section B.2).
- 3: Compute $U^{\text{3DP}}(\hat{\mathcal{A}}(K), K)$ via Projected Stochastic Gradient Descent.
- 4: Let $F(K)$ be the the objective value of the inner problem in of Problem (17).
- 5: **end for**
- 6: **return** $K^* = \arg \min_{K \in \mathcal{K}} F(K)$

- i) Computing an unbiased estimate of a subgradient (with respect to $\mathbf{q}$) of the objective of Problem (3) given a sample of $(\mathbf{D}, \mathbf{s})$.
- ii) Projecting an arbitrary point $\mathbf{q}$ to the positive orthant. This projection is given in closed-form by $\mathbf{q}^+$ (applied componentwise), see Boyd and Vandenberghe (2004).

Thus, it remains to compute a noisy, unbiased subgradient. To this end, we can reformulate $V^{\text{3DP}}(\mathbf{y}, K)$ (see Problem (4)) as a linear optimization problem and use duality theory to obtain:

PROPOSITION 5 (Subgradients of V^{3DP}). *Let $\bar{\mathbf{q}}^{\text{3DP}}$ denote the optimal solution to Problem (4) under fixed $\mathbf{y}$ and K . For $j \in \mathcal{A}$, a subgradient $\eta_j^{\mathbf{y}}$ of V^{3DP} with respect to y_j is given by:*

$$\eta_j^{\mathbf{y}} = \begin{cases} -h_j & \text{if } y_j \leq 0, \\ v_j & \text{if } y_j > 0 \text{ and } \bar{q}_j^{\text{3DP}} < y_j, \\ c_j^{\text{3DP}} - \eta^K & \text{otherwise.} \end{cases}$$

Here, η^K is a subgradient of V^{3DP} with respect to K , defined as

$$\eta^K = \begin{cases} -\max \{v_j - c_j^{\text{3DP}} \mid y_j > 0, \bar{q}_j^{\text{3DP}} < y_j\} & \text{if } \sum_{j \in \mathcal{A}} \bar{q}_j^{\text{3DP}} = K, \\ 0 & \text{otherwise.} \end{cases}$$

Finally, letting $\mathbf{y} = \mathbf{D} - \mathbf{q} \circ \mathbf{s}$ above, we have that the vector $(\mathbf{c} - \boldsymbol{\eta}^{\mathbf{y}}) \circ \mathbf{s}$ is an unbiased estimate of the subgradient of the objective in Problem (3) with respect to $\mathbf{q}$.

In words, the noisy subgradient can be efficiently computed in closed form if we can quickly identify an optimizer $\bar{\mathbf{q}}^{\text{3DP}}$ for $V^{\text{3DP}}(\mathbf{y}, K)$ given any $\mathbf{y}$ and K . Fortunately, this reduces to solving a fractional-knapsack problem, as shown next:

PROPOSITION 6 (Solving Problem (4)). *Fix some $\mathbf{y}$ and $K \geq 0$. Then, any solution $\bar{\mathbf{q}}^{3DP}$ to (4) also solves*

$$\begin{aligned} \bar{q}_j^{3DP} \in \arg \max_{0 \leq \mathbf{q}^{3DP} \leq \mathbf{y}^+} \sum_{j \in \mathcal{A}} (v_j - c_j^{3DP}) q_j^{3DP} \\ \text{s.t.} \quad \sum_{j \in \mathcal{A}} q_j^{3DP} \leq K. \end{aligned} \quad (19)$$

Problem (19) can be solved very efficiently by sorting $v_j - c_j^{3DP}$ (c.f. Algorithm 2). Combined with Proposition 5, we can thus construct unbiased estimates of the subgradient very efficiently, ensuring projected SGD remains highly scalable in Algorithm 1. See Section B.3 for additional implementation details.

6. Piloting 3DP as a Backup Strategy

In this section, we consider the perspective a firm not yet engaged in 3DP that is piloting its use as a backup strategy. Such a firm is likely unwilling to alter its existing dedicated backup strategies, but likely *would* consider backing up currently unprotected products $\mathcal{A}_0$. Thus, we focus on understanding the benefits of backing up only $\mathcal{A}_0$. As observed already in Proposition 1, $\mathcal{A}_0 \subseteq \mathcal{A}^*$.

Specifically, we ask when there exists a 3DP capacity investment $K > 0$ with positive cost savings, i.e., a K such that

$$U^{\text{tot-saved}}(\mathcal{A}_0, K) \equiv \sum_{j \in \mathcal{A}_0} U_j^0 - U^{3DP}(\mathcal{A}_0, K) - C^{3DP}(K) > 0. \quad (20)$$

Here, $U^{\text{tot-saved}}(\mathcal{A}_0, K)$ is the total cost we could save on $\mathcal{A}_0$, where U_j^0 is the operational cost of product j when unprotected. Since $C^{3DP}(K)$ is convex and increasing and Lemma 3 shows $K \mapsto U^{3DP}(\mathcal{A}_0, K)$ is convex, we have that $U^{\text{tot-saved}}$ is concave in K .

Moreover, since $C^{3DP}(0) = 0$, we have $U^{\text{tot-saved}}(\mathcal{A}_0, 0) = 0$, i.e., for $j \in \mathcal{A}_0$, investing in no 3DP capacity is tantamount to leaving j unprotected. Thus, as K increases, $U^{\text{tot-saved}}(\mathcal{A}_0, K)$ either decreases monotonically from 0 or initially increases before decreasing. In the second case, there is a positive cost saving for any small enough value of K . We can distinguish these two cases by examining the derivative of $U^{\text{tot-saved}}$ at $K = 0$.

THEOREM 6 (Necessary and Sufficient Conditions for 3DP Cost Savings under $\mathcal{A}_0$). *Let C^{3DP} be increasing, convex, differentiable at 0 and $\bar{q}_j^0$ as defined in (6). Then there exists $K > 0$ such that $U^{\text{tot-saved}}(\mathcal{A}_0, K) > 0$ if and only if*

$$0 < \mathbb{E}_{\mathbf{D}, \mathbf{s}} \left[\max_{j \in \mathcal{A}_0} (v_j - c_j^{3DP}) \mathbb{I} \left\{ [D_j - \bar{q}_j^0 s_j]^+ > 0 \right\} \right] - c^{\text{cap}}. \quad (21)$$

If such K exists, then $U^{\text{tot-saved}}(\mathcal{A}_0, K') > 0$ for all $K' \in [0, K]$.

Notice that Eq. (21) as written depends on the joint distribution of all pairs (D_j, s_j) . This aligns with the insight of Theorem 3, that 3DP costs savings depends on the correlations between suppliers.

The conclusion that $U^{\text{tot-saved}}(\mathcal{A}_0, K') > 0$ for all sufficiently small K' is crucial; it indicates that for a firm piloting 3DP, if Eq. (21) holds, then *any* small investment in 3DP capacity yields cost savings. This provides a natural way for a firm to scale-up the use of 3DP in their operations as they build internal expertise and experience.

7. Empirical Case Study

In this section, we conduct a case study inspired by Mattel, a major toy company. Mattel's products are primarily plastic with a density similar to common 3DP materials like resin and PLA, making them ideal candidates for printing. To better align our model with the case study, we make one technical adjustment; we replace the capacity constraint in Problem (4) with a "weighted" analogue: $\sum_{j \in \mathcal{A}} w_j q_j^{3\text{DP}} \leq K$. Here w_j is the material weight (in grams) required to print one unit of product j . This technical change requires minimal adjustments to our algorithms, but reflects the fact that different products may require different amounts of time to print (based on their weight). In the same spirit, we model the sourcing cost from the primary supplier as a proportional to the weight of a product, specifically at \$0.006 per gram, incorporating both the raw plastic cost and a 20% markup.


Finally, as is standard in the literature, we interpret our objective in Problem (1) as regret relative to an oracle seller who knew demand rather than realized costs. Hence, the back-order cost v_j is the lost profit on a unit not sold, h_j is unit cost of sourcing from the primary supplier, $c^{3\text{DP}}$ (resp. c^{DB}) is the additional cost printing (resp. expediting) relative to sourcing from the primary supplier, and $c_j = 0$. We assume a zero salvage value for all products throughout.

7.1. Calibration to Real-World Data

Since detailed supply chain data on Mattel's operations is not available, we calibrate our model to several sources of public data, including i) engineering specifications of state of the art printers, ii) scraped data from Mattel's website and the Mattel Store on Amazon.com, and iii) bill-of-lading (BOL) data on imports. We summarize the overall calibration process in Fig. 1, and describe in detail below. Additional details available in Section C.

7.1.1. 3DP Costs. We model unit 3DP production cost as twice the unit production cost from the primary supplier, and provide sensitivity analysis to this parameter below, letting it range from 1 to 4 times unit production cost at the primary supplier. We choose the baseline factor of "2" by considering the per gram cost of non-metal printing materials like resin and PLA, and an infill density of 50%.

Figure 1 Overview of our calibration procedure.



Note. We use e-commerce data from the Mattel Store on Amazon.com and Mattel.com to learn major product categories and identify products in those categories, their weights, and sales prices. We then use natural language processing on bill-of-lading (BOL) data from suppliers importing to Mattel to construct synthetic products – one per supplier – that we use as the primitives in our model. Finally, we calibrate printing costs to publicly available technical specifications and quotes from Sintratec (2024).

We also assume that cost of printing capacity is linear, i.e., $C^{3DP}(K) = c^{cap}K$, and let $c^{cap} = \frac{Q^{3DP}}{M^{3DP}}$, where Q^{3DP} and M^{3DP} represent the monthly per printer depreciation cost (in \$) and the monthly material output per printer (in grams). At present, industrial 3DPs typically range from \$5,000 to \$50,000, with higher-end models offering larger build volumes and advanced material capabilities rather than superior speed. For Mattel, which prints small plastic toys, a \$5,000 printer provides sufficient speed and material compatibility without excess cost. Thus, we set \$5,000 per printer as our baseline which, assuming a 10-year lifespan, yields $Q^{3DP} = \$41.67$ per month (All3DP 2024, Fusion3 Design 2024). We approximate the monthly material output M^{3DP} using published technical specifications of common printers (nozzle width, layer thickness, nozzle movement speed, infill density, and working hours in a month.) See Section C for details. Overall, this yields an estimate of $c^{cap} = \$0.0023$ per gram. Again, we provide sensitivity to this value below.

7.1.2. Learning Product Categories. As will be seen, we heavily rely on BOL data to infer a supply chain that approximates Mattel’s chain. Unfortunately, BOL data is incomplete in some respects. This necessitates some approximations based on product categories. Hence, as a first step, we discuss learning these categories.

We scrape public websites including Mattel.com and the Mattel Store on Amazon.com to retrieve a list of Mattel’s product categories (dolls, dinosaurs, cars and trucks, etc.) (Prices and descriptions in both settings are determined by Mattel, not by third party sellers.) Overall, we find seven, and within each category, we identify a constituent list of products. For each product, we

identify the products weight (in grams), sales price, and a text description. For each category ℓ , we then calculate the average unit weight w_{ℓ}^{ave} and average sales price $\kappa_{\ell}^{\text{ave}}$ within the category. Finally, using TF-IDF scoring, we generate a keyword list from the product descriptions meant to describe the category. We will use these average weights and keyword list in our next step of generating synthetic products.

Bill of Lading	Carrier SCAC / Vessel Code (Name)	Voyage / Container Size / Type
XXXX333333333333	Eglv, Eagle Van Lines Inc / 3333333 (Ever Envoy)	3333x / 3000*333*333 / 33X3
Destination Port Name / Code	Company Name / Address	Arrival Date
Los Angeles / 2704	Mattel Import Services Corp / 333 Continental Blvd, El Segundo, CA, 90245, USA	XX/XX/20XX
Departure Port Name / Code	Supplier Name / Address	
Yantian / 57078	Shenzhen Hutchison Inland / Container Depots Co Ltd Xintian Hutchison Warehouse Distripark Huanguan Rd Sth GI Town Baoan Disct Sz O/B	
HS Code / Full Description	Product Description	Quantity (Unit) / TEU / Weight
950611 / Toys, games and sports requisites; parts and accessories ...	American Girl Dolls Plastic Toys Dinotrux Diecast Hot Wheels Heavy Etc Market Mdm	333 (CTN) / 33 / 3333 kilogram

Table 1 Sample bill of lading. One such report is submitted to the Federal Customs and Border Protection Agency for every import by a US Firm. As highlighted, these forms provide partial information about Tier 1 suppliers, types of goods imported, and, demand for that supplier.

7.1.3. Identify Suppliers and Constructing Synthetic Products from Bill of Lading Data.

A bill of lading (BOL) is a report submitted by U.S. firms to the Federal Customs and Border Protection Agency, detailing shippers, consignees, goods, quantities, weights, and other specifics of every international maritime shipment (see Table 1 for a portion of a sample report). We obtained BOLs from ImportYeti.com, a third-party service that compiles BOL data across companies, including Mattel, for over the past decade. In what follows, we limit attention to the 55 suppliers with over 100 maritime shipments to Mattel in the past decade and neglect smaller, “one-off” suppliers. Suppliers are identified by the “Supplier Name” field in the BOLs.

In principle, BOLs might provide sufficient details to identify Mattel’s primary suppliers and their respective demands, but, as can be seen from the example, the reports are often vague in places:

1. The firm typically sources multiple products from each supplier, but BOLs provide only generic descriptions like “plastic toys”, with no specific product names or their respective quantities.
2. For each shipment, only the total weight is provided, while total unit counts are either missing or recorded in broad terms (e.g., “CTN” for cartons).


Hence, we will create a single “synthetic” product for each supplier and calibrate this synthetic product to the observed BOL data. Implicitly, this construction assumes that if a supplier is disrupted,

all of its products are disrupted (proportionately), and that products provided by different suppliers are not substitutable.

More specifically, we leverage our learned keywords for product categories. For each supplier j , let b_j^ℓ denote the proportion of bills containing at least one keyword from category ℓ . These values approximate the proportion of different product categories in a typical shipment from supplier j . Finally, we construct a synthetic product for supplier j with a per-unit weight of equal to $w_j = \sum_\ell b_j^\ell w_\ell^{\text{ave}}$ and a per-unit sales price of $\kappa_j^{\text{sell}} = \sum_\ell b_j^\ell \kappa_\ell^{\text{ave}}$. To compute a demand distribution for this synthetic product, we compute the total weight of a shipment each month and convert them to unit counts by dividing by w_j . To simplify our analysis, we then bin this discrete distribution using K -means clustering into a three point distribution for each product. Demand is assumed independent across suppliers, so with 55 suppliers, there are already $3^{55} > 10^{26}$ possible demand realizations.

Although somewhat involved, we do believe this procedure yields a reasonable approximation to the supply chain for a company like Mattel. Figure 2 gives an overview of the demand and price of these synthetic products.

Figure 2 Synthetic Products Calibrated to BOL Data.




Note. Our calibration procedure results in 55 synthetic products, one per supplier. Panel (a) shows that there is a correlation between product price and demand (maker size proportional to per-unit weight), with many low-volume, low-cost products. Panel (b) shows a fair amount of heterogeneity in weight (and hence cost) among products. Panel (c) presents aggregate demand for our firm, which exhibits significant variability.

To complete our calibration, we require yield data for each of these 55 suppliers. Yield data are rarely shared externally due to their proprietary nature. Thus, we model random yield explicitly as $s_j \in \{\alpha_j, 1\}$, where $1 - \alpha_j \in (0, 1)$ and $p_j = \mathbb{P}(s_j = 1 - \alpha_j)$ are assumed known. By default, $\alpha_j = 0.05$, $p_j = 0.05$ and s_j 's are independent, unless specified otherwise.

7.2. How Accurate is the SuperMod Approx Heuristic?

As a first step, we study how suboptimal solutions obtained from the SuperMod Approx method (c.f. Algorithm 1) are relative to the full-information optimum obtained by solving the mixed-binary optimization problem of Theorem 4.

Figure 3 Comparison of supermodular approximation and MIO for optimizing $\mathcal{A}$ 

As mentioned, solving this mixed-binary problem is challenging at scale, so in Fig. 3, we study the scalability of SuperMod Approx as n grows. For benchmarking, we compare to the solution time to solving the mixed-binary problem in Theorem 4 using Bender's cuts. Unfortunately, as n grows, the number of scenarios in each of the expectations grows exponentially fast. Hence, we replace these expectations by sample average approximations computed using 5% of the possible scenarios. We call this heuristic MIO-SAA and cap its run time at 598 seconds, matching the total computation time of SuperMod Approx across all cases for $n = 1 \dots 55$.

From Fig. 3a, at $n = 10$, the MIO-SAA approach already hits the time limit, whereas the SuperMod Approx scales very gradually with increasing suppliers. This highlights its usefulness for large, disaggregated chains.

Furthermore, Fig. 3b attempts to compare the suboptimality gap of SuperMod Approx and MIO-SAA for these larger chains. Since the full-information optimum is unavailable, we present the gap to best lower bound computed in the course of the Bender's algorithm. (This is the traditional stopping criteria for Benders.) One can see that across n , SuperMod Approx finds a solution which is no worse than the MIO-SAA approach at a much smaller computational cost.


7.3. How much value does 3DP offer?

Of course, a key question is if 3DP provides any cost savings over traditional back up strategies. In Panels (a) and (b) of Fig. 4 we plot the cost savings relative to a system without 3DP as a function of the amount of capacity K purchased. As our theory suggests, the optimal K is a relatively small fraction of the maximum demand, but the savings are modest. For our baseline parameters, at the optimal K , we see a savings of just under 4%, but, at larger costs of capacity or costs of printing, this benefit quickly disappears.

Perhaps more striking are panels (c) and (d) of Fig. 4. Here we plot the distribution of the demand shortfall, i.e., the total unmet demand after utilizing the received primary orders and backup productions. Note, this quantity is random depending on the random realizations of D_j and s_j for all j .

We compute these boxplots using 10^5 random samples of the vector $((D_j, s_j) : j = 1, \dots, 55)$. Here we see that although cost improvements are modest, 3DP effectively reduces the amount of unmet demand, particularly in tail scenarios. This benefit persists even for large costs of capacity or printing. In many ways we see this as the primary argument for 3DP as a resilience strategy; although average benefits are small, the additional flexibility helps guard against large vales of unmet demand.

Figure 4 Cost savings and demand shortfalls: varying c^{cap} and c_j^{3DP}




Note. Panels (a) and (b): Cost savings (%) relative to a system with dedicated backup but no 3DP. Savings are modest, and the optimal K is a relatively small fraction of demand. Panels (c) and (d) display box plots of demand shortfalls. Here we see a more striking reduction in the tails of the distribution. The x-axes indicate multiples of the baseline c^{cap} and c_j^{3DP} . Box plots labeled “No 3DP” correspond to the system with dedicated backups only and no 3DP.

Fig. 5 ablates this improvement by comparing the optimal strategy to the “piloting” strategy of only backing up $\mathcal{A}_0$ with 3DP. We sample 100 subsets of size $n \in \{15, 30, 45\}$ and also include the full supply chain ($n = 55$) without sampling. The left panel shows the percentage of suppliers backed up by dedicated resources under: i) “No 3DP”: the traditional system with no 3DP ii) “Piloting 3DP”: the traditional system where we additionally back up all unprotected products by 3DP iii) “Full 3DP”: the optimal policy. We see that at our baseline values, a substantive number of suppliers switch from dedicated backups to the flexible 3DP backup. In the middle panel, we see that a little over half of the cost savings comes from backing up the unprotected items. Finally, the right panel shows that nearly all reductions in average unmet demand result from the piloting 3DP strategy.


Thus, the piloting 3DP strategy – i.e. covering unprotected products – already seems to capture the principal benefits of reducing shortfalls, and captures most (but not all) of the cost benefits.

7.4. What kinds of products are the best candidates for 3DP backup?

As discussed, in our model, any product that is unprotected when only considering dedicated backups (i.e. $\mathcal{A}_0$) should be backed up by 3DP if possible. What is less clear is which products previously backed up by a dedicated resource should switch to a 3DP backup in an optimal policy $\mathcal{A}^*$. To better

Figure 5 Backup switch after introducing 3DP.

Note. From left to right: percentage of n assigned to dedicated backups, cost savings (relative to the cost without 3DP) and mean demand shortfall (relative to max demand) for $n \in \{15, 30, 45, 55\}$. We compare the dedicated backup-only strategy (blue), the piloting 3DP strategy (red, backing up only unprotected products), and the full 3DP strategy (yellow, potentially switching dedicated backups).

Figure 6 Key factors driving product switching.

Note. We fit a Decision tree to identify key drivers that govern which products will switch from dedicated to 3DP backup at the optimal capacity. The most likely products to switch have a moderate demand shortfall and a favorable 3DP Profit-to-Cost ratio.

understand the characteristics of products that switch, we sample 1,000 sets of $n = 10$ products, and simulate heterogeneity in the s distributions by randomly assigning $\alpha_j \in (0, 0.75]$ and $p_j \in (0, 0.5]$. We also randomly assign the ratio c_j^{DB}/c_j^{3DP} within $(0, 1)$. We then use a CART decision tree to predict if a product will switch based on several, normalized features:

$$\begin{aligned} \text{Mean Shortfall} &\equiv \frac{\mathbb{E}[(D_j - \bar{q}_j^0 s_j)^+]}{\mathbb{E}[D_j]}, & \text{3DP Profit-to-Cost} &\equiv \frac{(v_j - c_j^{3DP})/w_j}{c^{\text{cap}}} \\ \text{DB Critical Quantile} &\equiv \frac{c_j^{DB} - c_j}{c_j^{DB} + h_j}, & \text{Primary Critical Quantile} &\equiv \frac{v_j - c_j}{v_j + h_j}. \end{aligned} \quad (22)$$

where $\bar{q}_j^0$ is the optimal first-stage order without any backup (see (6)), and w_j is the unit weight. In essence, “Mean Shortfall” quantifies the unmet demand that backup needs to cover, while “3DP Profit-to-Cost” captures the profitability of a product (when 3D-printed) relative to 3DP unit capacity cost. The terms “DB-” and “Primary Critical Quantile” represent the critical quantiles in the

newsvendor calculation for a product when it is backed up by the dedicated strategy and when it is unprotected, respectively. All hyperparameters are tuned with 5-fold cross-validation.

The resulting decision tree is shown in Fig. 6. In each leaf node the first row is the majority label, the second row is the proportion of products in this leaf that “switch” from dedicated to flexible backup, and the third row shows the leaf’s proportion of the total data.


Figure 6 suggests that products with very high demand shortfall (≥ 0.34) or very low shortfall (≤ 0.031) do not switch; these products should either use a dedicated backup or none at all. But products with moderate shortfalls and a good 3DP Profit-to-Cost ratio are more likely to switch. While clearly only a heuristic model, this tree nonetheless provides some managerial insight into which products to consider switching.

7.5. Correlated Disruptions

Theorem 3 shows that 3DP will perform worst (relative to dedicated backups) when demand shortfalls are comonotonic. In this section, we further explore this phenomenon by considering a sequence of models where disruptions become more correlated and assess the drop in value of 3DP.

Specifically, we introduce a latent factor $X_0 \sim \text{Bernoulli}(p_0)$ with $p_0 \leq 0.05$ representing a global failure that affects all suppliers. We then take $X_j \sim \text{Bernoulli}(\frac{p-p_0}{1-p_0})$ independently. Finally, we let $s_j = 1 - \alpha \max(X_j, X_0)$. Thus, each supplier fails with probability $\mathbb{P}(s_j = 1 - \alpha) = p = 0.05$. However, by adjusting the ratio p_0/p from 0 to 1, we can control the correlation between failures via the common factor. Specifically, when $p_0/p = 0$, the s_j are independent. When $p_0/p = 1$, the s_j are comonotonic. In what follows, we also consider two values of $\alpha \in \{.05, 1\}$ to contrast the case of small and large disruptions.

Figure 7 Cost savings and demand shortfalls: varying correlations among disruptions



Note. Panel (a): Cost savings (relative to no 3DP) as p_0/p increases from 0 (independent s) to 1 (comonotonic s). Panel (b): Demand shortfall (relative to total demand) vs. p_0/p . Both panels compare 5% (“Small”) and 100% (“Large”) yield loss disruption regimes.

From the blue curve in Fig. 7a, for small disruption levels, as correlation increases, cost savings decrease by only 0.2%. This decrease is larger for the large disruption case, but even at comonotonic failures, 3DP offers some cost savings. From the left panel in Fig. 7b, 3DP consistently reduces demand shortfalls across all correlation levels considered, and for small disruptions, the magnitude of this reduction is relatively stable. For larger disruptions, however, the benefit from reducing shortfalls wanes as correlation increases. At very large correlations, if many suppliers fail (entirely) and simultaneously, the shortfalls significantly exceed the 3DP capacity, and the no 3DP system may perform better on average.

Altogether, these observations provide evidence that stronger disruption correlations weaken the 3DP resilience strategy's effectiveness in cost savings and demand shortfall reduction, particularly when disruption-induced yield losses are large, but that there still may be some value for high-values of correlation.

8. Conclusion

In this paper we explore the possibility of using 3DP a flexible backup resource as part of a larger portfolio of supply chain resilience strategies. Doing so requires solving a host of operational problems around which products to backup with 3DP, how much 3DP capacity to acquire, and how to structure both primary orders from suppliers and recourse printing actions. To that end, we formulate a mixed-binary stochastic optimization problem describing the setting, and derive several properties of its optimal solution that help assess the potential value and costs of adopting a 3DP strategy. We also provide a scalable heuristic for solving this problem based on constructing a supermodular approximation. Through an empirical case study inspired by Mattel, we establish that with current technology, 3DP offers only a modest savings over dedicated backup (3-4%), but that it offers a much more significant reduction in the amount of unmet demand in the system. For risk-sensitive firms or settings where qualitative branding risks are serious, such benefits might justify the investment.

There are a number of interesting direction for extension. We have considered a single firm and its Tier 1 suppliers. One could study other supply chain networks, such as an assembly network, and ask where topologically on the network 3DP might be most valuable. Moreover, we have taken a deliberately conservative viewpoint, neglecting potential benefits of 3DP from co-location with demand or faster lead times. Modeling these features appropriately might reveal additional benefits.

Overall, 3DP and additive manufacturing have long been considered inviable for large-scale manufacturing and supply chain, and relegated to low-volume applications like prototyping. While there are very real practical challenges to widespread 3DP adoption including quality assurance, upskilling

a workforce, and navigating regulatory requirements, there are also very real advantages of incorporating flexible back-up strategies in a firm's supply-chain resilience plan. We hope our work inspires academics to think more broadly about the potential of this technology.

9. Data Availability and Reproducibility

All code for reproducing our experiments can be found at https://github.com/ziyuhe/3DP_resilience_experiments/. In the spirit of reproducibility of research, we also provide the complete dataset of 55 suppliers with details on their synthetic products within this repository.

Data obtained from [ImportYeti.com](https://importyeti.com) is proprietary and available for purchase. We cannot share this data directly. Interested researchers should contact this firm.

References

- All3DP (2024) Best professional 3D printers for small businesses. URL <https://all3Dp.com/1/best-professional-3D-printer-small-business/>.
- Arbaban ME, Wagner SM (2020) The impact of 3D printing on manufacturer–retailer supply chains. *European Journal of Operational Research* 285(2):538–552.
- Bassamboo A, Randhawa RS, Van Mieghem JA (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.
- Bertsekas DP, Nedic A, Ozdaglar AE (2003) *Convex Analysis and Optimization*, volume 1 (Athena Scientific).
- Birge JR, Louveaux F (2011) *Introduction to Stochastic Programming* (New York, NY: Springer Science & Business Media).
- Bottou L, Curtis FE, Nocedal J (2018) Optimization methods for large-scale machine learning. *SIAM Review* 60(2):223–311.
- Boyd S, Vandenberghe L (2004) *Convex Optimization* (Cambridge University Press).
- Chen L, Cui Y, Lee HL (2021) Retailing with 3D printing. *Production and Operations Management* 30(7):1986–2007.
- Dada M, Petrucci NC, Schwarz LB (2007) A newsvendor's procurement problem when suppliers are unreliable. *Manufacturing & Service Operations Management* 9:9–32.
- Demirel S, Kapuscinski R, Yu M (2017) Strategic behavior of suppliers in the face of production disruptions. *Management Science* 64:533–551.
- Dong L, Shi D, Zhang F (2022) 3D printing and product assortment strategy. *Management Science* 68(8):5724–5744.
- Feige U, Mirrokni VS, Vondrák J (2011) Maximizing non-monotone submodular functions. *SIAM Journal on Computing* 40(4):1133–1153.
- Fusion3 Design (2024) How much does a 3D printer cost? URL <https://www.fusion3Design.com/how-much-does-a-3D-printer-cost/>.

- Goovaerts MJ, Dhaene J (1999) Supermodular ordering and stochastic annuities. *Insurance: Mathematics and Economics* 24(3):281–290.
- Jordan WC, Graves SC (1995) Principles on the benefits of manufacturing process flexibility. *Management Science* 41(4):577–594.
- Krause A, Golovin D (2014) Submodular function maximization. *Tractability*, volume 3, 71–104 (Cambridge University Press).
- Lopez J (2020) General motors opens dedicated 3D printing center. URL <https://gmauthority.com/blog/2020/12/general-motors-opens-dedicated-3D-printing-center/>, accessed: 2025-01-21.
- Materialise (2023) How cnh industrial incorporated 3D printing into their supply chain strategy. URL <https://www.materialise.com/en/inspiration/cases/cnh-additive-manufacturing-supply-chain-strategy>.
- McEachern S (2022) Gm 3D printed 60,000 parts to keep 2022 chevy tahoe production rolling. URL <https://gmauthority.com/blog/2022/06/gm-3d-printed-60000-parts-to-keep-2022-chevy-tahoe-production-rolling/>, accessed: 2025-01-21.
- McKinsey & Company (2021) Reimagining industrial supply chains. URL <https://www.mckinsey.com/industries/industrials-and-electronics/our-insights/reimagining-industrial-supply-chains>.
- McKinsey & Company (2023) Taking the pulse of shifting supply chains. URL <https://www.mckinsey.com/capabilities/operations/our-insights/taking-the-pulse-of-shifting-supply-chains>.
- Saghafian S, Van Oyen MP (2012) The value of flexible backup suppliers and disruption risk information: Newsvendor analysis with recourse. *IIE Transactions* 44(10):834–867.
- Saghafian S, Van Oyen MP (2016) Compensating for dynamic supply disruptions: Backup flexibility design. *Operations Research* 64(2):390–405.
- Sethuraman N, Parlaktürk AK, Swaminathan JM (2023) Personal fabrication as an operational strategy: Value of delegating production to customer using 3D printing. *Production and Operations Management* 32(7):2362–2375.
- Shapiro A, Dentcheva D, Ruszczyński A (2021) *Lectures on Stochastic Programming: Modeling and Theory* (Society for Industrial and Applied Mathematics).
- Simchi-Levi D, Schmidt W, Wei Y, Zhang PY, Combs K, Ge Y, Gusikhin O, Sanders M, Zhang D (2015a) Identifying risks and mitigating disruptions in the automotive supply chain. *Interfaces* 45(5):375–390.
- Simchi-Levi D, Schmidt W, Wei YD, Zhang PY, Combs K, Ge Y, Gusikhin O, Sanders M, Zhang D (2015b) Identifying risks and mitigating disruptions in the automotive supply chain. *Interfaces* 45:375–390.

- Simchi-Levi D, Wang H, Wei YD (2017) Increasing supply chain robustness through process flexibility and inventory. *Production and Operations Management* 27:1476 – 1491.
- Simchi-Levi D, Wei Y (2012) Understanding the performance of the long chain and sparse designs in process flexibility. *Operations Research* 60(5):1125–1141.
- Sintratec (2024) Sintratec kit: Sls 3D printer. URL <https://sintratec.com/sls-3d-printer/sintratec-kit/>.
- Song JS, Zhang Y (2020) Stock or print? impact of 3-d printing on spare parts logistics. *Management Science* 66(9):3860–3878.
- Stevens T (2022) Gm 3D prints 60,000 parts to keep tahoe deliveries on time. URL <https://www.cnet.com/roadshow/news/chevy-tahoe-3d-parts/>, accessed: 2025-01-21.
- Tomlin B (2006) On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science* 52:639–657.
- Topkis DM (1998) *Supermodularity and Complementarity* (Princeton, NJ: Princeton University Press).
- Tsitsiklis JN, Xu K (2013) On the power of (even a little) resource pooling. *Stochastic Systems* 2(1):1–66.
- Wang Y, Webster S (2022) Product flexibility strategy under supply and demand risk. *Manufacturing & Service Operations Management* 24(3):1779–1795.
- Westerweel B, Basten R, den Boer J, van Houtum GJ (2021) Printing spare parts at remote locations: Fulfilling the promise of additive manufacturing. *Production and Operations Management* 30(6):1615–1632.
- Yang ZB, Aydin G, Babich V, Beil DR (2009) Supply disruptions, asymmetric information, and a backup production option. *Management Science* 55:192–209.
- Zhang Y, Westerweel B, Basten R, Song JS (2022) Distributed 3D printing of spare parts via ip licensing. *Manufacturing & Service Operations Management* 24(5):2685–2702.

Appendix A: Omitted Proofs

In this section, we present the proofs of the key results discussed in this paper.

A.1. Background Results

First, we present several background results that will be used later in the proof of our main results:

1. We first reformulate V^{3DP} (defined in Problem (4)) as a fractional knapsack problem in Eq. (23).
2. We use this reformulation to provide an efficient subroutine (Algorithm 2) for evaluating V^{3DP} .
3. We use this reformulation and subroutine to provide subgradients of V^{3DP} in closed-form (see Lemma 5).
4. We establish Lemma 6, a technical lemma that characterizes the minimal-magnitude optimal solution of a univariate convex problem. This result will later be used to analyze the optimal 3DP capacity and first-stage order.

We begin with a reformulation of $V^{3DP}(\mathbf{y}, K)$:

$$V^{3DP}(\mathbf{y}, K) = \sum_{j \in \mathcal{A}} v_j y_j^+ + h_j y_j^- + f^{3DP}(\mathbf{y}^+, K), \quad (23)$$

where $\mathbf{y}^+$ is the vector consists of demand shortfall y_j^+ 's, and function f^{3DP} is defined by

$$\begin{aligned} f^{3DP}(\mathbf{z}, K) &\equiv \min_{\mathbf{q}^{3DP}} \sum_{j \in \mathcal{A}} -(v_j - c_j^{3DP}) q_j^{3DP} \\ \text{s.t. } &0 \leq q_j^{3DP} \leq z_j, \forall j \in \mathcal{A} \quad \text{and} \quad \sum_{j \in \mathcal{A}} q_j^{3DP} \leq K. \end{aligned} \quad (24)$$

This reformulation results from the following manipulations:

- For all $j \in \mathcal{A}$ such that $y_j < 0$, there is no demand shortfall to handle, hence the corresponding $q_j^{3DP} = 0$ in Problem (4). Thus, $[y_j - q_j^{3DP}]^-$ equals y_j^- at optimality.
- For all $j \in \mathcal{A}$ such that $y_j \geq 0$, it is never optimal to use 3DP to back up more than the demand shortfall y_j . Therefore, $0 \leq q_j^{3DP} \leq y_j^+$, and $[y_j - q_j^{3DP}]^+$ equals $y_j^+ - q_j^{3DP}$.

The first two terms of the summand in Eq. (23) do not depend on $\mathbf{q}^{3DP}$. This decision variable only occurs in Problem (24), which is a fractional knapsack problem. Therefore, the optimal $\mathbf{q}^{3DP}$ for Problem (4) can be obtained using the standard solution procedure for fractional knapsack problems, as summarized in the following algorithm. For convenience, we assume that the indices in $\mathcal{A}$ are sorted in descending order of $v_j - c_j^{3DP}$ for the remainder of this section.

Algorithm 2 Optimal $\mathbf{q}^{3DP}$ for V^{3DP}

Require: Demand gap y_j for all $j \in \mathcal{A}$ and capacity K . The remaining capacity K' is initially set to K .

- 1: **for** $j = 1 \dots |\mathcal{A}|$ **do**
 - 2: Set $\bar{q}_j^{3DP} = \min\{K', y_j^+\}$, and update K' as $K' - \bar{q}_j^{3DP}$.
 - 3: **end for**
 - 4: **return** $\bar{\mathbf{q}}^{3DP}$ as the optimal solution for Problem (4) which defines $V^{3DP}(\mathbf{y}, K)$
-

In other words, the optimal fulfillment $\bar{q}_j^{3DP}$ of demand shortfalls y_j^+ is obtained by iteratively filling y_j^+ in descending order of the per unit fulfillment reward $v_j - c_j^{3DP}$, until all shortfalls are addressed or the capacity K is exhausted.

Next, we use the structure of $\bar{\mathbf{q}}^{3DP}$, output from Algorithm 2, to derive a closed-form expression for the subgradients of V^{3DP} . This result enables efficient implementation of the first-order method discussed in Section 5.4 and plays a crucial role in subsequent proofs, where it is applied to characterize the optimal first stage order $\mathbf{q}$ and capacity K .

LEMMA 5 (Convexity and Subgradients of V^{3DP}). *The following properties hold for V^{3DP} :*

(i) $V^{3DP}(\mathbf{y}, K)$ is jointly convex in $\mathbf{y}$ and K .

(ii) Let $\bar{\mathbf{q}}^{3DP}$ be the output of Algorithm 2 for given $\mathbf{y}$ and K , and define $j^* \equiv \min\{j \in \mathcal{A} : y_j > 0, \bar{q}_j^{3DP} < y_j\}$ as the first positive shortfall that is not fully filled. Then, a subgradient λ of $V^{3DP}(\mathbf{y}, K)$ with respect to K is given by

$$\lambda = \begin{cases} -(v_{j^*} - c_{j^*}^{3DP}) & \text{if } \sum_{j \in \mathcal{A}} y_j^+ > K \\ 0 & \text{otherwise.} \end{cases} \quad (25)$$

Moreover, a subgradient μ_j of $V^{3DP}(\mathbf{y}, K)$ with respect to y_j is given by

$$\mu_j = \begin{cases} -h_j & \text{if } y_j \leq 0, \\ v_j & \text{if } y_j > 0 \text{ and } \bar{q}_j^{3DP} < y_j, \\ c_j^{3DP} - \lambda & \text{otherwise.} \end{cases} \quad (26)$$

Proof of Lemma 5 This proof relies on an epigraphical reformulation of $V^{3DP}(\mathbf{y}, K)$:

$$\begin{aligned} \min_{\mathbf{q}^{3DP} \geq 0, \zeta^+, \zeta^-} \quad & \sum_{j \in \mathcal{A}} (c_j^{3DP} q_j^{3DP} + v_j \zeta_j^+ + h_j \zeta_j^-) \\ \text{s.t.} \quad & \sum_{j \in \mathcal{A}} q_j^{3DP} \leq K, \end{aligned} \quad (27a)$$

$$y_j - q_j^{3DP} \leq \zeta_j^+, \quad 0 \leq \zeta_j^+, \quad \forall j \in \mathcal{A}, \quad (27b)$$

$$-y_j + q_j^{3DP} \leq \zeta_j^-, \quad 0 \leq \zeta_j^-, \quad \forall j \in \mathcal{A}. \quad (27c)$$

where we have replaced $[y_j - q_j^{3DP}]^+$ and $[y_j - q_j^{3DP}]^-$ in the objective of Problem (4) with nonnegative variables ζ_j^+ and ζ_j^- , subject to constraints $y_j - q_j^{3DP} \leq \zeta_j^+$ and $y_j - q_j^{3DP} \geq -\zeta_j^-$. The dual of this optimization problem is:

$$\begin{aligned} \max_{\lambda \leq 0, \mu^+, \mu^-} \quad & K\lambda + \mathbf{y}^\top (\boldsymbol{\mu}^- - \boldsymbol{\mu}^+) \\ \text{s.t.} \quad & \lambda - \mu_j^+ + \mu_j^- \leq c_j^{3DP}, \quad \forall j \in \mathcal{A}, \end{aligned} \quad (28a)$$

$$-v_j \leq \mu_j^+ \leq 0, \quad \forall j \in \mathcal{A}, \quad (28b)$$

$$-h_j \leq \mu_j^- \leq 0, \quad \forall j \in \mathcal{A}. \quad (28c)$$

Thus, $V^{3DP}(\mathbf{y}, K)$ is the maximum of linear functions and jointly convex in $\mathbf{y}, K$, proving statement (i).

To prove statement (ii), we will apply Danskin's theorem (Proposition 4.5.1 in Bertsekas et al. (2003)), which provides the subgradients of V^{3DP} in terms of an optimal solution to Problem (28). Hence, we first construct a primal and dual pair of optimal solutions to Problems 27 and 28. Namely,

- Let $\mathbf{q}^{3DP} = \bar{\mathbf{q}}^{3DP}$ where $\bar{\mathbf{q}}^{3DP}$ is the output of Algorithm 2.
- Let $\zeta_j^+ = [y_j - \bar{q}_j^{3DP}]^+$ and $\zeta_j^- = [y_j - \bar{q}_j^{3DP}]^-$, for all $j \in \mathcal{A}$.
- Let λ be defined as in Eq. (25).
- For all $j \in \mathcal{A}$ such that $y_j \leq 0$, set $\mu_j^+ = 0$ and $\mu_j^- = -h_j$.
- For all $j \in \mathcal{A}$ such that $y_j > 0$ and $\bar{q}_j^{3DP} = y_j$, set $\mu_j^+ = \lambda - c_j^{3DP}$ and $\mu_j^- = 0$.
- For all $j \in \mathcal{A}$ such that $y_j > 0$ and $\bar{q}_j^{3DP} < y_j$, set $\mu_j^+ = -v_j$ and $\mu_j^- = 0$.

Since $\bar{\mathbf{q}}^{3DP}$ is optimal for Problem (4), it follows that $(\mathbf{q}^{3DP}, \zeta^+, \zeta^-)$ are optimal to Problem (27).

To check dual feasibility, we consider each $j \in \mathcal{A}$ following the cases outlined above:

- For all $j \in \mathcal{A}$ such that $y_j \leq 0$: μ_j^+ and μ_j^- satisfy Eqs. (28b) and (28c) by construction. They also satisfy Eq. (28a), because $\lambda \leq 0$, and hence, the left side of Eq. (28a) is non-positive while the right side is non-negative.
- For all $j \in \mathcal{A}$ such that $y_j > 0$ and $\bar{q}_j^{3DP} = y_j$: Eq. (28c) is satisfied by construction. Since $\bar{q}_j^{3DP} = y_j$, product j cannot have been added to the knapsack after j^* in Algorithm 2. Hence, $v_j - c_j^{3DP} \geq v_{j^*} - c_{j^*}^{3DP} \geq -\lambda$. Rearranging shows that $-v_j \leq \lambda - c_j^{3DP}$, which implies the left inequality in Eq. (28b). The right inequality follows immediately because $\lambda \leq 0$ and $c_j^{3DP} \geq 0$. Finally, by substitution, the Eq. (28a) holds with equality.
- For all $j \in \mathcal{A}$ such that $y_j > 0$ and $\bar{q}_j^{3DP} < y_j$: Eqs. (28b) and (28c) hold by construction. Moreover, because $\bar{q}_j^{3DP} < y_j$, we have j could not have been added to the knapsack before j^* in Algorithm 2, and so $v_j - c_j^{3DP} \leq v_{j^*} - c_{j^*}^{3DP}$, which proves the Eq. (28a) holds.

Thus, our constructed dual solution is dual feasible.

It remains to show complementary slackness:

$$\begin{aligned} \left(\sum_{j \in \mathcal{A}} q_j^{3DP} - K \right) \lambda &= 0, \\ (-q_j^{3DP} - \zeta_j^+ + y_j) \mu_j^+ &= 0, \quad (-q_j^{3DP} + \zeta_j^- + y_j) \mu_j^- = 0, \quad \forall j \in \mathcal{A}, \\ (\lambda - \mu_j^+ + \mu_j^- - c_j^{3DP}) q_j^{3DP} &= 0, \quad (-\mu_j^+ - v_j) \zeta_j^+ = 0, \quad (-\mu_j^- - h_j) \zeta_j^- = 0, \quad \forall j \in \mathcal{A}. \end{aligned} \quad (29)$$

These constraints can again be checked by cases in exactly the same manner as above for dual feasibility, which in turn confirms that $\mu_j^+, \mu_j^-, \lambda$, constitute a dual optimal solution.

Hence, by Danskin's theorem, subgradients with respect to y_j and K are given by $\mu_j = \mu_j^- - \mu_j^+$ and λ , respectively, concluding the proof. $\square$

REMARK 1. We note that the subgradient of V^{3DP} is unique except when $(\mathbf{y}, K)$ falls into one of the following scenarios: (i) $\sum_{j \in \mathcal{A}} y_j^+ > K$ and $\bar{q}_{j^*}^{3DP} = y_{j^*}$; or (ii) $\sum_{j \in \mathcal{A}} y_j^+ = K$. In these cases, the optimal λ for the dual problem Problem (28) forms an interval, and we choose its upper bound as the λ presented in Lemma 5. Accordingly, in the third case of Eq. (26), this choice of λ makes $c_j^{3DP} - \lambda$ the smallest μ_j under this scenario.

This choice of λ and μ_j allows us to directly apply Eq. (26) and Eq. (25) to derive the right derivatives of $K \mapsto V^{3DP}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, K)$ and $q_j \mapsto V^{3DP}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, K)$, which we use in the proofs of Theorem 1 and Theorem 2 to characterize the optimal 3DP capacity and first-stage order. $\square$

Recall that in Theorem 1 and Theorem 2, ties among multiple optima are resolved by selecting the one with the smallest ℓ_1 -norm. When the decision variable is a positive scalar, the following lemma characterizes the least ℓ_1 -norm solution using only the right derivative of the objective function.

LEMMA 6 (Minimal Solution of Univariate Convex Function). *Let $f : \mathbb{R} \mapsto \mathbb{R}$ be convex and let $\bar{x} \equiv \min\{x \geq 0 : f'_+(x) \geq 0\}$, where $f'_+(x)$ denotes the right derivative of f . Then, $\bar{x}$ is the smallest optimal solution to $\min_{x \geq 0} f(x)$.*

Proof of Lemma 6 We first note that $\bar{x}$ is well-defined since $f'_+(x)$ is right-continuous by Bertsekas et al. (2003, Proposition 4.1.1(e)). Consequently, the set $\{x \geq 0 : f'_+(x) \geq 0\}$ attains a minimum.

Next, we observe that any optimal solution must be at least as large as $\bar{x}$. Indeed, by the first order optimality condition, any optimal solution x must satisfy $f'_+(x) \geq 0$, and since $\bar{x}$ is the minimal such value, it is less than or equal to any optimal solution. Thus, it only remains to show $\bar{x}$ is, itself, optimal.

We consider two cases: First suppose that $\bar{x} = 0$. Then, since it is on the boundary, the first order optimality condition directly gives the optimality of $\bar{x}$.

Now suppose $\bar{x} > 0$. Then, it suffices to show that $f'_-(\bar{x}) \leq 0$, where $f'_-(\cdot)$ denotes the left derivative. Suppose by contradiction $f'_-(\bar{x}) > 0$. By Bertsekas et al. (2003, Proposition 4.1.1. e), $f'_-(x)$ is left continuous. Hence, there exists $0 < x_0 < \bar{x}$ with $f'_-(x_0) > 0$. Bertsekas et al. (2003, Proposition 4.1.1. a) further gives $f'_+(x_0) \geq f'_-(x_0)$, which, together, imply $f'_+(x_0) > 0$. But $\bar{x}$ was the minimal value such that $f'_+(x) \geq 0$, a contradiction. Thus, $\bar{x}$ must be an optimal solution. $\square$

A.2. Proofs from Section 2

Proof of Lemma 2. First, note that the objective function of Problem (2), which defines U_j^{DB} , is given by:

$$\mathbb{E} [c_j q_j s_j + c_j^{\text{DB}} [D_j - q_j s_j]^+ + h_j [D_j - q_j s_j]^-] \quad (30)$$

This function is convex since both $[D_j - q_j s_j]^+$ and $[D_j - q_j s_j]^-$ are convex in q_j , and expectation preserves convexity. Lemma 6 thus implies that $\bar{q}_j$ is the smallest q_j for which the right derivative of the function in (30) is nonnegative. This right derivative is given by:

$$\begin{aligned} \frac{d}{dq_j^+} \mathbb{E} [c_j q_j s_j + c_j^{\text{DB}} [D_j - q_j s_j]^+ + h_j [q_j s_j - D_j]^+] &= \mathbb{E} \left[\frac{d}{dq_j^+} (c_j q_j s_j + c_j^{\text{DB}} [D_j - q_j s_j]^+ + h_j [q_j s_j - D_j]^+) \right] \\ &= \mathbb{E} [(c_j s_j - c_j^{\text{DB}} s_j \mathbb{I}\{D_j > q_j s_j\} + h_j s_j \mathbb{I}\{D_j \leq q_j s_j\})], \end{aligned}$$

where we can reverse the derivative and integration in the first equality because the integrand is uniformly Lipschitz in q_j with parameter at most $c_j + c_j^{\text{DB}} + h_j$. Simplifying shows that the right derivative is nonnegative if and only if

$$\mathbb{E} [s_j \mathbb{I}\{D_j \leq q_j s_j\}] \geq \left(\frac{c_j^{\text{DB}} - c_j}{c_j^{\text{DB}} + h_j} \right) \mathbb{E} [s_j].$$

Therefore, the optimal $\bar{q}_j$ is the smallest q_j that satisfies this inequality, which completes the proof. $\square$

Proof of Lemma 3. Throughout this proof, we assume that $\mathcal{A}$ is fixed.

To prove part (i), note that we have already shown in Lemma 5 that $V^{\text{3DP}}(\mathbf{y}, K)$ is jointly convex in $\mathbf{y}$ and K . Since composition with affine functions preserves convexity, it follows that the mapping $\mathbf{q} \mapsto V^{\text{3DP}}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, K)$ is convex.

To prove part (ii), note that expectation preserves convexity (Shapiro et al. 2021, Theorem 7.46), hence the objective function of Problem (3) is convex. Since the feasible region of this problem is also convex, the proof is complete.

Finally, for part (iii), note that by establishing part (i), we have shown that $\mathbf{q} \mapsto V^{\text{3DP}}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, K)$, and thus the objective function of Problem (3)), is jointly convex in $\mathbf{q}$ and K . Since U^{3DP} is derived from partial minimization over $\mathbf{q} \geq 0$, and partial minimization preserves convexity (see Proposition 2.3.6 from Bertsekas et al. (2003)), it follows that $K \mapsto U^{\text{3DP}}(\mathcal{A}, K)$ is convex. $\square$

A.3. Proofs from Section 3

Proof of Proposition 1 To establish $\mathcal{A}_0 \subseteq \mathcal{A}^*$, it suffices to show that the optimal $\mathcal{A}$ for any fixed $K \geq 0$ includes $\mathcal{A}_0$. This holds because adding an unprotected product to $\mathcal{A}$ and choosing to never print it in the second stage is costless.

To show that the inclusion can be strict, it suffices to construct an example where $\mathcal{A} = \mathcal{N}$ is strictly better than $\mathcal{A} = \mathcal{A}_0$. Then, $\mathcal{A}_0$ cannot be optimal. To that end, consider a setting where $0 \leq D_j \leq \bar{D}_j$ almost surely for all $j \in \mathcal{N}$, and $C^{\text{3DP}}(K) = 0$. In this case, we effectively have infinite 3DP capacity for any $\mathcal{A}$, since it is always optimal to set $K = \sum_{j \in \mathcal{A}} \bar{D}_j$. Recall that Problem (6) and Problem (7) define the optimal operational costs for product j when K

is zero and infinite, respectively. We denote these costs as U_j^0 and U_j^∞ . The total cost under $\mathcal{A} = \mathcal{N}$ is thus given by $\sum_{j \in \mathcal{N}} U_j^{3DP}$, while the total cost under $\mathcal{A} = \mathcal{A}_0$ is:

$$\sum_{j \in \mathcal{A}_0} U_j^0 + \sum_{j \in \mathcal{A}_0^c} (U_j^{DB} + C_j^{DB}),$$

where U_j^{DB} (as defined in Problem (2)) and C_j^{DB} represent the operational and fixed costs of dedicated backup, respectively. Comparing the costs for $\mathcal{A} = \mathcal{N}$ and $\mathcal{A} = \mathcal{A}_0$, and simplifying, we find that $\mathcal{A} = \mathcal{N}$ yields a lower cost if:

$$\sum_{j \in \mathcal{A}_0^c} U_j^\infty < \sum_{j \in \mathcal{A}_0^c} (U_j^{DB} + C_j^{DB}).$$

We can force this inequality to be true by taking C_j^{DB} to be sufficiently large. Fix any such sufficiently large value for the remainder of the proof. The only remaining challenge is to assert that $\mathcal{A}_0 \neq \mathcal{N}$, which will occur if

$$U_j^{DB} + C_j^{DB} < U_j^0$$

holds for some j . Notice, that neither of U_j^{DB} nor U_j^∞ depend on v_j (the stock out cost), because DB is uncapacitated and we effectively have infinite 3DP capacity. Hence, by increasing v_j we can increase U_j^0 and ensure that $\mathcal{A}_0 \neq \mathcal{N}$.

Thus, we have constructed an example where $\mathcal{A} = \mathcal{A}_0$ is not optimal, hence the optimal $\mathcal{A}$ can strictly include $\mathcal{A}_0$.

□

Proof of Theorem 1 Recall that $q^*(\mathcal{A}, K)$ is an optimal solution to Problem (3) (which defines $U^{3DP}(\mathcal{A}, K)$) with the minimal ℓ_1 -norm. We first note that $q_j^*(\mathcal{A}, K)$ is the smallest optimal solution to Problem (3) when minimizing only over $q_j \geq 0$, while keeping $q_i = q_i^*(\mathcal{A}, K)$ fixed for all $i \neq j$, i.e., $q_j^*(\mathcal{A}, K)$ solves the following univariate optimization problem:

$$\min_{q \geq 0} c_j q \mathbb{E}[s_j] + \sum_{i \in \mathcal{A}, i \neq j} c_i q_i^*(K) \mathbb{E}[s_i] + \mathbb{E}[V^{3DP}(\mathbf{D} - \hat{\mathbf{q}} \circ \mathbf{s}, \mathcal{A}, K)],$$

which is equivalent to

$$\min_{q \geq 0} \mathbb{E}[c_j q s_j + V^{3DP}(\mathbf{D} - \hat{\mathbf{q}} \circ \mathbf{s}, \mathcal{A}, \kappa)]. \quad (31)$$

In both problems, we let

$$\hat{q}_i = \begin{cases} q_i^*(\mathcal{A}, K) & \text{if } i \neq j, \\ q & \text{otherwise.} \end{cases} \quad (32)$$

Moreover, $q_j^*(\mathcal{A}, K)$ is the minimal solution this univariate problem. Indeed, if this were not the case, we could use it to construct an optimal solution to Problem (3) with a smaller ℓ_1 -norm than $q^*(\mathcal{A}, K)$, contradicting its definition.

Note that Problem (31) is a univariate convex optimization problem. Thus, by Lemma 6, $q_j^*(\mathcal{A}, K)$ is the smallest $q \geq 0$ for which the right derivative of its objective function is nonnegative. This right derivative is given by:

$$\frac{d}{dq^+} \mathbb{E}[c_j q s_j + V^{3DP}(\mathbf{D} - \hat{\mathbf{q}} \circ \mathbf{s}, \mathcal{A}, \kappa)] = \mathbb{E}\left[c_j s_j + \frac{d}{dq^+} V^{3DP}(\mathbf{D} - \hat{\mathbf{q}} \circ \mathbf{s}, \mathcal{A}, K)\right] \quad (33a)$$

$$= \mathbb{E}[c_j s_j - s_j \mu_j(\mathbf{D}, \mathbf{s}, q)] \quad (33b)$$

$$= \mathbb{E}[c_j s_j + h_j s_j \mathbb{I}\{D_j \leq q s_j\} - s_j \mu_j(\mathbf{D}, \mathbf{s}, q) \mathbb{I}\{D_j > q s_j\}], \quad (33c)$$

where $\mu_j(\mathbf{D}, \mathbf{s}, q)$ denotes the μ_j from Eq. (26) after setting $\mathbf{y} = \mathbf{D} - \hat{\mathbf{q}} \circ \mathbf{s}$ in Lemma 5. Specifically, Eq. (33a) swaps differentiation and integration since the integrand is uniformly Lipschitz in q with parameter at most $c_j + v_j + h_j$.

Eq. (33b) applies Proposition 4.2.2 of Bertsekas et al. (2003), which states that the right derivative of $q \mapsto V^{3DP}(D - \hat{q} \circ s, \mathcal{A}, K)$ is its largest subgradient. By the chain rule, this subgradient is given by $-s_j \mu_j(D, s, q)$, since $\mu_j(D, s, q)$ is constructed as the smallest possible subgradient of $y_j \mapsto V^{3DP}(y, \mathcal{A}, K)$ at $y = D - \hat{q} \circ s$ (see Remark 1). Finally, Eq. (33c) follows from the specific construction of $\mu_j(D, s, q)$ shown in Eq. (26).

In summary, we derive that

$$q_j^*(\mathcal{A}, K) = \min \left\{ q \geq 0 : H(q, \mu_j(D, s, q)) \geq 0 \right\} \quad (34)$$

where

$$H(q, z) \equiv c_j s_j + h_j s_j \mathbb{I}\{D_j \leq q s_j\} - s_j z \mathbb{I}\{D_j > q s_j\} \quad (35)$$

Using an argument similar to the proof of Lemma 2, we characterize $\bar{q}_j^\infty$ and $\bar{q}_j^0$, the smallest optimal solutions to Problem (6) and Problem (7), respectively, as follows:

$$\bar{q}_j^\infty = \min\{q \geq 0 : H(q, c_j^{3DP}) \geq 0\} \quad \text{and} \quad \bar{q}_j^0 = \min\{q \geq 0 : H(q, v_j) \geq 0\} \quad (36)$$

Thus, to establish $\bar{q}_j^\infty \leq q_j^{3DP}(\mathcal{A}, K) \leq \bar{q}_j^0$, it suffices to show that for all D, s and $q \geq 0$ such that $D_j > q s_j$, the following inequalities hold:

$$c_j^{3DP} \leq \mu_j(D, s, q) \leq v_j. \quad (37)$$

To see why Eq. (37) is sufficient, observe that $H(q, z)$ is non-increasing in z , thus Eq. (37) yields:

$$H(q, v_j) \leq H(q, \mu_j(D, s, q)) \leq H(q, c_j^{3DP}). \quad (38)$$

Consequently, we obtain

$$0 \leq H(q_j^*(\mathcal{A}, K), \mu_j(D, s, q_j^*(\mathcal{A}, K))) \leq H(q_j^*(\mathcal{A}, K), c_j^{3DP}),$$

where the first inequality follows from Eq. (34), and the second follows from Eq. (38). In other words, $q = q_j^*(\mathcal{A}, K)$ satisfies the inequality $H(q, c_j^{3DP}) \geq 0$. Since $\bar{q}_j^\infty$ is the smallest such q (cf. Eq. (36)), it follows that $\bar{q}_j^\infty \leq q_j^*(\mathcal{A}, K)$. By a similar argument, we also obtain $q_j^*(\mathcal{A}, K) \leq \bar{q}_j^0$.

Thus, the proof is complete once we establish Eq. (37) for all D, s and $q \geq 0$ with $D_j > q s_j$. We consider the following cases, setting $y = D - \hat{q} \circ s$ for simplicity, where $\bar{q}^{3DP}$ and j^* are defined as in Lemma 5 under this y :

- i) If $\bar{q}_j^{3DP} < y_j$, then from Eq. (26), $\mu_j(D, s, q) = v_j$, satisfying Eq. (37).
- ii) If $\bar{q}_j^{3DP} = y_j$ and $\sum_{j \in \mathcal{A}} y_j^+ \leq K$, then from Eq. (26), $\mu_j(D, s, q) = c_j^{3DP}$, satisfying Eq. (37).
- iii) Finally, when $\bar{q}_j^{3DP} = y_j$ and $\sum_{j \in \mathcal{A}} y_j^+ > K$, Eq. (26) gives $\mu_j(D, s, q) = c_j^{3DP} + (v_{j^*} - c_{j^*}^{3DP})$, hence we have $\mu_j(D, s, q) \geq c_j^{3DP}$ since $v_{j^*} - c_{j^*}^{3DP} \geq 0$. Moreover, in this scenario, shortfall j is fully filled, so by construction, $j^* > j$. Since $v_j - c_j^{3DP} \leq v_{j^*} - c_{j^*}^{3DP}$, it follows that $\mu_j(D, s, q) \leq v_j$ after rearranging terms.

In summary, we have shown that Eq. (37) holds in all cases, therefore proving $\bar{q}_j^\infty \leq q_j^*(\mathcal{A}, K) \leq \bar{q}_j^0$. $\square$

Proof of Theorem 2 For ease of reference, let $r_j \equiv v_j - c_j^{3DP}$ and denote the objective function of Problem (3) (which defines U^{3DP}) as $F^{3DP}(K)$. We begin by characterizing $K^*(\mathcal{A}, q)$, which is key to proving statements (i) and (ii).

Recall that $K^*(\mathcal{A}, q)$ is the smallest optimal capacity for a given $\mathcal{A}$ and a fixed first-stage order $q \geq 0$, making it the smallest minimizer of $C^{3DP}(K) + F^{3DP}(K)$ for $K \geq 0$, where $C^{3DP}(K) = c^{\text{cap}} K$. Since this is a univariate convex problem, it follows from Lemma 6 that

$$K^*(\mathcal{A}, q) = \min \{ K \geq 0 : H^*(K, q) \geq 0 \}, \quad (39)$$

where $H^*(K, \mathbf{q})$ is the right derivative of $c^{\text{cap}}K + F^{\text{3DP}}(K)$. We derive $H^*(K, \mathbf{q})$ in closed form as follows:

$$H^*(K, \mathbf{q}) = \frac{d}{dK^+} (c^{\text{cap}}K + F^{\text{3DP}}(K)) = c^{\text{cap}} + \frac{d}{dK^+} \mathbb{E} [V^{\text{3DP}}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, K)] \quad (40a)$$

$$= c^{\text{cap}} + \mathbb{E} \left[\frac{d}{dK^+} V^{\text{3DP}}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, K) \right] \quad (40b)$$

$$= c^{\text{cap}} + \mathbb{E} [\lambda(\mathbf{D}, \mathbf{s}, K)] \quad (40c)$$

$$= c^{\text{cap}} - \mathbb{E} \left[r_{j^*} \mathbb{I} \left\{ \sum_{j \in \mathcal{A}} [D_j - q_j s_j]^+ > K \right\} \right]. \quad (40d)$$

In (40b), $\lambda(\mathbf{D}, \mathbf{s}, K)$ is the λ from Eq. (25), defined under $\mathbf{y} = \mathbf{D} - \mathbf{q} \circ \mathbf{s}$; the index j^* in (40d) is defined in Lemma 5 under the same setting. Equation (40a) holds because V^{3DP} is the only term in F^{3DP} that depends on K . In Eq. (40b), differentiation and integration are interchanged, which is valid since the integrand is uniformly Lipschitz in q with a parameter at most $\max_{j \in \mathcal{N}} (v_j - c_j^{\text{3DP}})$. Next, Eq. (40b) applies Proposition 4.2.2 of Bertsekas et al. (2003), which states that the right derivative of $K \mapsto V^{\text{3DP}}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, K)$ is its largest subgradient. By construction (see Remark 1), this subgradient is precisely $\lambda(\mathbf{D}, \mathbf{s}, K)$, as given by λ in Eq. (26), whose explicit form yields Eq. (40d).

To prove part (i) of the theorem, we substitute $v_j - c_j = r$ for all $j \in \mathcal{A}$ into Eq. (40), which simplifies Eq. (39) to

$$K^*(\mathcal{A}, \mathbf{q}) = \inf \left\{ K \geq 0 : c^{\text{cap}} - r^{\min} \mathbb{P} \left(\sum_{j \in \mathcal{A}} [D_j - q_j s_j]^+ > K \right) \geq 0 \right\} \quad (41)$$

$$= \inf \left\{ K \geq 0 : \mathbb{P} \left(\sum_{j \in \mathcal{A}} [D_j - q_j s_j]^+ > K \right) \leq \frac{c^{\text{cap}}}{r} \right\}. \quad (42)$$

By the definition, $K^*(\mathcal{A}, \mathbf{q})$ is thus the $(1 - \frac{c^{\text{cap}}}{r})$ -quantile of $\sum_{j \in \mathcal{A}} [D_j - q_j s_j]^+$, thereby proving part (i).

We now prove part (ii) of the theorem, where r_j values are heterogeneous. Recall that $K^*(\mathcal{A})$ is the smallest optimal capacity under $\mathcal{A}$, so we have $K^*(\mathcal{A}) = K^*(\mathcal{A}, \mathbf{q}^*(\mathcal{A}))$, where $\mathbf{q}^*(\mathcal{A})$ is the optimal first-stage order with the minimal ℓ_1 -norm. For convenience, we define the following functions

$$H^{\min}(K, \mathbf{q}) \equiv c^{\text{cap}} - r^{\min} \mathbb{P} \left(\sum_{j \in \mathcal{A}} [D_j - q_j s_j]^+ > K \right) \quad \text{and} \quad H^{\max}(K, \mathbf{q}) \equiv c^{\text{cap}} - r^{\max} \mathbb{P} \left(\sum_{j \in \mathcal{A}} [D_j - q_j s_j]^+ > K \right).$$

Since $r^{\min} \leq r_{j^*} \leq r^{\max}$ almost surely, it follows that

$$H^{\min}(K, \mathbf{q}) \leq H^*(K, \mathbf{q}) \leq H^{\max}(K, \mathbf{q}). \quad (43)$$

The proof of part (ii) is completed as follows:

$$\left(1 - \frac{c^{\text{cap}}}{r^{\max}}\right)\text{-quantile of } \sum_{j \in \mathcal{A}} [D_j - \bar{q}_j^{\infty} s_j]^+ \leq \left(1 - \frac{c^{\text{cap}}}{r^{\max}}\right)\text{-quantile of } \sum_{j \in \mathcal{A}} [D_j - q_j^*(\mathcal{A}) s_j]^+ \quad (44a)$$

$$= \inf \left\{ K \geq 0 : \mathbb{P} \left(\sum_{j \in \mathcal{A}} [D_j - q_j^*(\mathcal{A}) s_j]^+ > K \right) \leq \frac{c^{\text{cap}}}{r^{\max}} \right\} \quad (44b)$$

$$= \inf \{ K \geq 0 : H^{\max}(K, \mathbf{q}^*(\mathcal{A})) \geq 0 \} \quad (44c)$$

$$\leq \inf \{ K \geq 0 : H^*(K, \mathbf{q}^*(\mathcal{A})) \geq 0 \} \quad (44d)$$

$$\leq \inf \{ K \geq 0 : H^{\min}(K, \mathbf{q}^*(\mathcal{A})) \geq 0 \} \quad (44e)$$

$$= \inf \left\{ K \geq 0 : \mathbb{P} \left(\sum_{j \in \mathcal{A}} [D_j - q_j^*(\mathcal{A}) s_j]^+ > K \right) \leq \frac{c^{\text{cap}}}{r^{\min}} \right\} \quad (44f)$$

$$= \left(1 - \frac{c^{\text{cap}}}{r^{\text{min}}}\right)\text{-quantile of } \sum_{j \in \mathcal{A}} [D_j - q_j^*(\mathcal{A})s_j]^+ \quad (44g)$$

$$\leq \left(1 - \frac{c^{\text{cap}}}{r^{\text{min}}}\right)\text{-quantile of } \sum_{j \in \mathcal{A}} [D_j - \bar{q}_j^0 s_j]^+. \quad (44h)$$

Specifically, Eq. (44a) and Eq. (44h) follow from the fact that the quantiles of $\sum_{j \in \mathcal{A}} [D_j - q_j s_j]^+$ are non-increasing in q_j , and from the bounds $\bar{q}_j^\infty \leq q_j^*(\mathcal{A}) \leq \bar{q}_j^0$ given in Theorem 1.

Eq. (44b) and Eq. (44g) arise from the definitions of their respective quantiles, while Eq. (44c) and Eq. (44f) follow from the definitions of H^{min} and H^{max} with some rearrangement.

Moreover, the right-hand sides of Eq. (44d) and Eq. (44e) represent two optimization problems with the same objective. By Eq. (43), we have $H^*(K, \mathbf{q}^*(\mathcal{A})) \leq H^{\text{max}}(K, \mathbf{q}^*(\mathcal{A}))$, so the feasible region in Eq. (44d) is larger, yielding a smaller optimal value. Hence, the inequality in Eq. (44d) holds. The same argument applies to Eq. (44e).

Finally, noting that the right-hand side of Eq. (44d) is $K^*(\mathcal{A}, \mathbf{q}^*(\mathcal{A}))$, and using $K^*(\mathcal{A}) = K^*(\mathcal{A}, \mathbf{q}^*(\mathcal{A}))$, we conclude that statement (ii) is proven. $\square$

Proof of Theorem 3 In this proof, we fix the marginal distributions of $(\mathbf{D}, \mathbf{s})$ as well as $\mathcal{A}, K$, and $\mathbf{q}$, omitting them from the notation. Without loss of generality, we assume the indices in $\mathcal{A}$ are ordered in descending $r_j \equiv v_j - c_j^{\text{3DP}}$.

The key idea of the proof is to show that $f^{\text{3DP}}(\mathbf{z})$ in Problem (24) is supermodular on $\mathbb{R}_+^m$. Recall, this function is submodular if for all $\mathbf{z}, \bar{\mathbf{z}} \geq 0$,

$$f^{\text{3DP}}(\mathbf{z}) + f^{\text{3DP}}(\bar{\mathbf{z}}) \leq f^{\text{3DP}}(\min\{\mathbf{z}, \bar{\mathbf{z}}\}) + f^{\text{3DP}}(\max\{\mathbf{z}, \bar{\mathbf{z}}\}),$$

where $\min\{\mathbf{z}, \bar{\mathbf{z}}\}$ and $\max\{\mathbf{z}, \bar{\mathbf{z}}\}$ denote the elementwise minimum and maximum, respectively.

To see why it is sufficient to show that this function is supermodular, note that by Goovaerts and Dhaene (1999, Theorem 6), among all random vectors $\mathbf{y}^+ \equiv [\mathbf{D} - \mathbf{q} \circ \mathbf{s}]^+$ with fixed marginal distributions, the comonotonic $\mathbf{y}^+$ dominates all others in the supermodular order. Thus, if $f^{\text{3DP}}(\mathbf{z})$ is supermodular, then from the definition of supermodular order, it holds that $\mathbb{E}[f^{\text{3DP}}(\mathbf{y}^+)]$ is maximized when $\mathbf{y}^+$ is comonotonic. Since this result holds for any $\mathcal{A} \subseteq \mathcal{N}$, $K \geq 0$, and $\mathbf{q} \geq 0$, and since $\mathbb{E}[f^{\text{3DP}}(\mathbf{y}^+)]$ is the only term in the objective function of Problem (3) (and thus Problem (1)) that depends on the correlation structure of $\mathbf{y}^+$, it follows that the optimal value of Problem (1) is maximized when $\mathbf{y}^+$ is comonotonic.

Hence, for the remainder of this proof, we focus on establishing the supermodularity of $f^{\text{3DP}}(\mathbf{z})$. By Corollary 2.6.1 in Topkis (1998), it suffices to show that $f^{\text{3DP}}(\mathbf{z})$ satisfies the *increasing differences* property. Specifically, for any $i, j \in \mathcal{A}$, fix z_ℓ for all $\ell \in \mathcal{A} \setminus \{i, j\}$ and express f^{3DP} solely as a function of z_i and z_j , i.e., $f^{\text{3DP}}(z_i, z_j)$, then

$$f^{\text{3DP}}(\bar{z}_i, z'_j + \Delta) - f^{\text{3DP}}(\bar{z}_i, z'_j) \geq f^{\text{3DP}}(z'_i, z'_j + \Delta) - f^{\text{3DP}}(z'_i, z'_j), \quad (45)$$

for all nonnegative scalars $z'_i, \bar{z}_i, z'_j$ and Δ with $z'_i \leq \bar{z}_i$.

We first prove Eq. (45) for sufficiently small $\Delta > 0$ such that $\lambda(z'_i, z'_j) = \lambda(z'_i, z'_j + \Delta)$, where $\lambda(z_i, z_j)$ is the optimal λ in the following dual reformulation of $f^{\text{3DP}}(z_i, z_j)$, with z_ℓ fixed for all $\ell \notin \{i, j\}$:

$$\begin{aligned} f^{\text{3DP}}(z_i, z_j) &= \max_{\lambda \leq 0, \mu} K\lambda + \sum_{\ell \in \mathcal{A}} \mu_\ell z_\ell &= \max_{\lambda \leq 0} K\lambda - \sum_{\ell \in \mathcal{A}} [r_\ell + \lambda]^+ z_\ell. \\ \text{s.t. } &\mu_\ell \leq -\lambda - r_\ell, \mu_\ell \leq 0, \forall \ell \in \mathcal{A} \end{aligned} \quad (46)$$

The second equality in Problem (46) follows from the dual constraint $\mu_\ell \leq \min\{0, -\lambda - r_\ell\} = -[r_\ell + \lambda]^+$, hence maximizing the dual objective yields $\mu_\ell = -[r_\ell + \lambda]^+$ at optimality.

By the same argument in Lemma 5, after replacing y_ℓ^+ with z_ℓ , the λ defined in Eq. (25) remains optimal for Problem (46). Thus, applying Danskin's theorem, the subgradient of $z_j \mapsto f^{3DP}(z_i, z_j)$ is given by $-[r_j + \lambda(z_i, z_j)]^+$, and:

$$f^{3DP}(z'_i, z'_j + \Delta) - f^{3DP}(z'_i, z'_j) \leq -\Delta[r_j + \lambda(z'_i, z'_j + \Delta)]^+ = -\Delta[r_j + \lambda(z'_i, z'_j)]^+, \quad (47)$$

$$f^{3DP}(\bar{z}_i, z'_j + \Delta) - f^{3DP}(\bar{z}_i, z'_j) \geq -\Delta[r_j + \lambda(\bar{z}_i, z'_j)]^+. \quad (48)$$

The inequalities in Eq. (47) and Eq. (48) follow from subgradient inequality of convex function f^{3DP} , while the equality in Eq. (47) holds due to the $\lambda(z'_i, z'_j) = \lambda(z'_i, z'_j + \Delta)$ assumption. Since $t \mapsto [t]^+$ is non-decreasing, for Eq. (45) to hold, it suffices to show that $\lambda(\bar{z}_i, z'_j) \leq \lambda(z'_i, z'_j)$.

Recall that the optimal λ for Problem (46) (as given in Eq. (25)) is either $-r_{j^*}$, where j^* is the first positive shortfall that is not fully filled, or 0 when all shortfalls are fully filled. To unify these two scenarios, we introduce an artificial shortfall $z_{m+1} = K$ with reward $r_{m+1} = 0$, so that we can always write the optimal λ as $-r_{j^*}$. Under this framework, the index j^* is smaller in the $z_i = \bar{z}_i$ case since it requires filling more shortfalls than $z_i = z'_i$. Consequently, $-r_{j^*}$ is smaller for $z_i = \bar{z}_i$, leading to $\lambda(\bar{z}_i, z'_j) \leq \lambda(z'_i, z'_j)$. Therefore, Eq. (45) holds under $\lambda(z'_i, z'_j) = \lambda(z'_i, z'_j + \Delta)$.

To extend this result to general $\Delta > 0$, note that $z_j \mapsto f^{3DP}(z)$ is convex piecewise affine from Problem (46). Thus, the interval $[z'_j, z'_j + \Delta]$ can be partitioned into subintervals $[z'_j + \Delta_{\ell-1}, z'_j + \Delta_\ell]$ for $\ell = 1, \dots, L$, where $\Delta_0 = 0$ and $\Delta_L = \Delta$, such that within the interior of each subinterval, $\lambda(z'_i, z_j)$ remains constant and takes a unique value λ_ℓ^* .

Although $\lambda(z'_i, z'_j + \Delta_{\ell-1})$ and $\lambda(z'_i, z'_j + \Delta_\ell)$ are not unique, Bertsekas et al. (2003)[Proposition 4.2.3 (b)] ensures that both can take the value λ_ℓ^* . Thus, when restricted to each subinterval $[z'_j + \Delta_{\ell-1}, z'_j + \Delta_\ell]$, we may assume without loss of generality that $\lambda(z'_i, z'_j + \Delta_{\ell-1}) = \lambda(z'_i, z'_j + \Delta_\ell)$ and apply Eq. (45) under this condition to conclude that

$$f^{3DP}(z'_i, z'_j + \Delta_\ell) - f^{3DP}(z'_i, z'_j + \Delta_{\ell-1}) \leq f^{3DP}(\bar{z}_i, z'_j + \Delta_\ell) - f^{3DP}(\bar{z}_i, z'_j + \Delta_{\ell-1}). \quad (49)$$

Finally, summing Eq. (49) over all $\ell = 1, \dots, L$ establishes Eq. (45) for any $\Delta > 0$, thus completing the proof. $\square$

A.4. Proofs from Section 4

Proof of Lemma 4 First, recall that $y = D - q \circ s$ for some fixed D, s , and $q \geq 0$. We can without loss of generality assume $q_j = 0$ for all $j \in \mathcal{N}$ such that $x_j = 0$, since all terms indexed by such j in Problem (9) and (8) are effectively eliminated. On the other hand, we can assume that $\bar{q}_j^\infty \leq q_j \leq \bar{q}_j^0$ for all $j \in \mathcal{N}$ where $x_j = 1$, as the optimal solution lies within this range (see Theorem 1).

For convenience, we denote the right hand side of (10) as $\hat{V}^{3DP}$, so verifying $V^{3DP} = \hat{V}^{3DP}$ establishes the equivalence in (10). It suffices to show that setting M_j^1, M_j^2 , and M_j^3 as in Eq. (11) ensures Problem (10) includes the optimal solution $\bar{q}^{3DP}$ of Problem (9) and that their objective functions simplify to the same expression.

Specifically, for all $j \in \mathcal{N}$ such that $x_j = 1$, inequality (10b), (10c), and (10d) in $\hat{V}^{3DP}$ simplify to:

$$(10b) \Rightarrow 0 \leq q_j^{3DP} \leq M_j^1, \quad (10c) \Rightarrow z_j^{3DP} = q_j^{3DP}, \quad (10d) \Rightarrow |z_j^{3DP} - y_j| \leq M_j^3. \quad (50)$$

To show that $\bar{q}_j^{3DP}$ satisfies (50), it suffices to verify that M_j^1 and M_j^3 bound $|\bar{q}_j^{3DP}|$ and $|D_j - q_j s_j - \bar{q}_j^{3DP}|$, respectively. Indeed, note that $\bar{q}_j^{3DP} \leq [D_j - q_j s_j]^+$, our boundedness assumptions on D_j and s_j ensures that $0 \leq \bar{q}_j^{3DP} \leq \bar{D}_j - \bar{q}_j^\infty s_j^{\min}$, which justifies M_j^1 . Similarly, we derive $-\bar{q}_j^0 s_j^{\max} \leq D_j - q_j s_j - \bar{q}_j^{3DP} \leq \bar{D}_j - \bar{q}_j^\infty s_j^{\min}$, justifying M_j^3 .

On the other hand, for $j \in \mathcal{N}$ such that $x_j = 0$, inequality (10b), (10c) and (10d) in $\hat{V}^{3DP}$ can be simplified as

$$(10b) \Rightarrow q_j^{3DP} = 0, \quad (10c) \Rightarrow |z_j^{3DP}| \leq M_j^2, \quad (10d) \Rightarrow z_j^{3DP} - y_j = 0. \quad (51)$$

Note that when $x_j = 0$, the optimal $\bar{q}_j^{3DP}$ for V^{3DP} is arbitrary and can be set to 0, trivially satisfying (51). It remains to verify that $|z_j^{3DP}| \leq M_j^2$. Since $z_j^{3DP} = y_j = D_j - q_j s_j = D_j$, where the last equality follows from $q_j = 0$ in this case, choosing $M_j^2 = \bar{D}_j$ is sufficient for $|z_j^{3DP}| \leq M_j^2$ to hold.

Combining (50) and (51) simplifies (10e) in $\hat{V}^{3DP}$ to $\sum_{j \in \mathcal{N}: x_j=1} q_j^{3DP} \leq K$, which is equivalent to $\sum_{j \in \mathcal{N}} q_j^{3DP} x_j \leq K$ in V^{3DP} . Thus, we have shown that any vector $\bar{q}^{3DP}$ optimal for V^{3DP} satisfies all the constraints in $\hat{V}^{3DP}$.

Finally, substituting (50) and (51) into the objective of Problem (10) confirms its equivalence to the objective of Problem (9), thereby establishing that $\hat{V}^{3DP} = V^{3DP}$.

The convexity of $\hat{V}^{3DP}(\mathbf{y}, \mathbf{x})$ then follows by penalizing the constraints in Problem (10) with an indicator function and applying Proposition 2.3.6 from Bertsekas et al. (2003). $\square$

Proof of Theorem 4 By Lemma 4 and Theorem 1, it holds that:

$$\begin{aligned} U^{3DP}(\mathbf{x}, K) = \min_{\mathbf{q}} \quad & \sum_{j=1}^n c_j q_j \mathbb{E}[s_j] + \mathbb{E}_{\mathbf{D}, \mathbf{s}} [V^{3DP}(\mathbf{D} - \mathbf{q} \circ \mathbf{s}, \mathbf{x}, K)] \\ \text{s.t.} \quad & \bar{q}_j^\infty x_j \leq q_j \leq \bar{q}_j^0 x_j, \quad \forall j = 1 \dots n. \end{aligned} \quad (52)$$

where $V^{3DP}(\mathbf{y}, \mathbf{x}, K)$ is given by Problem (10). Note that this is sufficient to establish the equivalence between Problem (12) and Problem (1). The convexity of the objective function in Problem (12) follows from the convexity of V^{3DP} (as established in Lemma 4) and the fact that the expectation operator preserves convexity. $\square$

A.5. Proofs from Section 5

Proof of Proposition 2. First, note that $\sum_{j \in \mathcal{A}} c_j q_j \mathbb{E}[s_j]$ is modular in $\mathcal{A}$, and since expectation preserves supermodularity, it suffices to show that V^{3DP} is supermodular in $\mathcal{A}$. Using the reformulation in Eq. (23), this reduces to proving that $f^{3DP}(\mathbf{z}, \mathcal{A}, K)$ is supermodular in $\mathcal{A}$ for any $\mathbf{z} \in \mathbb{R}_+^n$ and $K \geq 0$, where f^{3DP} is given in Problem (24) and explicitly written here as a function of $\mathcal{A}$. Specifically, for all $\mathcal{S}, \mathcal{T} \subseteq \mathcal{N}$ with $\mathcal{S} \subset \mathcal{T}$ and $i \notin \mathcal{T}$, we aim to show that:

$$f^{3DP}(\mathbf{z}, \mathcal{S} \cup \{i\}, K) - f^{3DP}(\mathbf{z}, \mathcal{S}, K) \leq f^{3DP}(\mathbf{z}, \mathcal{T} \cup \{i\}, K) - f^{3DP}(\mathbf{z}, \mathcal{T}, K). \quad (53)$$

Note that by the construction of f^{3DP} in Problem (24), we have $f^{3DP}(\mathbf{z}, \mathcal{A}, K) = f^{3DP}(\mathbf{z}(\mathcal{A}), \mathcal{N}, K)$, where

$$\mathbf{z}(\mathcal{A}) \equiv \begin{cases} z_j, & \text{if } j \in \mathcal{A}, \\ 0, & \text{otherwise.} \end{cases}$$

Substituting this identity into Eq. (53), it suffices to prove the equivalent inequality:

$$f^{3DP}(\mathbf{z}(\mathcal{S} \cup \{i\})) - f^{3DP}(\mathbf{z}(\mathcal{S})) \leq f^{3DP}(\mathbf{z}(\mathcal{T} \cup \{i\})) - f^{3DP}(\mathbf{z}(\mathcal{T})). \quad (54)$$

From the proof of Theorem 3, we know that f^{3DP} satisfies the *increasing differences* property:

$$f^{3DP}(\boldsymbol{\zeta} + \beta \mathbf{e}_\ell) - f^{3DP}(\boldsymbol{\zeta}) \leq f^{3DP}(\boldsymbol{\zeta} + \alpha \mathbf{e}_k + \beta \mathbf{e}_\ell) - f^{3DP}(\boldsymbol{\zeta} + \alpha \mathbf{e}_k), \quad (55)$$

for any $\boldsymbol{\zeta} \in \mathbb{R}_+^n$ and $\alpha, \beta > 0$, where $\mathbf{e}_k$ and $\mathbf{e}_\ell$ denote the k^{th} and ℓ^{th} coordinate vectors.

We now apply this *increasing differences* property of f^{3DP} to prove Eq. (54) by induction. Reordering the indices so that $\mathcal{S} \setminus \mathcal{T} = \{1, \dots, L\}$ with $L = |\mathcal{S} \setminus \mathcal{T}|$, the base case follows as:

$$\begin{aligned} f^{3DP}(\mathbf{z}(\mathcal{S} \cup \{i\})) - f^{3DP}(\mathbf{z}(\mathcal{S})) &= f^{3DP}(\mathbf{z}(\mathcal{S}) + z_i \mathbf{e}_i) - f^{3DP}(\mathbf{z}(\mathcal{S})) \\ &\leq f^{3DP}(\mathbf{z}(\mathcal{S}) + z_1 \mathbf{e}_1 + z_i \mathbf{e}_i) - f^{3DP}(\mathbf{z}(\mathcal{S}) + z_1 \mathbf{e}_1), \end{aligned} \quad (56)$$

where the equality follows from the definition of $z(\mathcal{S})$, and the inequality follows by applying Eq. (55) with $\zeta = z(\mathcal{S})$, $\alpha = z_1$, $\beta = z_i$, and setting $k = 1$, $\ell = i$. Now, assume that for some $j' < L$, the following holds:

$$f^{3DP}(z(\mathcal{S} \cup \{i\})) - f^{3DP}(z(\mathcal{S})) \leq f^{3DP}\left(z(\mathcal{S}) + \sum_{j=1}^{j'} z_j e_j + z_i e_i\right) - f^{3DP}\left(z(\mathcal{S}) + \sum_{j=1}^{j'} z_j e_j\right). \quad (57)$$

Applying Eq. (55) with $\zeta = z(\mathcal{S}) + \sum_{j=1}^{j'} z_j e_j$, $\alpha = z_{j'+1}$, $\beta = z_i$, and setting $k = j' + 1$, $\ell = i$, we obtain:

$$f^{3DP}\left(z(\mathcal{S}) + \sum_{j=1}^{j'} z_j e_j + z_i e_i\right) - f^{3DP}\left(z(\mathcal{S}) + \sum_{j=1}^{j'} z_j e_j\right) \leq f^{3DP}\left(z(\mathcal{S}) + \sum_{j=1}^{j'+1} z_j e_j + z_i e_i\right) - f^{3DP}\left(z(\mathcal{S}) + \sum_{j=1}^{j'+1} z_j e_j\right).$$

Combining this with Eq. (57), we establish:

$$f^{3DP}(z(\mathcal{S} \cup \{i\})) - f^{3DP}(z(\mathcal{S})) \leq f^{3DP}\left(z(\mathcal{S}) + \sum_{j=1}^{j'+1} z_j e_j + z_i e_i\right) - f^{3DP}\left(z(\mathcal{S}) + \sum_{j=1}^{j'+1} z_j e_j\right).$$

Thus, by induction, Eq. (57) holds for all $1 \leq j' \leq L$. In particular, substituting the constructions

$$z(\mathcal{T}) = z(\mathcal{S}) + \sum_{j=1}^L z_j e_j, \quad z(\mathcal{T} \cup \{i\}) = z(\mathcal{S}) + \sum_{j=1}^L z_j e_j + z_i e_i, \quad (58)$$

into Eq. (57) with $j' = L$, we establish Eq. (54), hence completing the proof. $\square$

Proof of Proposition 3 This result is proven by counterexample. Let $n = 3$ with $s_j \sim \text{Bernoulli}(0.9)$ (all-or-nothing disruptions), assuming all uncertainties are independent across suppliers and let D follow a two-point distribution. The table below presents the scenarios and probabilities of D_j , along with the cost parameters.

	c_j	c_j^{3DP}	v_j	h_j	D_j scenario 1	D_j probability 1	D_j scenario 2	D_j probability 2
$j = 1$	0.4	0.1	1	0.1	150	0.2	340	0.8
$j = 2$	0.3	0.2	0.6	0.2	140	0.1	150	0.9
$j = 3$	0.6	0.5	1	0.5	70	0.1	140	0.9

Under this setting, we have $U^{3DP}(\{1, 2, 3\}) - U^{3DP}(\{1, 2\}) = 89.47$ and $U^{3DP}(\{1, 3\}) - U^{3DP}(\{1\}) = 91.07$, violating the supermodularity condition in (13). Similarly, $U^{3DP}(\{1, 2, 3\}) - U^{3DP}(\{2, 3\}) = 150.44$ and $U^{3DP}(\{1, 2\}) - U^{3DP}(\{1\}) = 149.52$, contradicting submodularity. $\square$

Proof of Theorem 5 For simplicity, we will occasionally express U^{3DP} and L^{3DP} only as functions of K , with the understanding that $\mathcal{A}$ is fixed.

Part i) and ii) With the convexity of U^{3DP} in K established in Lemma 3, we derive a lower bound on U^{3DP} as follows:

$$U^{3DP}(K) \geq U^{3DP}(0) + K\eta = \sum_{j \in \mathcal{A}} U_j^0 + K \mathbb{E}[\lambda(\mathbf{D}, \mathbf{s}, 0)] \quad (59a)$$

$$= \sum_{j \in \mathcal{A}} U_j^0 - K \mathbb{E}_{\mathbf{D}, \mathbf{s}} \left[\max_{j \in \mathcal{A}} (v_j - c_j^{3DP}) \mathbb{I} \{ [D_j - \bar{q}_j^0 s_j]^+ > 0 \} \right]. \quad (59b)$$

Here, η in Eq. (59a) is the largest subgradient of $\kappa \mapsto U^{3DP}(\kappa)$ at $\kappa = 0$, and the inequality in Eq. (59a) follows from the subgradient inequality of U^{3DP} . Meanwhile, $\lambda(\mathbf{D}, \mathbf{s}, 0)$ in Eq. (59a) corresponds to the λ from Eq. (25), defined under $\mathbf{y} = \mathbf{D} - \bar{\mathbf{q}}^0 \circ \mathbf{s}$ and $K = 0$ for some $\mathbf{D}$ and $\mathbf{s}$, making it the largest subgradient of $\kappa \mapsto V^{3DP}(\mathbf{D} - \bar{\mathbf{q}}^0 \circ \mathbf{s}, \kappa)$ at $\kappa = 0$ by construction (see Remark 1). Applying Danskin's theorem to U^{3DP} , we obtain $\eta = \mathbb{E}[\lambda(\mathbf{D}, \mathbf{s}, 0)]$, validating the equality in Eq. (59a), while Eq. (59b) follows directly from the construction of $\lambda(\mathbf{D}, \mathbf{s}, 0)$ in Eq. (25).

Finally, defining the right-hand side of (59b) as $L^{3DP}(K)$ completes the proof of part (i). Part (ii) then follows taking the limit $K \rightarrow 0$ on both sides of Eq. (59b).

Part iii) The first term in (14) is trivially modular in $\mathcal{A}$, therefore to prove supermodularity of L^{3DP} in $\mathcal{A}$, it is suffice to show that the following function is submodular in $\mathcal{A}$:

$$E(\mathcal{A}) \equiv \mathbb{E}_{D,s} \left[\max_{j \in \mathcal{A}} r_j^{3DP} \mathbb{I} \{ \bar{D}_j^* > 0 \} \right] \quad (60)$$

where $r_j^{3DP} \equiv v_j - c_j^{3DP}$ and $\bar{D}_j^* \equiv [D_j - \bar{q}_j^0 s_j]^+$. Note that:

$$E(\mathcal{A}) = \mathbb{E}_{D,s} \left[\sum_{\mathcal{A}' \subseteq \mathcal{A}} \mathbb{I} \left\{ \begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{A}', \\ \bar{D}_j^* = 0, \forall j \in \mathcal{A} \setminus \mathcal{A}' \end{array} \right\} \left(\max_{j \in \mathcal{A}'} r_j^{3DP} \right) \right] = \sum_{\mathcal{A}' \subseteq \mathcal{A}} \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{A}', \\ \bar{D}_j^* = 0, \forall j \in \mathcal{A} \setminus \mathcal{A}' \end{array} \right) \left(\max_{j \in \mathcal{A}'} r_j^{3DP} \right)$$

Let $\mathcal{B} \subseteq \mathcal{N}$ be an arbitrary set that $\mathcal{A} \subseteq \mathcal{B}$, and let $i \in \mathcal{N} \setminus \mathcal{B}$, first we have

$$E(\mathcal{A}) = \sum_{\mathcal{A}' \subseteq \mathcal{A}} \left[\mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{A}' \cup \{i\}, \\ \bar{D}_j^* = 0, \forall j \in \mathcal{A} \setminus \mathcal{A}' \end{array} \right) + \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{A}', \\ \bar{D}_j^* = 0, \forall j \in (\mathcal{A} \setminus \mathcal{A}') \cup \{i\} \end{array} \right) \right] \left(\max_{j \in \mathcal{A}'} r_j^{3DP} \right)$$

Note that

$$\begin{aligned} E(\mathcal{A} \cup \{i\}) &= \sum_{\mathcal{A}' \subseteq \mathcal{A}} \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{A}' \cup \{i\}, \\ \bar{D}_j^* = 0, \forall j \in \mathcal{A} \setminus \mathcal{A}' \end{array} \right) \max \left\{ r_i^{3DP}, \max_{j \in \mathcal{A}'} r_j^{3DP} \right\} \\ &\quad + \sum_{\mathcal{A}' \subseteq \mathcal{A}} \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{A}', \\ \bar{D}_j^* = 0, \forall j \in (\mathcal{A} \setminus \mathcal{A}') \cup \{i\} \end{array} \right) \left(\max_{j \in \mathcal{A}'} r_j^{3DP} \right) \end{aligned}$$

Hence

$$E(\mathcal{A} \cup \{i\}) - E(\mathcal{A}) = \sum_{\mathcal{A}' \subseteq \mathcal{A}} \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{A}' \cup \{i\}, \\ \bar{D}_j^* = 0, \forall j \in \mathcal{A} \setminus \mathcal{A}' \end{array} \right) \max \left\{ 0, r_i^{3DP} - \max_{j \in \mathcal{A}'} r_j^{3DP} \right\} \geq 0 \quad (61)$$

We can similarly obtain

$$E(\mathcal{B} \cup \{i\}) - E(\mathcal{B}) = \sum_{\mathcal{B}' \subseteq \mathcal{B}} \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{B}' \cup \{i\}, \\ \bar{D}_j^* = 0, \forall j \in \mathcal{B} \setminus \mathcal{B}' \end{array} \right) \max \left\{ 0, r_i^{3DP} - \max_{j \in \mathcal{B}'} r_j^{3DP} \right\} \geq 0 \quad (62)$$

To show the submodularity of H , we need to prove that

$$E(\mathcal{A} \cup \{i\}) - E(\mathcal{A}) \geq E(\mathcal{B} \cup \{i\}) - E(\mathcal{B})$$

Notice that from (61) we have

$$\begin{aligned} &E(\mathcal{A} \cup \{i\}) - E(\mathcal{A}) \\ &= \sum_{\mathcal{A}' \subseteq \mathcal{A}} \sum_{\mathcal{E} \subseteq \mathcal{B} \setminus \mathcal{A}} \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{A}' \cup \mathcal{E} \cup \{i\}, \\ \bar{D}_j^* = 0, \forall j \in (\mathcal{A} \setminus \mathcal{A}') \cup [(\mathcal{B} \setminus \mathcal{A}) \setminus \mathcal{E}] \end{array} \right) \max \left\{ 0, r_i^{3DP} - \max_{j \in \mathcal{A}'} r_j^{3DP} \right\} \\ &= \sum_{\mathcal{B}' \subseteq \mathcal{B}} \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{B}' \cup \{i\}, \\ \bar{D}_j^* = 0, \forall j \in \mathcal{B} \setminus \mathcal{B}' \end{array} \right) \max \left\{ 0, r_i^{3DP} - \max_{j \in \mathcal{B}' \cap \mathcal{A}} r_j^{3DP} \right\} \end{aligned}$$

where the second equality follows from the fact that any $\mathcal{B}' \subseteq \mathcal{B}$ can be uniquely decomposed into two subsets, one contained in $\mathcal{A}$ and the other in $\mathcal{B} \setminus \mathcal{A}$. Combine this result with (62), we obtain

$$\begin{aligned} &E(\mathcal{A} \cup \{i\}) - E(\mathcal{A}) - [E(\mathcal{B} \cup \{i\}) - E(\mathcal{B})] \\ &= \sum_{\mathcal{B}' \subseteq \mathcal{B}} \mathbb{P} \left(\begin{array}{l} \bar{D}_j^* > 0, \forall j \in \mathcal{B}' \cup \{i\}, \\ \bar{D}_j^* = 0, \forall j \in \mathcal{B} \setminus \mathcal{B}' \end{array} \right) \underbrace{\left(\max \left\{ 0, r_i^{3DP} - \max_{j \in \mathcal{B}' \cap \mathcal{A}} r_j^{3DP} \right\} - \max \left\{ 0, r_i^{3DP} - \max_{j \in \mathcal{B}'} r_j^{3DP} \right\} \right)}_{\geq 0, \text{ since } \mathcal{B}' \cap \mathcal{A} \subseteq \mathcal{B}'} \geq 0 \end{aligned}$$

Thus we have proven the submodularity of $E(\mathcal{A})$ hence the supermodularity of L^{3DP} in $\mathcal{A}$. $\square$

Proof of Proposition 4 For convenience, throughout this proof, we use the following notations:

$$r_j^{3DP} \equiv v_j - c_j^{3DP}, \quad \bar{D}_j^* \equiv [D_j - \bar{q}_j^0 s_j]^+, \quad E(\mathcal{A}) \equiv \mathbb{E}_{D,s} \left[\max_{j \in \mathcal{A}} r_j^{3DP} \mathbb{I} \{ \bar{D}_j^* > 0 \} \right]$$

Part i) Let $\mathcal{A} \subseteq \mathcal{N}$ and $K \geq 0$ be given, then

$$E(\mathcal{A}) \leq \left(\max_{j \in \mathcal{N}} r_j^{3DP} \right) \mathbb{E}_{D,s} \left[\max_{j \in \mathcal{A}} \mathbb{I} \{ \bar{D}_j^* > 0 \} \right] = \left(\max_{j \in \mathcal{N}} r_j^{3DP} \right) \mathbb{P}(\exists j \in \mathcal{A} : \bar{D}_j^* > 0), \quad (63)$$

which show that $\hat{L}^{3DP}(\mathcal{A}, K)$ is a lower bound of $L^{3DP}(\mathcal{A}, K)$.

Part ii) Denote the right-hand-side of (63) as $\hat{E}(\mathcal{A})$, then to show the supermodularity of $\hat{L}^{3DP}(\mathcal{A}, K)$, it suffices to prove that $\hat{E}(\mathcal{A})$ is submodular. Let $\mathcal{A}$ and $\mathcal{B}$ be subsets of $\mathcal{N}$ so that $\mathcal{A} \subseteq \mathcal{B}$ and let $i \in \mathcal{N} \setminus \mathcal{B}$, then

$$\begin{aligned} & \hat{E}(\mathcal{A} \cup \{i\}) - \hat{E}(\mathcal{A}) - [\hat{E}(\mathcal{B} \cup \{i\}) - \hat{E}(\mathcal{B})] \\ &= \left(\max_{j \in \mathcal{N}} r_j^{3DP} \right) \mathbb{P} \left(\{ \bar{D}_j^* = 0, \forall j \in \mathcal{A} \} \cap \{ \bar{D}_i^* = 0 \} \cap \{ \exists j \in \mathcal{B} \setminus \mathcal{A} \text{ s.t. } \bar{D}_j^* > 0 \} \right) \geq 0 \end{aligned}$$

This verifies that $\hat{E}$ is submodular, hence $\hat{L}^{3DP}$ is supermodular in $\mathcal{A}$.

Part iii) From (63), it is trivial to show that $L^{3DP} = \hat{L}^{3DP}$ if r_j^{3DP} are identical across $j \in \mathcal{N}$.

Part iv) For ease of notation, we define:

$$\tilde{E}(\mathcal{A}, \lambda) \equiv \frac{1}{\lambda} \log \left(\sum_{j \in \mathcal{A}} [\mathbb{P}(\bar{D}_j^* = 0) + \mathbb{P}(\bar{D}_j^* > 0) e^{\lambda r_j}] \right) \quad (64)$$

To prove that $L^{3DP}(\mathcal{A}, K) \geq \tilde{L}^{3DP}(\mathcal{A}, K, \lambda)$ for all $\mathcal{A} \subseteq \mathcal{N}, K \geq 0$ and $\lambda > 0$, it suffices to show that $\tilde{E}(\mathcal{A}, \lambda) \geq E(\mathcal{A})$ for all $\lambda > 0$. Specifically, first note that for any collection of random variables $\{X_j : j \in \mathcal{A}\}$, it holds that:

$$\exp \left(\lambda \mathbb{E} \left[\max_{j \in \mathcal{A}} X_j \right] \right) \leq \mathbb{E} \left[\exp \left(\lambda \max_{j \in \mathcal{A}} X_j \right) \right] = \mathbb{E} \left[\max_{j \in \mathcal{A}} e^{\lambda X_j} \right] \leq \sum_{j \in \mathcal{A}} \mathbb{E} [e^{\lambda X_j}]$$

Applying logarithm for both sides gives us:

$$\mathbb{E} \left[\max_{j \in \mathcal{A}} X_j \right] \leq \frac{1}{\lambda} \log \left(\sum_{j \in \mathcal{A}} \mathbb{E} [e^{\lambda X_j}] \right) \quad (65)$$

Substituting X_j in Eq. (65) with $r_j^{3DP} \mathbb{I} \{ \bar{D}_j^* > 0 \}$, the left-hand side becomes $E(\mathcal{A})$ as defined in Eq. (60), while the right-hand side corresponds to $\tilde{E}(\mathcal{A}, \lambda)$. Therefore, we have shown that $\tilde{E}(\mathcal{A}, \lambda) \geq E(\mathcal{A})$.

Part v) To prove that $\tilde{L}^{3DP}(\mathcal{A}, K, \lambda)$ is supermodular in $\mathcal{A}$, it suffices to show that $\tilde{E}(\mathcal{A}, \lambda)$ defined in Eq. (64) is submodular in $\mathcal{A}$. Note that for arbitrary $\mathcal{S} \subseteq \mathcal{N}$ and $i \notin \mathcal{S}$, it holds that:

$$\tilde{E}(\mathcal{S} \cup \{i\}, \lambda) - \tilde{E}(\mathcal{S}, \lambda) = \frac{1}{\lambda} \log \left(1 + \frac{\mathbb{P}(\bar{D}_i^* = 0) + \mathbb{P}(\bar{D}_i^* > 0) e^{\lambda r_i^{3DP}}}{\sum_{j \in \mathcal{S}} [\mathbb{P}(\bar{D}_j^* = 0) + \mathbb{P}(\bar{D}_j^* > 0) e^{\lambda r_j^{3DP}}]} \right). \quad (66)$$

It follows that the right-hand side of Eq. (66) decreases with $|\mathcal{S}|$, implying that $\tilde{E}(\mathcal{A}, \lambda)$ is submodular in $\mathcal{A}$. $\square$

Proof of Proposition 5 and Proposition 6. Proposition 5 has already been established in Lemma 5, while Proposition 6 directly follows from the reformulation of V^{3DP} in Lemma 5. $\square$

Appendix B: Implementation Details

In this section, we provide the following algorithms whose main steps are omitted:

- Benders decomposition for solving MIO (12), introduced in Section 4.2.
- Local search algorithm for the heuristic scheme Algorithm 1 outlined in Section 5.2.
- Projected SGD for U^{3DP} under fixed $\mathcal{A}$, introduced in Section 5.4).

Algorithm 3 Benders Cut Generation for Algorithm 4

Require: From Algorithm 4, input $\mathcal{A}(t), K(t), \mathbf{x}(t), \mathbf{q}^{3DP}(t, i), \bar{D}^{ti}$ and $(D^i, \mathbf{s}^i)$ for all $i = 1 \dots S$.

1: **for** $i=1, \dots, S$ **do**

$$\gamma^{ti} = \begin{cases} -\max \{v_j - c_j^{3DP} \mid \bar{D}_j^{ti} > 0, q_j^{3DP}(t, i) < \bar{D}_j^{ti}, x_j(t) = 1\} & \text{if } \sum_{j \in \mathcal{A}(t)} q_j^{3DP}(t, i) = K(t) \\ 0 & \text{otherwise.} \end{cases}$$

2: **for** $j = 1, \dots, |\mathcal{A}|$ **do**

$$\begin{aligned} & (\alpha_j^{ti}, \beta_j^{ti}, \sigma_j^{ti}) \\ &= \begin{cases} (0, 0, 0) & \text{if } x_j(t) = 0 \\ (h_j M, h_j s_j^i, -h_j(D_j^i + M)) & \text{if } x_j(t) = 1, \bar{D}_j^{ti} \leq 0 \\ (v_j M, -v_j s_j^i, v_j(D_j^i - M)) & \text{if } x_j(t) = 1, \bar{D}_j^{ti} > 0, q_j^{3DP}(t, i) < \bar{D}_j^{ti} \\ (- (\gamma^{ti} - c_j^{3DP}) M, (\gamma^{ti} - c_j^{3DP}) s_j^i, - (\gamma^{ti} - c_j^{3DP})(D_j^i - M)) & \text{otherwise.} \end{cases} \end{aligned}$$

3: **end for**

4: **end for**

5: **return** $(\gamma^{ti}, \alpha^{ti}, \beta^{ti}, \sigma^{ti})$ for $i = 1, \dots, S$.

Algorithm 4 Benders Decomposition for MIO (12)

Require: A collection of scenarios $\{(s^i, D^i) : i = 1 \dots S\}$ each assigned with probability P_i . Stopping criteria $\varepsilon > 0$.

1: Initialize with some $K(0) \geq 0, \mathbf{x}(0) \in \{0, 1\}^n$ and $\mathbf{q}(0) \geq 0$. Set $\text{UB}(0) = \infty$ and $\text{LB}(0) = -\infty$.

2: **for** $t = 0, 1, \dots$ **do**

/— — — *Solve Master Problem* — — — /

3: Let $\{(\gamma^\tau, \alpha^\tau, \beta^\tau, \sigma^\tau) : \tau \leq t-1\}$ be given so that we can solve the master problem

$$\begin{aligned} & \min_{\mathbf{x} \in \{0, 1\}^n, K \geq 0, \mathbf{q}, \theta} C^{3DP}(K) + \sum_{j=1}^n (1 - x_j) U_j^{\text{DB}} + \sum_{j=1}^n c_j q_j \mathbb{E} s_j + \theta \\ & \text{s.t.} \quad 0 \leq q_j \leq M x_j, \quad \forall j \in \mathcal{N} \\ & \quad \sigma^\tau + \gamma^\tau K + \alpha^\tau + \mathbf{q}^\top \beta^\tau \leq \theta, \quad \forall \tau = 1 \dots t-1 \end{aligned} \tag{67}$$

4: Denote its optimal solution as $(K(t), \mathbf{x}(t), \mathbf{q}(t), \theta(t))$ and its optimal value as $\text{LB}(t)$.

/— — — *Update Global Upper Bound* $\text{UB}(t)$ — — — /

5: Let $\mathcal{A}(t) \equiv \{j \in \mathcal{N} : x_j(t) = 1\}$ and denote $\bar{D}_j^{ti} \equiv D_j^i - q_j(t) s_j^i$.

6: **for** $i = 1, \dots, S$ **do**

7: Apply Algorithm 3 with input $\mathcal{A} = \mathcal{A}(t), K = K(t)$ and $y_j = \bar{D}_j^{ti}$.

8: Denote the output $\bar{\mathbf{q}}^{3DP}$ as $\mathbf{q}^{3DP}(t, i)$ and the output V^{3DP} value as V_{ti}^{3DP} .

9: **end for**

10: Update $\text{UB}(t) = \min \left\{ \text{UB}(t-1), C^{3DP}(K(t)) + \sum_{j=1}^n [(1 - x_j(t)) U_j^{\text{DB}} + c_j \mathbb{E} s_j q_j(t)] + \sum_{i=1}^S P_i V_{ti}^{3DP} \right\}$.

/— — — *Check Termination* — — — /

11: **if** $\text{UB}(t) - \text{LB}(t) < \varepsilon$ **then**

12: Terminate the algorithm.

13: **else**

/— — — *Add a New Benders Cut* — — — /

14: Apply Algorithm 3 to compute $(\gamma^{ti}, \alpha^{ti}, \beta^{ti}, \sigma^{ti})$ for $i = 1, \dots, S$ and denote

$$\gamma^t \equiv \sum_{i=1}^S P_i \gamma^{ti}, \quad \alpha^t \equiv \sum_{i=1}^S P_i \alpha^{ti}, \quad \beta^t \equiv \sum_{i=1}^S P_i \beta^{ti}, \quad \sigma^t \equiv \sum_{i=1}^S \sum_{j=1}^n P_i \sigma_j^{ti}.$$

15: **end if**

16: **end for**

17: **return** The latest $K(t), \mathbf{q}(t), \mathbf{x}(t)$ at termination.

B.1. Benders Methods in Section 4.2

We outline the Benders decomposition for the mixed-integer program (12), beginning with a subroutine to generate Benders cuts for the main algorithm.

B.2. Local Search Algorithm in Section 5.2

In our heuristic framework for optimizing $\mathcal{A}$ and K (Algorithm 1), one of the two key steps is approximating the “inner-min” problem in Problem (18) by replacing U^{3DP} with L^{3DP} :

$$\min_{\mathcal{A} \subseteq \mathcal{N}} \left\{ L^{3DP}(\mathcal{A}, K) + \sum_{j \in \mathcal{A}^c} (C_j^{DB} + U_j^{DB}) \right\}. \quad (68)$$

Here, we detail an efficient solution approach, starting with a reformulation of Problem (68). This reformulation interprets the objective as the firm’s total profit, which we reasonably assume to be positive—a common requirement for approximation guarantees in combinatorial optimization. Specifically, we obtain it with the following steps:

$$\begin{aligned} & -L^{3DP}(\mathcal{A}, K) - \sum_{j \in \mathcal{A}^c} (C_j^{DB} + U_j^{DB}) + \sum_{j \in \mathcal{N}} v_j \mathbb{E}[D_j] \quad \leftarrow \text{Reverse the sign and add term } \sum_{j \in \mathcal{N}} v_j \mathbb{E}[D_j] \\ &= -\sum_{j \in \mathcal{A}} U_j^0 + E(\mathcal{A})K - \sum_{j \in \mathcal{N}} (C_j^{DB} + U_j^{DB}) + \sum_{j \in \mathcal{A}} (C_j^{DB} + U_j^{DB}) + \sum_{j \in \mathcal{N}} v_j \mathbb{E}[D_j] \quad \leftarrow \text{Apply the specific form of } L^{3DP} \\ &= \sum_{j \in \mathcal{N}} (v_j \mathbb{E}D_j - C_j^{DB} - U_j^{DB}) + \sum_{j \in \mathcal{A}} (C_j^{DB} + U_j^{DB} - U_j^0) + E(\mathcal{A})K \quad \leftarrow \text{Re-arrange terms} \end{aligned} \quad (69)$$

where we let $E(\mathcal{A})$ to be defined as follows (recall that this is submodular in $\mathcal{A}$ from our analysis of Theorem 5)

$$E(\mathcal{A}) \equiv \mathbb{E}_{D,s} \left[\max_{j \in \mathcal{A}} (v_j - c_j^{3DP}) \mathbb{I} \{ [D_j - \bar{q}_j^0 s_j]^+ > 0 \} \right].$$

Additionally, denote:

$$I \equiv \sum_{j \in \mathcal{N}} (v_j \mathbb{E}D_j - U_j^{DB} - C_j^{DB}) - C^{3DP}(K), \quad \Delta_j \equiv U_j^{DB} + C_j^{DB} - U_j^0. \quad (70)$$

Finally, we arrive at the following reformulation of Problem (68), which maximizes a non-monotone positive-valued submodular function:

$$\max_{\mathcal{A} \subseteq \mathcal{N}} I + \sum_{j \in \mathcal{A}} \Delta_j + E(\mathcal{A})K \quad (71)$$

The algorithm shown below is known to terminate with a solution $\hat{\mathcal{A}}$ whose objective value is at least $\frac{1}{3}$ of the optimal value of Problem (71) (see (Feige et al. 2011)).

Algorithm 5 Approximate Optimal Policy $\mathcal{A}$ under Fixed $K > 0$

Require: $K > 0$, $\{\Delta_j\}_{j=1}^n$ (from Eq. (70)), initial $\mathcal{A}_0$, and stopping criterion $\epsilon > 0$. Initialize with $t = 0$.

- 1: **while** true **do**
- 2: If there exists $i \in \mathcal{N} \setminus \mathcal{A}_t$ such that

$$K [E(\mathcal{A}_t \cup \{i\}) - E(\mathcal{A}_t)] + \Delta_i > \frac{\epsilon}{n^2} \left(I + \sum_{j \in \mathcal{A}_t} \Delta_j + E(\mathcal{A}_t)K \right)$$

- then update $\mathcal{A}_{t+1} = \mathcal{A}_t \cup \{i\}$ and repeat this step with $t \leftarrow t + 1$. Otherwise, proceed to what follows.
- 3: If there exists $i \in \mathcal{A}_t$ such that:

$$K [E(\mathcal{A}_t \setminus \{i\}) - E(\mathcal{A}_t)] - \Delta_i > \frac{\epsilon}{n^2} \left(I + \sum_{j \in \mathcal{A}_t} \Delta_j + E(\mathcal{A}_t)K \right)$$

- then update $\mathcal{A}_{t+1} = \mathcal{A}_t \setminus \{i\}$, set $t \leftarrow t + 1$, and continue the loop. Otherwise, exit.
 - 4: **end while**
 - 5: **return** The latest $\mathcal{A}_t$ at termination.
-

B.3. Projected SGD in Section 5.4

We present the pseudo-code for the projected SGD method outlined in Section 5.4. For generality, the algorithm we choose to present here also optimizes K for Problem (1). If only $U^{3DP}(\mathcal{A}, K)$ under a fixed K is of interest, this can be achieved by setting $K(t) \equiv K$ in the pseudo-code.

Throughout the algorithm, we assume access to a data sequence $\{(s^t, D^t) : t \geq 0\}$, either sampled offline from the known distribution of (s, D) or collected in real-time. Additionally, C^{3DP} is assumed convex with subdifferential ∂C^{3DP} .

Algorithm 6 Projected SGD: evaluating U^{3DP} and optimizing K under fixed $\mathcal{A}$

Require: Initialize $q(0) \geq 0$ and $K(0) \geq 0$, with step size $\rho_t > 0$ at step t . Set $t = 0$.
1: **while** true **do**
2: Compute $q^{3DP}(t)$ as the output $\bar{q}^{3DP}$ of Algorithm 2 with input $y = D^t = q(t) \circ s^t$.
3: Compute $\eta^y(t)$ and $\eta^K(t)$ as the η^y and η^K defined in Proposition 5 under $\bar{q}^{3DP} = q^{3DP}(t)$, $y = D^t - q(t) \circ s^t$, and $K = K(t)$.
4: **for** $j = 1, \dots, |\mathcal{A}|$ **do**
5: Update $q_j(t+1) = [q_j(t) - \rho_t (c_j \mathbb{E} s_j - \eta_j^y(t) s_j^t)]^+$.
6: **end for**
7: Update $K(t+1) = [K(t) - \rho_t (\gamma^t + g^t)]^+$ where $g^t \in \partial C^{3DP}(K(t))$ is a subgradient of C^{3DP} .
8: Terminate when a stopping rule is met at step t , otherwise proceed to step $t \leftarrow t + 1$.
9: **end while**
10: **return** The latest $q(t)$ and $K(t)$ at termination.

Appendix C: Numerical Experiment Details

C.1. Additional Calibration Details

C.1.1. 3DP Costs. The unit sourcing cost for product j is estimated as $\kappa_j = 0.006w_j$, where w_j is the weight in grams. This calculation assumes a plastic material cost of \$0.005/gram, marked up by 1.2 to account for additional costs.

As mentioned in the main text, we let $c^{cap} = \frac{Q^{3DP}}{M^{3DP}}$, where Q^{3DP} and M^{3DP} represent the monthly per printer depreciation cost (in \$) and the monthly material output per printer (in grams). Based on a 10-year life-span and current printer costs, we let $Q = \$41.67$. The monthly material output M^{3DP} is calculated as $M^{3DP} = \delta^{3DP} S^{3DP} W^{3DP} H^{3DP} T / \rho^{3DP}$, where W^{3DP} is the nozzle width (typically 0.4 cm), H^{3DP} is the layer thickness (set to 0.01 cm), S^{3DP} is the nozzle movement speed (in cm/s), T is the total number of seconds in a month, and δ^{3DP} is the density of the non-metal printing material (typically 1 gram/cm³). This formula computes the total weight of material a printer can output during continuous operation for a month, adjusted by the infill density $\rho^{3DP} = 0.5$, to account for the fact that 3D-printed objects are typically not solid.

C.1.2. Dedicated Backup Costs. We assume the sourcing cost from a dedicated backup is $1.5 \times$ the unit sourcing cost from a primary supplier, ensuring it is a) cheaper than printing and b) more expensive than primary sourcing. We similarly let C_j^{DB} be 75% of the primary sourcing cost times the expected demand. The choice of 75% reflects conversations with industry professionals that reserving capacity in high-volume, low-margin products usually entails a large retainer fee to offset the lost revenue the backup supplier could have earned by serving another firm.