Digital Twin-Driven Supply Chain Enhancement to Support Direct-to-Consumer Growth

by

Siddhant Agrawal
B.A., Applied Mathematics
Harvard University, 2019

Submitted to the MIT Sloan School of Management and
Department of Civil and Environmental Engineering
in partial fulfillment of the requirements for the degrees of

Master of Business Administration
and
Master of Science in Civil and Environmental Engineering

in conjunction with the Leaders for Global Operations program at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
May 2024

Author . . . .
Siddhant Agrawal
MIT Sloan School of Management and Department of Civil and Environmental Engineering
May 30, 2024

Certified by . . . .
David Simchi-Levi, Thesis Supervisor
Professor of Engineering Systems

Certified by . . . .
Stephen C. Graves, Thesis Supervisor
Professor of Operations Management

Certified by . . . .
Vivek F. Farias, Thesis Supervisor
Professor of Operations Management

Accepted by . . . .
Heidi Nepf
Donald and Martha Harleman Professor of Civil and Environmental Engineering

Accepted by . . . .
Maura Herson
Assistant Dean, MBA Program, MIT Sloan School of Management

This page is intentionally left blank.

Digital Twin-Driven Supply Chain Enhancement to Support Direct-to-Consumer Growth

by

Siddhant Agrawal

Submitted to the MIT Sloan School of Management and
Department of Civil and Environmental Engineering
on May 30, 2024, in partial fulfillment of the
requirements for the degrees of
Master of Business Administration
and
Master of Science in Civil and Environmental Engineering

Abstract

In response to the rising trend of Direct-to-Consumer (D2C) sales, many traditional retailers, which have historically relied on wholesale business models, are now undertaking significant supply chain transformations. This thesis explores the strategic shift of a large retailer in the footwear and apparel sector, pseudonymously referred to as Iota in this study, as it transitions towards a D2C-focused supply chain. This transition, emblematic of a broader industry transformation, is aimed at enhancing alignment with the evolving expectations of customers in terms of service, cost-effectiveness, and sustainability.

Central to this research are the proposed enhancements by Iota's leadership to decentralize Iota's supply chain. These enhancements include adding both physical infrastructure, with the planned establishment of a cross-dock facility, and digital infrastructure, through the development of a decision engine that aids in efficiently routing products within the new decentralized supply chain network. The cross-dock facility is envisioned to provide an opportunity for decision postponement in the inventory flow from Asian factories to US distribution centers. Meanwhile, the decision engine, leveraging a heuristic-based algorithm, is set to unlock new inventory flows and enhance inventory distribution.

With the new infrastructure to decentralize the supply chain yet to be fully operational, a retrospective study was conducted using a digital twin of Iota's supply chain. Various push and pull-based inventory deployment strategies were simulated in the digital twin with the goal of alleviating pressure on the primary distribution center and increasing fulfillment from regional distribution centers. In the simulation process, challenges with forecast data and lumpiness of supply are discovered and subsequently addressed through the use of synthetic datasets, which emulate improved forecast coverage and smooth supply.

The key findings from simulations highlight that despite achieving a modest performance in meeting the goals for the decentralized network, valuable insights were obtained that could drive future supply chain enhancements. The research underscores the benefits of smoothing supply for network performance, the critical role of comprehensive and reliable forecast data, and the necessity for supplementary storage solutions to complement the cross-dock facility. For example, one pull-based scenario using a synthetic dataset to emulate enhanced forecast coverage and smoother supply tripled network performance while reducing network costs by 1% compared to the baseline pull-based scenario. Such cost savings could be substantial for a large-scale retailer.

Concluding with recommendations, the thesis advises Iota to re-evaluate purchasing practices, consider integrating multiple internal sources of forecast data into a single source, and continue with simulation analyses. These recommendations are designed to support Iota, and by extension, similar retailers, in their transition towards a robust and agile D2C supply chain, ensuring competitive advantage in the dynamic retail sector.

Thesis Supervisor: David Simchi-Levi
Title: Professor of Engineering Systems

Thesis Supervisor: Stephen C. Graves
Title: Professor of Operations Management

Thesis Supervisor: Vivek F. Farias
Title: Professor of Operations Management

This page is intentionally left blank.

Acknowledgments

I am profoundly grateful to my advisors David Simchi-Levi, Stephen Graves, and Vivek Farias for their invaluable guidance and support throughout this thesis. My gratitude also extends to my project managers at the host company for their insights and contributions to my research. The collaborative environment and resources provided by the host company were indispensable. On a personal note, my heartfelt thanks go to my mother, dad, sister, and my partner for their constant encouragement and support. Their belief in me has been a constant source of strength and motivation. To all, I offer my sincere thanks.

Note on Proprietary Information

To safeguard the host company's proprietary information, the data presented in this thesis has been adjusted to reflect relative, rather than absolute, values. Furthermore, the diagrams and data labels have been intentionally modified to preserve confidential competitive information. Therefore, they should be considered as representative illustrations and not as exact representations of real data.

Table of Contents

List of Figures

Acronyms

Acronym	Description
BOP	Beginning of Period. Start of a two-week decision period in the simulation.
CPU	Cost Per Unit.
CxDF	Cross-Dock Facility. A pass-through facility with no on-site storage.
E2EDT	End-To-End Digital Twin. A high-fidelity simulation environment for supply chain decision-making.
EOP	End of Period. End of a two-week decision period in the simulation.
GAC	Goods at Consolidator. A facility in Asia where the products from the manufacturers are assembled and purchased by Iota to be sent to the US.
IAS	Inventory Allocation Software. A decision engine to allocate inventory in first- and middle-mile of Iota's decentralized supply chain network.
JIT	Just in Time. An inventory deployment strategy that minimizes storage needs.
MVP	Minimum Viable Product. A software development approach with the goal of getting a product to users as quickly as possible to get feedback to iterate upon.
OST	Off-Site Storage. A facility to store excess inventory arriving at the cross-dock facility.
PDC	Primary Distribution Center. The main distribution center, originally designed to primarily serve the wholesale business.
PO	Purchase Order.
RDC	Regional Distribution Center. There are three of them in operation currently, and they are used to fulfill online orders only.

Chapter 1: Introduction

The retail landscape is experiencing a significant shift with the rise of the Direct-to-Consumer (D2C) sales channel. This channel encompasses both online customer orders and transactions from physical storefronts. This paradigm shift poses considerable challenges for many retailers whose supply chains were initially designed for wholesale operations. As such retailers transition their supply chains – traditionally centralized for infrequent, bulk shipments – to cater to the burgeoning D2C demand, which requires frequent, smaller shipments, a substantial reconfiguration becomes imperative. This involves moving from large-scale distribution to a more decentralized, agile system. The improved supply chain network must be capable of distributing smaller quantities nationwide cost-effectively, promptly, and sustainably. This restructuring is crucial as D2C sales are becoming an increasingly vital revenue source, and the existing supply chain models are ill-equipped to handle this change.

In this evolving landscape, companies are grappling with uncertainties as they consider making significant investments to adapt their supply chains. The advent of digital twins – virtual replicas of physical entities – has shown potential in aiding decision-making for these transitions. These digital tools enable simulations with high fidelity, allowing businesses to assess the impact of supply chain network changes on various metrics such as network costs, fulfillment time, and carbon footprint, thereby guiding decision-making processes.

This study conducts an empirical assessment of various supply chain network configurations for a large footwear and apparel retailer, anonymously named Iota in this thesis. It utilizes a digital twin to navigate the supply chain network's restructuring to achieve specific supply chain objectives amid growing D2C demand. The thesis offers insights into the application of digital twins in strategic decision-making and sheds light on the challenges encountered by traditionally wholesale-centric retailers transitioning to accommodate D2C requirements. The findings are pertinent to other industry players considering a similar shift.

1.1. Overview of Iota's Supply Chain

Iota's supply chain, initially tailored to accommodate the wholesale business model – which constituted the bulk of its operations – was centralized for efficiency. The Primary Distribution Center (PDC) in the South, once the world's largest of its kind at inception, served as the core hub for all products shipped from Asian factories. These products were typically dispatched in bulk from the PDC to wholesale partners.

However, over the past decade, D2C sales have surged, prompting Iota to begin distributing products directly to consumers from the PDC. It became evident that the PDC, originally designed for wholesale order fulfillment, was not optimally configured for handling D2C orders. To address this during the Covid-19 pandemic, Iota established three Regional Distribution Centers (RDCs) aimed at processing digital orders more effectively. These RDCs, located on the West Coast, East Coast, and Central US, played a crucial role in managing the spike in D2C sales during the pandemic.

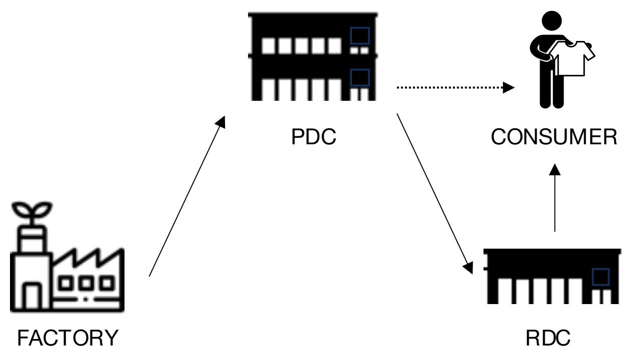

- RDC = Regional Distribution Center
- PDC = Primary Distribution Center

Figure 1: Iota's Current Supply Chain Network

Despite these adjustments, Iota's distribution network largely remains centralized as seen in Figure 1. All products from Asia initially enter through the port on the West Coast and are then transported to the PDC. Upon arrival at the PDC, these items undergo a distribution assessment. Based on specific needs, a portion of this inventory is redirected to the RDCs to fulfill digital orders. However, if a particular item required for an online order is not available at the corresponding RDC, it is then shipped directly to the customer from the PDC, subject to its availability. Such a centralized approach in the distribution network leads to operational inefficiencies. For example, consider a digital order from a West Coast customer. The item, if fulfilled from the West Coast RDC, would travel an extended route from Asia to the West Coast port, onward to PDC, and then back to the RDC on the West Coast. Such centralized operations – which likely result in increased costs, fulfillment time, and carbon emissions – are increasingly becoming a challenge in the context of Iota's steep D2C growth.

1.2. Project Drivers and Objectives

Iota is in the process of transforming its supply chain from a wholesale-focused model to a D2C framework, aiming to align with evolving customer expectations regarding service, cost, and sustainability. Iota is considering additional investments in both physical infrastructure and logical infrastructure to decentralize its network to support efficient fulfillment of its digital orders.

On the physical infrastructure front, Iota is considering establishing a cross-dock facility (**CxDF**) on the West Coast to intercept and strategically redirect inventory arriving from Asian factories. Instead of routing all goods to the PDC, the CxDF will provide a decision postponement opportunity and allow inventory to be sent directly to the RDCs or the PDC depending on the need.

On the digital infrastructure side, Iota is developing a new decision engine, which we will refer as Inventory Allocation Software (**IAS**) going forward in this thesis. This engine is tasked with managing the distribution of inventory across the first- and middle-mile of Iota's evolving decentralized network. In particular, it determines the allocation of inventory across PDC, CxDF, and RDCs at factories in Asia (first-mile) as well as the allocation of inventory arriving at the CxDF across RDCs and PDC (middle-mile). The IAS operates on a heuristic-based model[1] and is built upon a third-party software, anonymously referred to as SupplySoft in this thesis, which Iota is currently licensing. At the time of this study, the implementation of the **IAS/CxDF solution** was still in the preparatory phase and had not yet been fully operational.

Iota's leadership has established key strategic goals to enhance the efficiency of their supply chain. One such target is to maximize in-region fulfillment i.e., increase the percentage of digital orders fulfilled directly through the RDCs, as opposed to being dispatched from the PDC. Additionally, they aim to increase the share of inventory received by the RDCs from sources other than the PDC, thereby reducing the reliance and strain on the PDC to supply the RDCs' inventory needs. These other sources of inventory flows into the RDCs include direct shipment from factories in Asia and inventory arriving at the CxDF. The enhancements and investments being made in both the physical and digital aspects of Iota's infrastructure are directed towards reaching these supply chain targets.

The primary aim of this research was to evaluate the impact of the IAS/CxDF solution on attaining the strategic goals outlined by Iota's leadership. In addition to the strategic targets, other conventional supply chain metrics like inventory turns, network costs, and unutilized inventory at the CxDF are also used in the evaluation. Initially, the study sought to confirm that the introduction of the IAS/CxDF solution would not detrimentally affect the current performance of the supply chain, ensuring that these new infrastructural elements did not lead to a decline in the network performance. Following this assessment, the focus shifted towards exploring ways to improve the effectiveness of the IAS, with the objective of getting closer to Iota's strategic targets.

[1] It must be emphasized that IAS's algorithm is heuristics-driven and does not explicitly optimize an objective function unlike other decision engines used at Iota.

1.3. Overview of Approach

Given that the **IAS/CxDF solution** was not operational during the research period, obtaining concurrent empirical data to assess its impact was not feasible. Consequently, a retrospective study leveraging simulation results was used to study the performance of the IAS/CxDF solution.

Iota has built a sophisticated simulation platform, referenced in this thesis as **End-To-End Digital Twin (E2EDT)**, designed to offer a high-fidelity environment for testing various supply chain decisions and their implications on network performance. E2EDT integrates several decision engines, including ones that simulate product manufacturing across different locations to evaluate the consequences of production shifts. For distribution strategies, it utilizes two main decision engines that are responsible for determining the movement of inventory from the PDC to the RDCs (covering the middle-mile logistics) and deciding the best source for fulfilling orders, whether from the PDC or RDCs (covering the last-mile logistics). Unlike IAS, which leverages a heuristics-based model, these other decision engines already integrated into the E2EDT are set up to optimize an explicit objection function that considers a combination of network costs, service-level, and sustainability performance.

Ideally, the IAS/CxDF solution would have been integrated with E2EDT to directly simulate its impact and explore avenues for enhancing network performance. However, due to the third-party nature of the IAS/CxDF solution and its evolving status, integrating it into E2EDT was outside the scope of this project. The challenge was compounded by the substantial changes the IAS/CxDF solution introduced to Iota's distribution model, such as diversifying the first-mile delivery routes to include the CxDF, RDCs, or PDC, and reconfiguring the middle-mile logistics to accommodate new network flows enabled by the CxDF. These significant alterations necessitated more time for full integration of IAS into E2EDT.

To address these obstacles, the project entailed creating a streamlined digital twin of Iota's supply chain distribution, which will be referred to as **digital twin** moving forward in this thesis. This digital twin emulated the heuristic logic of the IAS for first- and middle-mile logistics and used West greedy algorithms to perform the roles of the other distribution decision engines, namely replenishing inventory at the RDCs and fulfilling consumer orders, used in the E2EDT. **This digital twin – a streamlined, standalone simulation environment – served as the basis for scenario analysis.** The scenario analysis aimed to evaluate and enhance network performance by simulating the hypothetical availability of the IAS/CxDF solutions during a recent historical timeframe. It involved varying model parameters (discussed later) to uncover opportunities for improving the IAS's performance.

Overview of Thesis Structure

This thesis is structured as follows: Chapter 2 delves into the construction of the digital twin, with a focus on two key decisions made by IAS: the allocation of shipments from the port across PDC, CxDF, and RDCs, and allocation of inventory arriving at the CxDF across PDC and RDCs. This chapter also examines a streamlined approach to middle-mile replenishment and last-mile fulfillment employed in the digital twin and concludes with a discussion on the validation process of the digital twin. Chapter 3 explores the various metrics employed to evaluate the supply chain network performance. These metrics extend beyond the two strategic targets, encompassing additional factors such as network costs and inventory turnover at the RDCs. In Chapter 4, I outline the three main categories of push and pull-based scenarios simulated using the digital twin. This chapter also identifies two critical challenges discovered in enhancing the IAS/CxDF solution's performance: the lumpy nature of Iota's supply and the suboptimal forecast coverage. Chapter 5 presents the conclusion and offers recommendations for enhancing performance of the IAS/CxDF solution and more effective supply chain management.

Chapter 2: Digital Twin of Iota's Supply Chain

In collaboration with Iota, a digital twin of Iota's supply chain was built. The digital twin models the distribution leg of the supply chain and focuses on replicating the decisions made by the IAS engine in the first- and middle-mile of the network. Some of the inventory allocation decisions in the middle- and last-mile logistics were simplified in the digital twin due to time constraints. The digital twin replicates two critical decisions made by the IAS. Firstly, it determines the inventory levels to be dispatched from factories to the RDCs, CxDF, and PDC. Secondly, it determines the allocation of volume between RDCs and PDC upon their arrival at the CxDF. These decisions are made on a bi-weekly basis, corresponding to a two-week time frame.

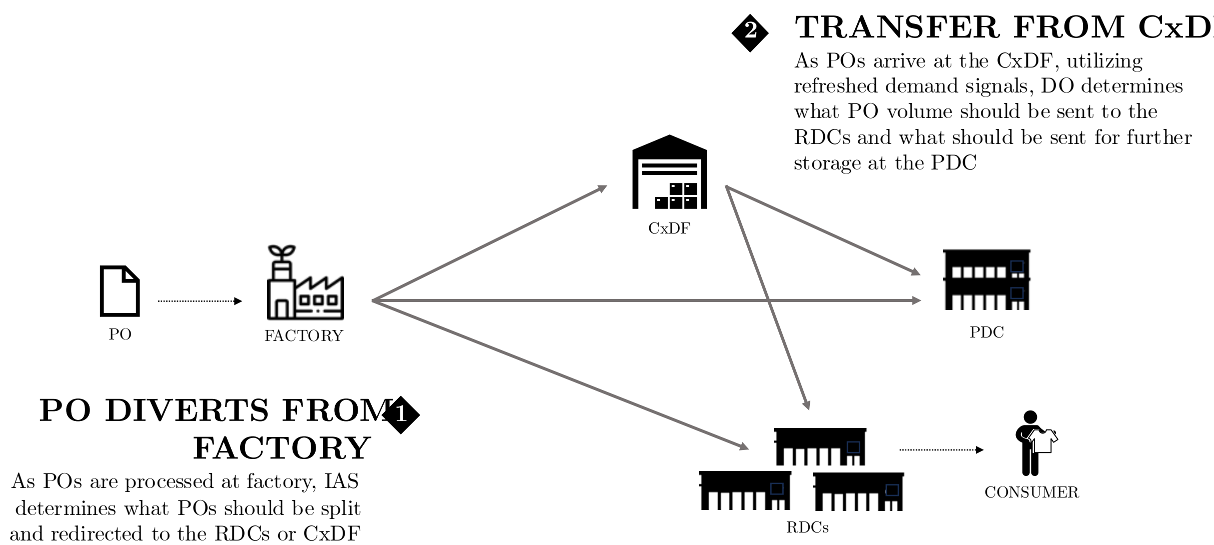

Figure 2: Key Decisions Made by Inventory Allocation Software (IAS)

The supply side is modeled based on historical POs available during the simulation timeframe. A PO is essentially an order placed by Iota with its contract manufacturers, specifying the total quantity of a SKU to be produced, along with the expected delivery date. While POs are designated for either digital or wholesale channels, it is noteworthy that non-digital POs may occasionally fulfill digital orders. This simulation exclusively considers POs allocated for digital orders on the supply side.

The unconstrained nature of the model enables us to assess the upper bounds for the performance improvements resulting from IAS/CxDF. Except for the supply constraint determined by historical PO volume when matching demand and supply, the model assumes unlimited storage and processing capacity at each of the nodes in the network.

The simulation covers demand from February 2022 to March 2023 for digital orders only. It can be run at the SKU level (defined at product-size level), and the simulations across the SKUs are independent. The simulation covers approximately 32 million[2] units, encompassing thousands of SKUs, shipped.

2.1. IAS Decision I: PO Diverts from Factory

Without the IAS/CxDF solution, all digital inventory from the factory is sent to the PDC. From the PDC, the inventory is either replenished at the RDCs to fulfill digital orders, or it is sent directly to the consumer from the PDC if inventory is not available at the RDCs to fulfill digital demand.[3] Following the implementation of the IAS/CxDF solution, inventory can be routed more efficiently from the factory.

The first decision made by IAS is to allocate supply from the factory across RDCs, CxDF and PDC as depicted in Figure 2. To comprehend how IAS makes this decision, one must consider the demand and supply dynamics of a specific SKU. The available supply for a SKU on a particular decision date is ascertained by the historical POs eligible[4] for diversion from the factory. It is common for there to be multiple eligible POs for an SKU, with the collective units from these POs contributing to the SKU's total supply. Moreover, the model operates under the assumption that no inventory is present at the PDC, CxDF, or RDCs to satisfy demand, meaning supply for this IAS decision is strictly defined by the available POs.

On the demand side, the model utilizes the long-range forecast of an SKU with a 13-week lead time, which is the anticipated duration for goods to travel from factories in Asia to the distribution centers in the US and become available to meet consumer demand. Historical forecast data snapshots are leveraged to establish the long-range forecast for each SKU. During this IAS decision-making process, any safety stock requirements are disregarded, and only the forecasted demand is considered.

Following the establishment of long-range demand forecast, it must be disaggregated to distribute the demand across the RDCs, CxDF, and PDC. This disaggregation process is further explained in section 2.1.1. With granular demand and supply data at hand, a greedy heuristic

---

[2] During this period, the actual digital sales were much higher. However, the simulation does not include sales from specialized distribution centers, orders fulfilled by nodes that are not yet distribution centers, launch products, or digital orders fulfilled by POs that were allocated for non-digital sales.
[3] The fulfillment decision is made by another decision engine, which may prioritize fulfilling from the PDC even if the inventory is available at the RDCs, in order to meet the overall objective, which includes minimizing split shipments.
[4] Eligible POs to be diverted are determined based on their expected arrival dates at the consolidator in Asia. The model assumes a 5 week lead time between decision date and arrivate date at the consolidator.

algorithm is employed to match demand with supply. The specifics of this algorithm are detailed in section 2.1.2.

The IAS in the digital twin runs on a bi-weekly cycle, assessing the demand and supply for each SKU within a two-week span. For example, during the first two weeks (weeks 0 and 1) of the simulation, the demand forecast for weeks 13 and 14 informs the decision-making process. The subsequent decision period covers weeks 2 and 3, which considers the demand forecast for weeks 15 and 16, and this pattern continues throughout the simulation. Figure 3 presents a detailed timeline of the demand and supply considered for each IAS decision period.

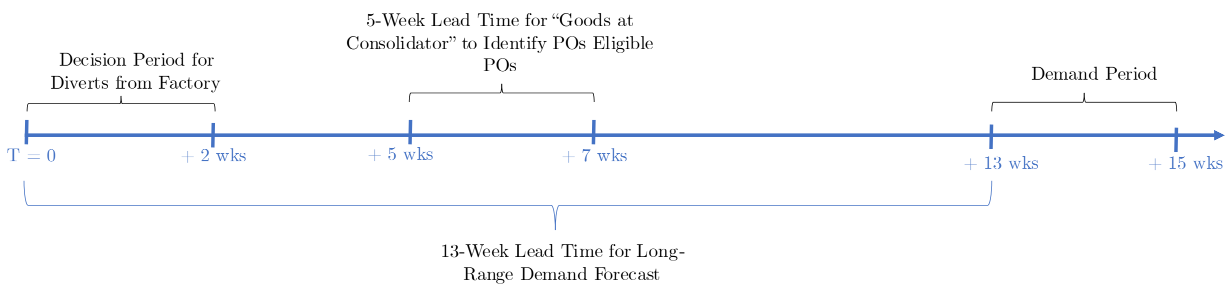

Figure 3: Decision Timeline in IAS Decision I

2.1.1. Disaggregating Long-Range Demand Forecast

The long-range forecast (13-week lead time), available at product level, is a key input in making the PO split decision at the factory. As mentioned earlier, the inventory on hand or the safety-stock targets are not factored in for this IAS decision. To match forecasted demand with supply (POs), it is necessary to disaggregate the long-range demand forecast to match the level of granularity of the PO data.

We begin by disaggregating the forecast to the product-size level, and subsequently splitting it even further to map to the three geographic regions served by the RDCs. Using Size Curves (an internally developed model that splits demand across sizes for each product) and historical proportion of digital sales across the three RDCs, the long-range forecast at the product level is split to be at the product-size-region level as shown in Figure 4.

Long-range Forecast For Two Week Demand Window

Product Code	Forecasted Quantity
X08720-A10	100

Splitting the Demand using Size Curves

Product Code	Size	Forecasted Quantity
X08720-A10	S (30%)	30
	M (50%)	50
	L (20%)	20

Splitting the Demand across 3 RDCs using Historical Digital Sales

Product Code	Size	RDC Node Code	Forecasted Quantity
X08720-A10	S	RDC 1 (25%)	7.5
		RDC 2 (40%)	12
		RDC 3 (35%)	10.5
	M	RDC 1 (25%)	12.5
		RDC 2 (40%)	20
		RDC 3 (35%)	17.5
	L	RDC 1 (25%)	5
		RDC 2 (40%)	8
		RDC 3 (35%)	7

Figure 4: Long-Range Forecast Data Splitting Process at Product-Size-Region Level (Note: Illustrative Proportions)

Next, to emulate the logic of the SupplySoft Solution (third-party software on top of which IAS was built), the forecasted demand, currently at product-size-region level, is further split into three confidence levels: high, medium, and low. This division is based on the confidence of demand actualization. These confidence levels are determined using an internally developed Flow Strategy model which is explained below.

The Flow Strategy model categorizes products into seven segments based on their forecast volume and accuracy. Then, for each segment, the model defines demand proportions for high, medium, and low categories, illustrated in Figure 5. By mapping products to specific segments, the same set of demand confidence proportions is applied uniformly to all products within each segment. This division results in high, medium, and low confidence proportions for the demand. A more detailed description of the Flow Strategy model can be found in Appendix C.

DEMAND CONFIDENCE PROPORTION

PRODUCT SEGMENT	H	M	L
A	0.4	0.2	0.4
B	0.3	0.3	0.4
C	0.2	0.3	0.5
T	0	0	1

Figure 5: Confidence Proportions in the Flow Strategy Model (Note: Illustrative Purposes)

Next, the Flow Strategy model specifies the lanes for diverting the supply based on forecasted demand. While medium confidence demand across all three RDCs is sent to the CxDF, and low confidence demand is sent to the PDC, high confidence demand is sent directly to the three RDCs. This implies that high confidence demand must be distributed among three nodes, while medium and low confidence demand consolidate at a single node.

To ensure a more equitable distribution of supply among the three RDCs, we further split the high-confidence demand into two[5] segments and queue them (see *H Quenue Position* column) in the order of the segements (a SupplySoft Solution logic) as shown in Figure 6. As explained in the next section, this process enables a more fair allocation of supply across the RDCs, instead of diverting the supply to a single RDC in cases where enough supply is not available to meet demand across all the RDCs.

Product Code	Size	RDC Node Code	Confidence Level	H Queue Position	Forecasted Quantity	Destination Node
X08720-A10	S (30%)	RDC 1 (25%)	H (40%)	1	7.5 x 0.4 x 0.5 = 1.5	RDC 1
				2	7.5 x 0.4 x 0.5 = 1.5	RDC 1
			M (20%)	-	7.5 x 0.2 x 1 = 1.5	CxDF
			L (40%)	-	7.5 x 0.4 = 3	PDC
		RDC 2 (40%)	H (40%)	1	12 x 0.4 x 0.5 = 2.4	RDC 2
				2	12 x 0.4 x 0.5 = 2.4	RDC 2
			M (20%)	-	12 x 0.2 x 1 = 2.4	CxDF
			L (40%)	-	12 x 0.4 x 1 = 4.8	PDC
		RDC 3 (35%)	H (40%)	1	10.5 x 0.4 x 0.5 = 2.1	RDC 3
				2	10.5 x 0.4 x 0.5 = 2.1	RDC 3
			M (20%)	-	10.5 x 0.2 x 1 = 2.1	CxDF
			L (40%)	-	10.5 x 0.4 x 1 = 4.2	PDC

Figure 6: Demand Split by Confidence Level With H-Demand Further Split into Two Segments for Equitable Distribution Across RDCs (Assuming X08720-A10 is in Product Segment A)

2.1.2. Matching Disaggregated Forecasted Demand with Supply

Emulating the SupplySoft Solution logic, a demand queue is established to match supply with the disaggregated forecasted demand as shown in Figure 7. Within this queue, high-confidence demand takes priority, followed by medium-confidence demand, and finally, low-confidence demand. The high-confidence demand segments are sorted by their position in the queue, starting with the first position, and moving on to subsequent ones. Within each high-confidence demand segments with the same queue position, the demand lines with the smallest quantities

[5] The choice of two segments was primarily to minimize runtime. A higher number of segments would enable more equitable distribution of supply across the RDCs when it is not sufficient to meet all forecasted demand.

are fulfilled first (1.5 followed by 2.1 followed by 2.4 within *H Queue Position* 1 in the example below).

Next, case pack requirements are implemented in the demand queue to ensure that only full cases are shipped as per business requirements. The case pack requirement[6] is 6 for footwear and 12 for the first pack of apparel, followed by 6 for subsequent packs. For example, if there is a demand for 10.5 pairs of shoes, they will be shipped in two cases of 6 pairs each. If there is a demand for 12.4 shirts, they will be shipped in one case of 12 shirts, plus another case of 6 shirts.

Product Code	Size	Confidence Level	H Queue Position	Raw Forecasted Quantity	Forecasted Quantity with Case Pack Rounding	Destination Node
X08720-A10	S (30 %)	H	1	1.5	6	RDC 1
		H	1	2.1	6	RDC 3
		H	1	2.4	6	RDC 2
		H	2	1.5	0	RDC 1
		H	2	2.1	0	RDC 3
		H	2	2.4	0	RDC 2
		M	-	6	6	CxDF
		L	-	12	12	PDC

Figure 7: Demand Queue for Supply-Forecasted Demand Matching with Case Pack Rounding (Assuming X08720-A10 is in a Footwear)

After establishing the demand queue, the supply is matched with the demand in a linear manner, considering the available supply. Figure 8 illustrates three scenarios regarding the available supply: when the supply equals the forecasted demand with the case pack requirement, when the supply is less than the demand, and when the supply exceeds the demand.

[6] The case pack requirements varied between adults and kids, which was ignored in the model.

Product Code	Size	Confidence Level	H Queue Position	Raw Forecasted Quantity	Forecasted Quantity with Case Pack Rounding	Destination Node	Supply Allocation (S = 36)	Supply Allocation (S = 12)	Supply Allocation (S = 42)
X08720-A10	S (30%)	H	1	1.5	6	RDC 1	6	6	6
		H	1	2.1	6	RDC 3	6	6	6
		H	1	2.4	6	RDC 2	6	0	6
		H	2	1.5	0	RDC 1	0	0	0
		H	2	2.1	0	RDC 3	0	0	0
		H	2	2.4	0	RDC 2	0	0	0
		M	-	6	6	CxDF	6	0	6
		L	-	12	12	PDC	12	0	18

Figure 8: Matching Supply with Forecasted Demand in the Queue

In the first scenario, we precisely match the forecasted demand at each node with the available supply. However, when the supply is insufficient compared to the demand, we match the supply with the demand until the supply is depleted. This approach showcases the advantage of further splitting the high-confidence demand into two segments and queuing them by segment, as it allows for a fairer allocation of supply across the RDCs. The first segment of of demand is fulfilled at each RDC before moving on to fulfill the second segment of demand for the RDC at the top of the demand queue.

In the case where the supply exceeds the forecasted demand, the surplus is directed to the PDC. For example, even though the demand was only 12 units, 18 units are sent to the PDC. Additionally, if a product has an eligible purchase order (PO) but is not covered in the Flow Strategy, it is sent directly to the PDC.

By the end of this process, we observe that originally, the PO would have been sent directly to the PDC. However, the SupplySoft Solution splits each PO to be sent to a maximum of five destinations: three RDCs, CxDF, and PDC.

2.2. IAS Decision II: Inventory Transfer from the CxDF

The IAS engine's second critical decision involves the allocation of inventory that is diverted from the factory and arrives at the CxDF. Upon arrival, the IAS model distributes the inventory between the RDCs and the PDC, with the CxDF serving as a juncture for decision postponement. This allows for the integration of updated demand signals to inform the allocation of inventory. The simulation spans 21 bi-weekly decision periods, stretching from May 2022 to February 2023.

To understand the IAS's second decision, it is helpful to consder the demand and supply for a SKU. Let's consider the supply side first. Although we made IAS's first decision estimating that POs would take 13 weeks to arrive at the CxDF from the factory, there is variability in the actual arrival dates. Thus, the historical arrival dates of POs are used to determine the actual inventory available at the CxDF. Figure 9 presents the timeline of these elements. For this decision, the inventory available at the RDCs is also considered as part of the supply to fulfill demand.

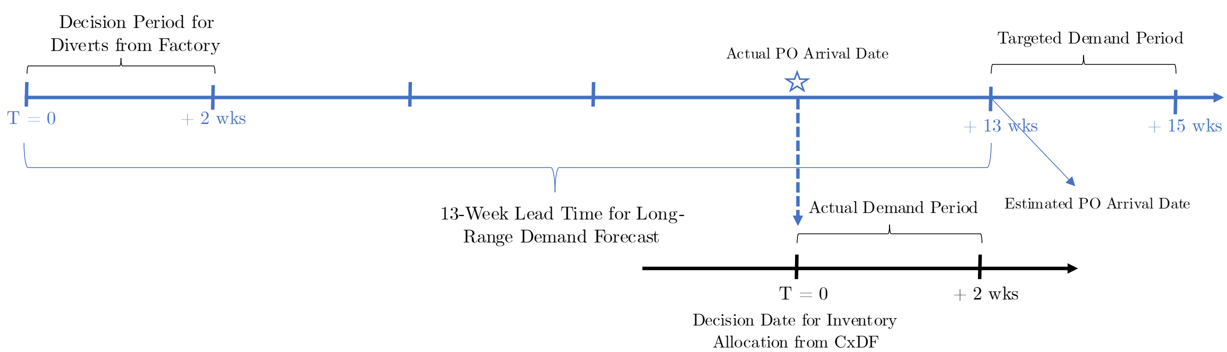

Figure 9: Decision Timeline for IAS Decision II

Regarding demand, the updated short-range demand forecast is taken into account. Moreover, for this IAS decision, safety stock levels are also factored into the supply and demand matching process. In the baseline scenario, safety stock requirements are set to correspond to one week of short-range demand.

The existing inventory at the RDCs is deducted from the short-range demand forecast and the safety stock target to determine the inventory needs at the RDCs. The goal is to fulfill these requirements from the CxDF's available supply. Should the CxDF's supply prove insufficient, the inventory at the PDC is then utilized. For the purposes of this decision, the model assumes zero lead time for transferring inventory from the CxDF to the RDCs/PDC and from the PDC to the RDCs. It also presumes there are no storage or processing capacity constraints at any of the network's nodes. The intricacies of these calculations, along with an illustrative example, are detailed in the forthcoming sub-sections.

2.2.1. Calculating Inventory Required at the RDCs

As of the time product reaches the CxDF and needs to be diverted, the model calculates the inventory required at the RDCs. The first input is the updated forecasted demand for the two-week decision period. To address the limitation of assuming zero lead time, the model uses the demand forecasted one week prior to the decision date (refer to Figure 9). This approach ensures the utilization of more realistic forecast data that would be available for determining the allocation of inventory arriving at the CxDF. The model calculates the safety-stock target, which is set at one week of demand. For simplicity in modeling, the safety stock is determined based on the forecasted two-week demand, rather than considering future projected demand for the third week.

With demand and safety stock calculated, the model considers the inventory available from the last decision period[7] and incoming supply from the factory to determine the inventory required at the RDCs for the decision period. For instance, taking RDC1 in Figure 10, if the forecasted demand for a SKU at an RDC is 30 units, the safety stock target is 15 units, 10 units were in stock at the RDC, and 0 units are arrived directly at the RDC from the factory, then the inventory required for that SKU at the RDCs is 35 (30+15-10-0) units in total, 20 units to meet forecasted demand and 15 units to meet the safety stock target.

Product Code	Size	RDC #	Beginning of Period Inv	Arriving Factory Diverts	Forecasted Demand for 2 Weeks	Safety Stock (SS) Target	Inv Required for Demand	Inv Required for SS
X08720-A10	S	RDC 1	10	0	30	15	20	15
		RDC 2	0	10	10	5	0	5
		RDC 3	110	20	80	40	0	0

Figure 10: Inventory Calculation for Meeting Demand and Safety Stock Targets across 3 RDCs for a SKU

Using this logic, the model calculates the required inventory to meet the forecasted demand and safety stock target for each SKU at the three RDCs.

2.2.2. Matching Demand and Supply at the CxDF and the PDC

The calculated inventory required to meet the forecasted demand and the safety stock target is used to create a demand queue for each SKU in the simulation. In addition to following the priorities applied for creating the demand queue for the divert decision from the factory, this

[7] For the first decision period, the model uses historical snapshot of inventory available. For the subsequent decision periods, the inventory balance at the end of the previous decision period as calculated in the simulation is used.

demand queue also prioritizes inventory required for forecasted demand before inventory required to meet safety stock targets. Recall that while making divert decisions at the factory, we only considered long-range demand forecast, ignoring any safety stock requirements.

Figure 11 shows the demand queue (column *Inventory Required)* for the SKU, which is used to match demand with supply that arrives at the CxDF from the factory. The splitting of high-confidence demand into two segments and queuing them ensures that the supply available at the CxDF is split across the RDCs in an equitable manner. The case pack rounding requirement is applied to the demand queue (column *Inventory Required with Case Pack Rounding)* before matching demand and supply.

Product Code	Size	Inventory Type	H Queue Position	Inventory Required	Inventory Required with Case Pack Rounding	Destination Node
X08720-A10	S	Demand	1	0	0	RDC 2
		Demand	1	0	0	RDC 3
		Demand	1	10	12	RDC 1
		Demand	2	0	0	RDC 2
		Demand	2	0	0	RDC 3
		Demand	2	10	12	RDC 1
		Safety Stock	1	0	0	RDC 3
		Safety Stock	1	2.5	6	RDC 2
		Safety Stock	1	7.5	12	RDC 1
		Safety Stock	2	0	0	RDC 3
		Safety Stock	2	2.5	0	RDC 2
		Safety Stock	2	7.5	6	RDC 1

Figure 11: Demand Queue for Inventory Allocation Decision at the CxDF - Prioritizing Forecasted Demand over Safety Stock

The supply available at the CxDF for the SKU is matched to the demand queue in linear order. Recall that because the CxDF is a cross-dock facility without storage, any inventory arriving at the CxDF must be allocated to the RDCs or PDC. If the supply at the CxDF is more than what is needed to fulfill demand and safety stock targets at the RDCs, then the excess is sent to the PDC.

Finally, if the inventory required at the RDCs cannot be fulfilled using the supply arriving at the CxDF, then the model creates another demand queue to fulfill the leftover requirement using inventory available at the PDC. The leftover requirement is calculated by netting out demand and safety stock requirements that were fulfilled using the inventory arriving at the CxDF.

In the example presented in Figure 11, 42 units are required at RDC 1 and 6 units at RDC 2. If 54 units of the SKU arrive at the CxDF from factory diverts, then the model will send 48 units to the RDCs, with 42 units going to RDC 1 and 6 units going to RDC 2. The excess 6 units (CxDF Overage) is sent to the PDC. However, if only 24 units arrive at the CxDF, then 24 units are sent to RDC 1 since it comes first in the demand queue. The model then tries to fulfill the remaining 24 units required from the inventory available at the PDC.

2.3. Validation of Digital Twin

Iota's existing supply chain, initially designed for wholesale transactions, operates by routing all inventory from the manufacturing sites directly to the PDC. Once at the PDC, this inventory is either directed towards replenishing stock at the RDCs or fulfilling digital (and wholesale) orders directly. To measure the accuracy of the simulation, the current operational model was replicated within the simulation to compare simulated inventory turnover rates at the RDCs with the historical acutal figures.

Simulated vs. Actual Bi-Weekly Inventory Turns at RDC1

Simulated vs. Actual Bi-Weekly Inventory Turns at RDC2

Simulated vs. Actual Bi-Weekly Inventory Turns at RDC3

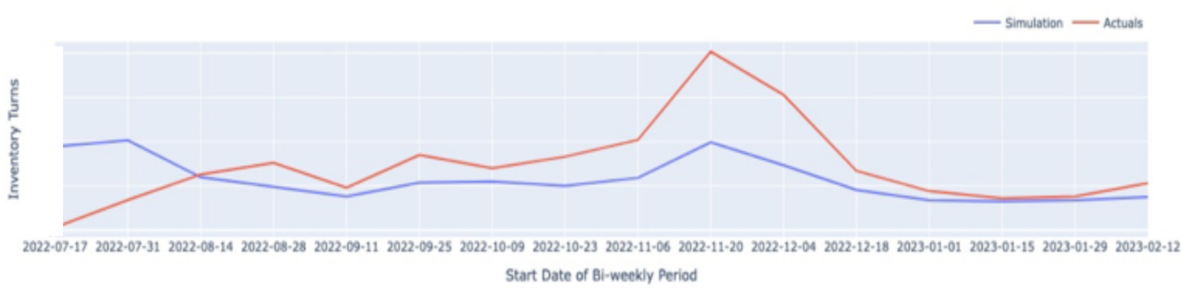

Figure 12: Simulated Status Quo Scenario vs Actual Inventory Turns at RDCs

Figure 12 showcases a side-by-side comparison of inventory turnover rates (calculated as annualized sales/average inventory over the two-week decision period)[8] from the simulation against actual rates for each decision period at the three RDCs. The X-axis presents the start date of two-week decision period and y-axis presents the inventory turns with y-scale hidden to maintain data privacy. The close alignment between the simulated and actual turnover rates is apparent, with the simulation successfully reflecting the spike in turnover during the Black Friday period and the subsequent decrease post-Holidays, mirroring the real-life trend. This congruence between the simulated and actual data affirms the reliability of the model in depicting the nuances of Iota's supply chain concerning digital orders.

[8] See Chapter 3 for a detailed description of how the RDC Inventory Turns are calculated in the simulation

Chapter 3: Metrics for Evaluating Supply Chain Network Performance

The following key metrics are calculated to evaluate the performance of the supply chain network:

- **Inventory Turns at RDCs**: The inventory turns at the RDCs represent the number of times the inventory at the RDCs is replaced within a year. This metric is calculated in the spirit of the formula below:

$$RDC\ Inventory\ Turns = \frac{Annual\ Sales\ (units)}{Average\ Inventory\ (units)}$$

Several intermediate variables in the simulation contribute to the calculation of RDC inventory turns. Figure 13 showcases three examples using a single SKU of calculating RDC inventory turns using these intermediate variables in the simulation. These simplified cases are explained below and subsequently generalized using equations to encompass all SKUs.

Case	EOP Inventory(t-1) [A]	Factory Diverts + CxDF(t) Inbound(t) + PDC Inbound(t) [B]	BOP Inventory(t) [C] = [A] + [B]	Actual Sales(t) [D]	Lost Regional Sales(t) [E] = [D] − [C] if [D] > [C[, otherwise 0	EOP Inventory (t) [F] = [C] + [E] − [D]	Average Inventory (t) [G] = ([C] + [F])/2	Regional Daily Fulfilled Quantity (t) [H] = ([D] − [E])/14	RDC Inventory Turns(t) [I] = [H]x365 /[G]
Case 1	50	20	70	40	0	30	50	40/14	20.9
Case 2	20	30	50	40	0	10	30	40/14	34.8
Case 3	10	0	10	40	30	0	5	10/14	52.1

Figure 13: Three Cases Showcasing Calculation of RDCs Inventory Turns Using a Single SKU

In all three cases, the *Average Inventory* over the two-week decision period is determined by calculating the average of the *Beginning of Period (BOP) Inventory*, which is the inventory level at the start of the decision period, and the *End of Period (EOP) Inventory*, which is the inventory level at the end of the decision period. For instance, in the first case, the *BOP Inventory* (70 units) is calculated by taking the *EOP Inventory* from the last period (50 units) and adding any arriving units at the RDCs from the factory, CxDF or PDC (20 units). The *EOP Inventory* (30 units) is determined by taking the *BOP Inventory*, subtracting the *Actual Sales*, and adding back any *Lost Regional Sales.* This convoluted definition of *EOP Inventory* is due to how the simulation is setup. In the simulation, *Lost Regional Sales* are defined as the difference between *Actual sales* and *BOP Inventory* if the former exceeded the latter, and 0 otherwise. The simulation assumes that these *Lost Regional Sales* are fulfilled from the PDC. In the first case, *Lost Regional Sales* are assumed to be 0, so the *EOP Inventory* is simply calculated as 70 units (BOP inventory) – 40 units (actual sales) + 0 (lost regional sales) = 30 units. As such, we end up with an *Average Inventory* of (70 + 30)/2 = 50 in period *t.*

In the third scenario, the *BOP Inventory* was 10 units, and *Actual Sales* were 40 units, indicating that 30 units must have been sourced from the PDC. Therefore, the *EOP Inventory* is calculated as 10 units (BOP) plus 30 units (lost regional sales) minus 40 units (actual sales), resulting in 0 units. The *Average Inventory* in this case is (10+0)/2 = 5.

To calculate the annual sales in the numerator, the *Regional Daily Fulfilled Quantity* is defined as the difference between *Actual Sales* and *Lost Regional Sales*, divided by 14 to convert the metric to a daily basis. This figure is then multiplied by 365 to project the annual volume sold from the RDCs. In the first case, the *Actual Sales* are 40 units and *Lost Regional Sales* are 0 units. Thus, the *Regional Daily Fulfilled Quantity* is calculated as (40-0)/14. This metric is annualized in the numerator when calculating the *RDCs Inventory Turnover.*

The three cases discussed above assume a single SKU. If these three cases were combined such that the simulation had just 3 of these SKUs, the *RDC Inventorty Turns* would be caluclated as:

$$RDC\ Inventory\ Turns_t = \frac{\left(\frac{40}{14} + \frac{40}{14} + \frac{10}{14}\right) x\ 365}{50 + 30 + 5}$$

Below, I present the generalized equations used in the simulation to calculate inventory turns at the RDCs across all SKUs.

$$RDC\ Inventory\ Turns_t = \frac{365 * \sum_{SKU} Regional\ Daily\ Fulfilled\ Quantity_{SKU,t}}{\sum_{SKU} Average\ Inventory\ Quantity_{SKU,t}}$$

Let’s deconstruct the metric in the numerator:

$$Regional\ Daily\ Fulfilled\ Quantity_{SKU,t} = \frac{Actual\ Sales_{SKU,t} - Lost\ Regional\ Sales_{SKU,t}}{14}$$

$Actual\ Sales_{SKU,t}$ is taken from a historical snapshot of data and $Lost\ Regional\ Sales_{SKU,t}$ is calculated as:

$$Lost\ Regional\ Sales_{SKU,t} = Actual\ Sales_{SKU,t} - Inventory\ at\ BOP_{SKU,t}$$
$$if\ Actual\ Sales\ >\ BOP\ Inventory,$$
$$0\ otherwise$$

Let's deconstruct the metric in the denominator:

$$Average\ Inventory\ Quantity_{SKU,t} = \frac{BOP\ Inventory_{SKU,t} + EOP\ Inventory_{SKU,t}}{2}$$

$$\begin{aligned}BOP\ Inventory_{SKU,t}&\\ &= EOP\ Inventory_{SKU,t-1} + Factory\ Diverts_{SKU,t} + CxDF\ Inbound_{SKU,t}\\ &+ PDC\ Inbound_{SKU,t}\end{aligned}$$

$$\begin{aligned}EOP\ Inventory_{SKU,t}&\\ &= BOP\ Inventory_{SKU,t} + Lost\ Regional\ Sales_{SKU,t} - Actual\ Sales_{SKU,t}\end{aligned}$$

- **Percentage of Non-PDC Inflows to RDCs**: This is the percentage of inventory arriving at the RDCs that did not come from the PDC. To relieve capacity and processing constraints for digital demand at the PDC, leadership has set a goal of maximizing the inventory arriving at the RDCs coming from non-PDC nodes. This metric can be measured reliably in the simulation and has become a key metric for identifying improvements to IAS.

- **CxDF Overage**: This is the percentage of *excess* inventory sent to the CxDF that is not utilized by the RDCs and thus sent to the PDC.

- **In-Region Fulfillment Percent**: This metric represents the proportion of demand at the RDCs that was satisfied using inventory already available at the RDCs. The model, however, operates under the assumption of zero lead time for the movement of inventory within and between nodes in the US, without considering any capacity or processing constraints. This assumption allows inventory from the PDC to be transferred to the RDCs too readily, which then contributes to the calculation of in-region fulfillment. Consequently, this metric does not provide a reliable measure for gauging progress towards Iota's strategic goal of maximizing in-region fulfillment. Due to this unreliability and despite its relevance to one of Iota's key objectives, this metric is not included in the reporting of the digital twin's findings.

- **Variable Network Costs**: This metric encompasses the expenses associated with processing, transporting, and splitting up POs. These costs are computed for each scenario to facilitate a comparative analysis of performance against costs incurred. Given that fixed costs do not vary across scenarios, the focus here is solely on variable costs. The components of the network costs include:

- **Processing Cost Per Unit**: This covers the variable costs incurred at various facilities, such as the RDCs, PDC, CxDF, and Off-Site Storage (OST)[9], for handling inventory. Processing costs specific to the port in Asia are omitted from this calculation since they are constant across all scenarios. Additionally, it is assumed that the per-unit processing cost at the CxDF represents an average of the costs at the PDC and OST.

- **Shipping Cost Per Unit**: This reflects the cost associated with transporting each unit across different segments of the supply chain, delineated as first-, middle-, and last-mile costs:
    - **First-mile**: The cost of shipping from the port in Asia to the RDCs, CxDF, or PDC.
    - **Middle-mile**: The cost of transporting goods from the CxDF to the PDC/RDCs/OST; from the OST to the RDCs; and from the PDC to the RDCs. For the purposes of this analysis, the OST is posited to be a national center strategically located the PDC. This central location enables the OST to efficiently serve all RDCs across the US.
    - **Last-mile**: The cost of delivery from the RDCs or PDC directly to the customer.

- **PO Splitting Costs**: Splitting a PO incurs additional expenses that are factored into the simulation. These costs are attributed to each new PO created from the division of an original PO. As a result, every new PO derived from splitting an original one bears a specific cost. For example, if an original PO is divided into y new POs, with each split incurring a cost of $\$x$, the total cost associated with splitting the original PO would amount to $\$xy$.

[9] OST is a buffer storage facility that will be described in detail later in the thesis.

Chapter 4: Scenario Analysis Using Digital Twin

The digital twin of Iota's supply chain, developed as part of this study, is an instrumental simulation tool for enhancing the performance of the IAS/CxDF solution. In this section, I will detail and examine the results from various scenario analyses conducted with the digital twin. These scenarios involved the application of different inventory deployment strategies in the first- and middle-mile segments of the distribution network. The specific parameters varied in these scenarios will be discussed later in this section. The choice of these variations was informed by insights from supply chain leaders at Iota, literature research, and intermediate simulation results. The simulation encompasses a total of 32 million units, comprising thousands of SKUs arriving between May 2022 and February 2023. Running each scenario in the simulation takes approximately 40 minutes. This large run time per scenario constrained the number of scenarios that could be run and evaluated in a timely manner. The scenarios analyzed through the digital twin fall into three main groups:

1. **Demand-Driven Inventory Deployment**: In these scenarios, inventory deployment decisions were based on forecasted demand patterns.
2. **Supply-Driven Inventory Deployment**: These scenarios focused on deploying inventory according to available supply, regardless of forecasted demand.
3. **Demand-Driven Inventory Deployment with Synthetic Data**: Similar to the first group, these scenarios also revolved around demand forecasts for inventory deployment. However, they incorporated synthetic data to evaluate certain policies for which actual data was not available.

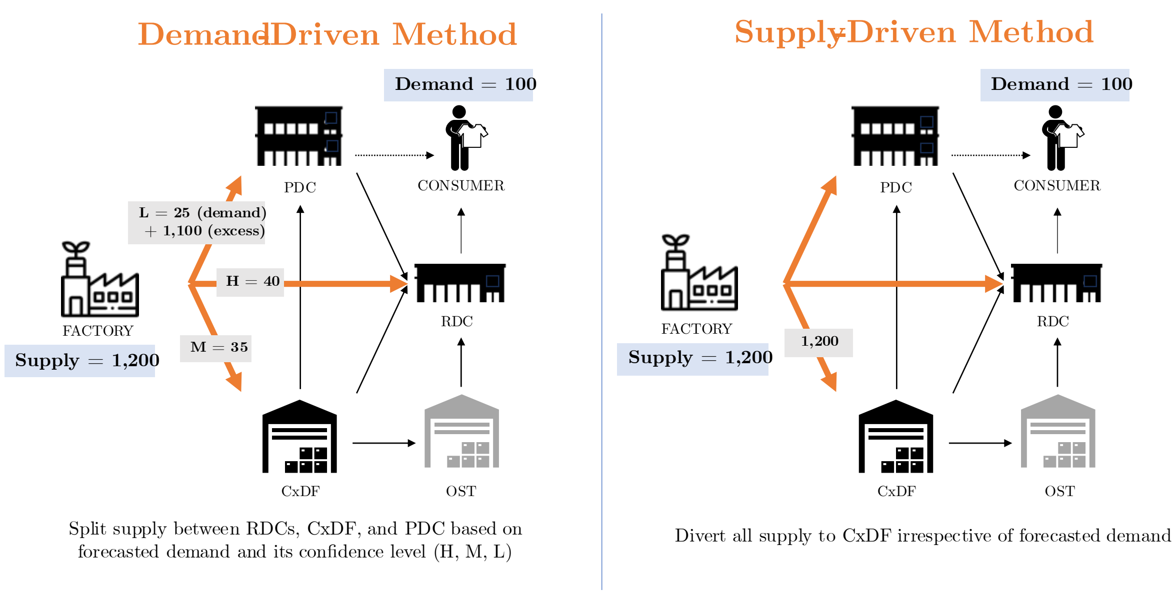

Figure 14: Demand vs Supply Driven Deployment

The terms “Demand-Driven” and “Supply-Driven” require additional clarification to better understand their application within this study. As depicted in Figure 14, the decision to divert inventory from the factory under a demand-driven strategy considers forecasted demand. For instance, with a forecasted demand of 100 units, the IAS strategically allocates a total of 100 units—divided as 40 high, 35 medium, and 25 low confidence levels—to the RDCs, CxDF, and PDC respectively. Any surplus, in this case, 1100 units, is directed to the PDC. The 35 units arriving at the CxDF are allocated to satisfy unfulfilled demand at the RDCs, guided by the updated demand forecast and the unmet requirements at the CxDF. Any surplus inventory is then forwarded to the PDC. The demand-driven scenarios can alternatively be conceptualized as **pull strategy**, where inventory distribution is dictated by forecasted demand, essentially pulling inventory through the supply chain.

Conversely, the supply-driven approach dispatches all 1200 units to the CxDF, disregarding forecasted demand. Here, inventory deployment is solely based on the available supply, with demand considerations playing no role. The substantial volume of inventory reaching the CxDF is distributed among the RDCs, OST, and PDC using various allocation strategies, which are explored within the digital twin and will be discussed later in this document. This supply-driven approach resembles a **push strategy**, wherein all inventory is pushed through to the next stage, irrespective of forecasted demand. This distinction highlights the underlying strategies of responding to actual demand versus prioritizing the movement of available supply.

The following three sections of this chapter delve into each of the three scenario groups in depth. They outline the specific parameters that were adjusted, explain the reasons behind these adjustments, and examine the effects on different metrics across the scenarios. The final section of this chapter synthesizes the findings from all scenarios to offer broader insights. Figure 16 presents an overview of the scenarios that will be discussed in the subsequent sections. It's important to note that while this chapter provides detailed analysis on 11 scenarios, the study encompassed over 50 scenarios in total.

Policy #	Inventory Policy	Rationale	Parameter Modified
1	Send 100% of supply to the PDC.	Modeling the status quo.	None
	Demand-Driven Inventory Deployment		
2	Baseline Flow Strategy (FS) Model	Intended logic of FS in IAS.	Use confidence proportions as determined by the baseline FS model. Send supply to cover two weeks of forecasted demand split across RDCs, CxDF, and PDC based on confidence proportions. Send excess supply to PDC.
3	Send supply to CxDF for High and Medium confidence demand.	Decision postponement for a larger volume of supply compared to (2).	Add the High proportion to Medium proportion in the baseline FS model.
4	Send supply to CxDF for High, Medium, and Low confidence demand.	Decision postponement for a larger volume of supply compared to (3).	Add the High and Low proportions to Medium proportion in the baseline FS model.
	Supply-Driven Inventory Deployment		
5	Send all POs allocated to digital to CxDF.	Decision postponement for all supply allocated for digital orders.	Bypass baseline FS model and send all inventory allocated to digital to the CxDF.
6	(5) + Increase SS target at RDCs from 1 week to 8 weeks.	Decision postponement for all supply allocated for digital orders and address "lumpiness" of Iota's supply by using existing RDCs.	Bypass baseline FS model and send all inventory allocated to digital to the CxDF. Increase SS target at RDCs.
7	(5) + Add a new storage node – OST -- to store some of the excess inventory arriving at the CxDF.	Decision postponement for all supply allocated for digital orders and address "lumpiness" of Iota's supply by building a new storage facility.	Bypass baseline FS model and send all inventory allocated to digital to the CxDF. Add a new storage node – OST – to the network.
8	(7) + Divert 30% of inventory to RDCs for	Decision postponement for all supply allocated for	Bypass baseline FS model and send all inventory

	products with no forecast signal when they arrive at the CxDF.	digital orders. Address "lumpiness" of Iota's supply by building a new storage facility as well as address the lack of forecast coverage.	allocated to digital to the CxDF. Add a new storage node – OST – to the network. Divert 30% of inventory for products with no forecast upon arriving at CxDF to RDCs.
Demand-Driven Inventory Deployment with Synthetic Data			
9	Baseline FS with Synthetic Forecast	Improve the forecast coverage and revert to demand-driven deployment.	Create synthetic forecast data and use it in lieu of original forecast.
10	Baseline FS with Synthetic POs	Address the lumpiness of supply by creating a synthetic smooth supply.	Create synthetic supply data and use it in lieu of original supply data.
11	Baseline FS with Synthetic Forecast and POs	Address the lumpiness of supply as well as lack of forecast coverage with synthetic data.	Create synthetic supply and forecast data and use it in lieu of original data.

Figure 15 Inventory Deployment Policies Simulated Using the Digital Twin

4.1. Scenario Analysis: Demand-Driven Inventory Deployment

The first set of simulations conducted with the digital twin examine the principles of demand-driven inventory management. Designed according to the specifications of the IAS/CxDF solution, the digital twin directs inventory at the Asian port based on forecasted demand (13 weeks out) and confidence levels determined by the Flow Strategy model. The inventory from the port is routed not just to the PDC but also to the RDCs and CxDF, diversifying the distribution channels beyond the original PDC-only approach. The allocation of inventory among these routes is guided by demand confidence proportions, which are specified for each product segment by the Flow Strategy model, a process detailed in Chapter 2.

In these scenarios, the primary parameter varied is the quantity of forecasted demand allocated to the CxDF. In the baseline scenario, inventory associated with medium-confidence demand forecasts is redirected to the CxDF. It is crucial to remember that the CxDF serves as a strategic point for delaying final distribution decisions; upon inventory's arrival at the CxDF, the IAS reevaluates updated demand forecasts to decide whether inventory should be forwarded to the RDCs or the PDC. Given that the decision to divert inventory from the Asian port is made 13 weeks in advance, there exists significant uncertainty within the demand forecast. Therefore, the CxDF's role in postponing decisions becomes pivotal in optimizing the distribution of inventory from Asia, ensuring more accurate and efficient routing.

Utilizing the CxDF as a hub for decision postponement, the analyzed scenarios vary the amount of inventory rerouted to the CxDF, thus broadening the scope of decision postponement to encompass a larger volume and wider range of products. Instead of confining these diversions to inventory linked with medium-confidence demand, the scenarios incrementally incorporate inventory associated with high-confidence (H) and, ultimately, low-confidence (L) demand into the diversion process, enhancing the flexibility and responsiveness of the supply chain to fluctuating demand forecasts.

As illustrated in Figure 16, employing this demand-driven approach, which resembles a pull strategy, resulted in only 5.4% of the inflows bypassing the PDC to reach the RDCs. When the strategy was adjusted to divert a greater volume of inventory to the CxDF, including the high-confidence segment of demand, this percentage further decreased to 3.1%. This decrease indicates that some inventory, initially directed to the RDCs and now rerouted to the CxDF due to its high-confidence demand status, turned out to be unnecessary at the RDCs based on subsequent demand assessments. The figure modestly improved to 4.1% when inventory associated with low-confidence demand was also redirected to the CxDF. Nonetheless, regardless of these adjustments, the performance metrics, ranging from 3.1% to 5.4%, are modest at best given the objective to maximize the non-PDC inflows to RDCs.

Scenario	% of PO Volume Diverted to CxDF	CxDF Overage	% of Non-PDC Inflows to RDCs	RDC Inventory Turns	Variable Network Costs (Scaled)
Baseline Flow Strategy Model	1.6%	58.0%	5.4%	11.6	100.0
H+M to CxDF, RDC SS 1wk	3.7%	31.7%	3.1%	11.8	100.1
H+M+L to CxDF, RDC SS 1wk	5.7%	39.7%	4.1%	11.8	100.2

Figure 16: Performance Metrics for Demand-Driven Inventory Deployment

Remarkably, a mere 1.6% of the inventory leaving the factory was rerouted to the CxDF from the Asian port. This figure slightly increases to 5.7% in the third scenario, where the entire inventory matching the two-week demand forecast is diverted to the CxDF. Put another way, when implementing a demand-driven deployment strategy, over 94% of the inventory ends up being directed to the PDC from the port, which is not much different from the previous design of routing all inventory to the PDC. This pattern suggests that due to certain complexities within the model, most of the inventory is channeled to the PDC.

A closer analysis of the data revealed that the suboptimal performance of the demand-driven deployment – espeically the 94% of inventory being diverted to the PDC from the port in Asia – was mainly influenced by two critical factors:

- **Lumpiness of Supply**: Iota's POs demonstrate a lumpy pattern, characterized by infrequent occurrences but in substantial volumes. This trait reflects Iota's procurement strategy of bulk purchasing for an entire season, a practice rooted in its traditional wholesale business model. Within the simulation, 25% of these large POs, each comprising more than 1000 units, accounted for 74% of the total PO volume. These POs typically covered 8 weeks of demand on average. Additionally, an average SKU receives fewer than 5 POs annually within the network. This pattern of lumpy supply typically directs the majority of inventory to the PDC from the port as the IAS weekly model seeks to distribute inventory based on two weeks of demand. As a result, the excess supply in the POs consistently ends up at the PDC. While Iota was aware of the lumpy nature of its POs, it had not posed a significant concern until the D2C channel began to capture a substantial share of sales. The simulation underscored the pronounced impact of this lumpy supply pattern on efficiently fulfilling D2C orders, revealing the need for adjustments in Iota's supply chain design to accommodate the growing importance of D2C sales.

- **Forecast Coverage**: The simulation also highlighted a significant number of products lacking long-range forecast data necessary for making PO diversion decisions in the digital twin. Products devoid of relevant forecast data were excluded from the flow strategy model, leading to their POs being automatically rerouted to the PDC. The issue of missing forecast data is complex and is partly due to the reliance on a singular source of forecast data in the digital twin (and by extension, in the SupplySoft Solution), which does not encompass all available forecast information available to Iota for diversion decisions. Another complexity is that the comprehensive forecast data intended for input into the SupplySoft Solution may not have been consistently maintained in the past. This lack of historical data maintenance could have contributed to the elevated level of missing forecasts observed in this simulation.

To address these challenges uncovered from the simulation, two strategies were pursued: (i) Transitioning from a demand-driven to a supply-driven deployment, and (ii) Reverting to a demand-driven approach but incorporating synthetic data to simulate both the smoothing of POs and the comprehensive availability of forecast data. This adjustment aimed to explore potential improvements in network performance if the identified obstacles in the initial demand-driven deployment could be overcome. The next two sections describe the scenarios explored and the results of these two appraoches.

4.2. Scenario Analysis: Supply-Driven Inventory Deployment

Given the reliance on accurate, long-term forecasts for making demand-driven inventory deployment decisions at the Asian port, this section explores supply-driven inventory deployment. This method employs a push strategy, allocating inventory according to available supply rather than anticipated demand, a concept elaborated upon earlier in this chapter. It is important to highlight that this push strategy is particularly pertinent to the decisions executed by the weekly IAS model. Once inventory reaches the CxDF, the inventory allocation decisions made by the IAS daily model are made with a one-week lead time, at which point it is assumed that demand forecasts are accessible, thus enabling a demand-driven deployment strategy from the CxDF onwards.

4.2.1. Supply-Driven Inventory Deployment: Pushing All Digital POs to CxDF

In this scenario, all POs intended for digital channels were directly diverted to the CxDF from the Asian port. This effectively directed the complete supply of a product to the CxDF, temporarily disregarding the CxDF's ability to handle this significant volume of inventory. By adopting a supply-driven strategy, this approach allowed for the postponement of decision-making across the entire inventory.

Scenario	% of PO Volume Diverted to CxDF	CxDF Overage	% of Non-PDC Inflows to RDCs	RDC Inventory Turns	Variable Network Costs (Scaled)
Baseline Flow Strategy Model	1.6%	58.0%	5.4%	11.6	100.0
H+M to CxDF, RDC SS 1wk	3.7%	31.7%	3.1%	11.8	100.1
H+M+L to CxDF, RDC SS 1wk	5.7%	39.7%	4.1%	11.8	100.2
100% digital POs to CxDF, RDC SS 1wk	**100.0%**	**89.1%**	**14.5%**	**11.8**	**110.4**

Figure 17: Performance Metrics for Supply-Driven Inventory Deployment I

As seen in Figure 17, the implementation of a supply-driven deployment strategy resulted in an increase in the percentage of non-PDC inflows to the RDCs to 14.5%, a significant rise from the maximum of 5.4% observed in demand-driven deployment scenarios. This represents a considerable increase in the flow of inventory directly to the RDCs without passing through the PDC. However, this scenario also led to a spike in CxDF overage to 89%, indicating that the vast majority of inventory sent to the CxDF exceeded the needs of the RDCs, thereby requiring redirection to the PDC. Moreover, variable network costs escalated from 100 in the baseline scenario to 110.4, indicating a 10% increase in variable costs under this supply-driven model. This increase is attributed mainly to the fact that inventory, which would have been sent directly to the PDC in a demand-driven scenario, now incurs extra processing and transportation costs due to the initial diversion to the CxDF followed by subsequent rerouting to the PDC.

An analysis of the 89% CxDF overage revealed that the primary cause was Iota's lumpy supply. Although the IAS daily model created STOs to cover two weeks of forecasted demand and an additional week of safety stock, it was inadequate in aligning with the large supply, which was typically planned for the entire season. As a result, a significant portion of the inventory still ended up being routed to the PDC after its interim handling at the CxDF. The subsequent two scenarios are designed to tackle the challenge of lumpy supply arriving at the CxDF.

4.2.2. Supply-Driven Inventory Deployment: Increasing SS at RDCs

To mitigate the issue of lump supply arriving at the CxDF, a scenario was implemented where the safety stock target at the RDCs was increased from 1 week to 8 weeks. This adjustment was intended to allow a portion of the surplus inventory at the CxDF to be stored at the RDCs, rather than being redirected to the PDCs. Such a strategy could prove advantageous if the products were later needed at the RDCs, as it would potentially lower transportation costs. However, should the inventory remain unused at the RDCs, it would negatively affect the inventory turnover rates there and eventually lead to additional expenses associated with clearing the surplus inventory from the RDCs.

As shown in Figure 18, this strategy led to an increase in non-PDC flows to RDCs from 14.5% to 16.9%, as depicted in Figure 19, and a reduction in CxDF overage from 89.1% to 83.5%. However, it's crucial to note that this approach significantly compromised inventory efficiency at the RDCs, with median inventory turns plummeting by over 50%, from 11.8 to 4.2. Maintaining inventory turnover rates at the RDCs was an implicit goal of the supply chain network design, aimed at doing "no harm." Additionally, variable network costs saw a 5% increase to 115 compared to the baseline scenario. While the increase in safety stock target at the RDCs did marginally improve non-PDC inflows and CxDF overage, these gains were negated by the adverse effects on inventory turnover and a rise in network costs. Therefore, raising the safety stock target at the RDCs is not deemed a viable option.

Scenario	% of PO Volume Diverted to CxDF	CxDF Overage	% of Non-PDC Inflows to RDCs	RDC Inventory Turns	Variable Network Costs (Scaled)
Baseline Flow Strategy Model	1.6%	58.0%	5.4%	11.6	100.0
H+M to CxDF, RDC SS 1wk	3.7%	31.7%	3.1%	11.8	100.1
H+M+L to CxDF, RDC SS 1wk	5.7%	39.7%	4.1%	11.8	100.2
100% digital POs to CxDF, RDC SS 1wk	100.0%	89.1%	14.5%	11.8	110.4
100% digital POs to CxDF, RDC SS 8wk	**100.0%**	**83.5%**	**16.9%**	**4.2**	**115**

Figure 18: Performance Metrics for Supply-Driven Inventory Deployment II

4.2.3. Supply-Driven Inventory Deployment: Adding OST Storage

To address the issue of lumpy supply arriving at the CxDF while preserving inventory turnover rates at the RDCs, the concept of establishing a new Off-Site Storage (OST) facility was explored. The OST is envisioned as an intermediary storage point within the network, designed to house a portion of the overflow inventory received at the CxDF, thereby reducing the necessity to reroute all surplus stock to the PDC.

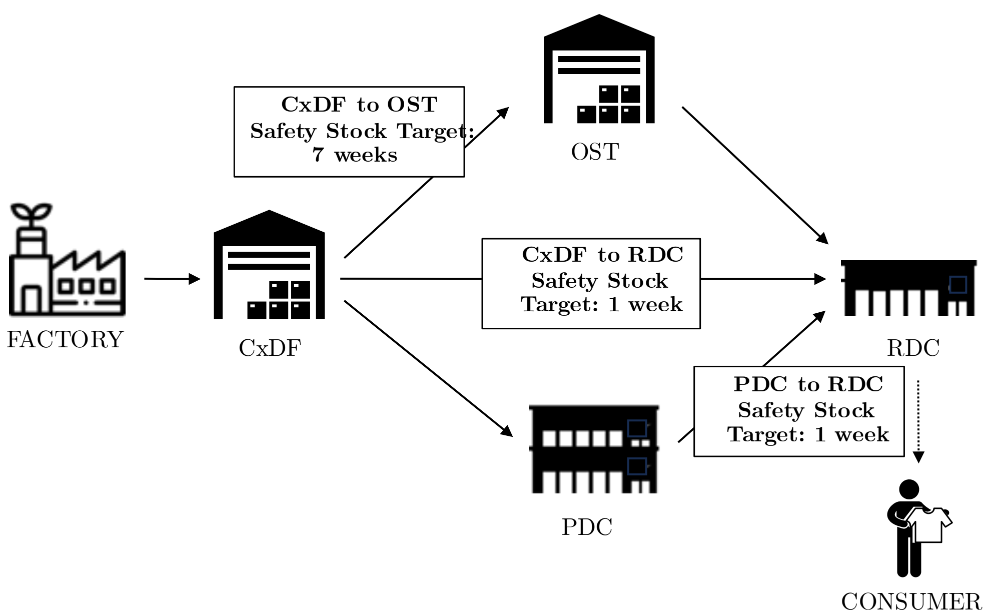

Figure 19: Design of OST Node in the Simulation

In Figure 19, the integration of the OST into the simulation is presented. The OST is designed to receive inventory exclusively from the CxDF based on the safety stock target defined at the

OST. A 7-week[10] safety stock target was chosen for OST considering that it takes on average 7 week of demand to consume a PO based on historical data in recent years. The model calculates OST's safety stock target by extrapolating the two-week forecasted demand at the RDCs.

If there is excess inventory at the CxDF after fulfilling the demand and safety stock target at the RDCs, this surplus is sent to the OST based on the safety stock need at the OST. Any remaining inventory at the CxDF, after meeting the requirements at the RDCs and OST, continues to be dispatched to the PDC. Regarding the matching of demand and supply, the model first utilizes the supply available at the CxDF, followed by the supply at the OST, and finally, the supply at the PDC. The model assumes that a single national OST serves all three RDCs.

Scenario	% of PO Volume Diverted to CxDF	CxDF Overage	% of Non-PDC Inflows to RDCs	RDC Inventory Turns	Variable Network Costs (Scaled)
Baseline Flow Strategy Model	1.6%	58.0%	5.4%	11.6	100.0
H+M to CxDF, RDC SS 1wk	3.7%	31.7%	3.1%	11.8	100.1
H+M+L to CxDF, RDC SS 1wk	5.7%	39.7%	4.1%	11.8	100.2
100% digital POs to CxDF, RDC SS 1wk	100.0%	89.1%	14.5%	11.8	110.4
100% digital POs to CxDF, RDC SS 8wk	100.0%	83.5%	16.9%	4.2	115.0
OST+RDC, 8wk, ignore no fcst	**100.0%**	**77.8%**	**25.7%**	**11.8**	**110.5**

Figure 20: Performance Metrics for Supply-Driven Inventory Deployment III

As illustrated in Figure 20, the addition of the OST with a seven-week safety stock target led to notable improvements. This approach increased the proportion of inventory flowing directly to the RDCs from 14.5% to 25.7%, while also reducing CxDF overage from 89.1% to 77.8%. Crucially, these improvements did not negatively impact the inventory turnover rates at the RDCs, and the costs were on par with the scenario where 100% of the supply was directed to the CxDF without addressing the issue of uneven supply. However, it's important to note that costs were still 10% higher than those observed in the baseline demand-driven scenario.

[10] I explored using a 3-week safety stock target at the OST but settled on a 7-week target given the large size of POs. Additionally, it would enable a fairer comparison with the scenario where 8 weeks of SS was diverted directly to the RDCs from the factory.

4.2.4. Supply-Driven Inventory Deployment: 30% Divert Policy

Despite efforts to improve forecast coverage at the Asian port and manage the lumpiness of supply reaching the CxDF through the addition of the OST, the non-PDC inflows to the RDCs stabilized at 27%. Notably, CxDF overage persisted at 78%, suggesting that a significant portion of the inventory arriving at the CxDF continued to be redirected to the PDCs, even after implementing strategies to route one week of safety stock (SS) to the RDCs and seven weeks of SS to the OST.

Analysis of the data revealed that even with the OST in operation, CxDF overage remained at 78%, largely because more than half of the products arriving at the CxDF lacked updated demand forecast data. This deficiency in forecast coverage was notably amplified by disruptions related to Covid-19 period covered in the simulation. With POs arriving prematurely at the U.S. port due to supply chain adjustments during the Covid era, forecast data was frequently unavailable in the period covered. It was observed that acquiring a forecast signal for products initially without such data upon reaching the CxDF averaged around six weeks. Consequently, the model consistently opted to redirect these POs to the PDC.

To enhance IAS daily model's functionality, improving both the breadth and precision of forecast data is essential. Thus, a specific policy was enacted within the simulation framework. This policy mandates that, in the absence of forecast data for a product arriving at the CxDF, 30% of its associated PO volume should be allocated to the OST.

Scenario	% of PO Volume Diverted to CxDF	CxDF Overage	% of Non-PDC Inflows to RDCs	RDC Inventory Turns	Variable Network Costs (Scaled)
Baseline Flow Strategy Model	1.6%	58.0%	5.4%	11.6	100.0
H+M to CxDF, RDC SS 1wk	3.7%	31.7%	3.1%	11.8	100.1
H+M+L to CxDF, RDC SS 1wk	5.7%	39.7%	4.1%	11.8	100.2
100% digital POs to CxDF, RDC SS 1wk	100.0%	89.1%	14.5%	11.8	110.4
100% digital POs to CxDF, RDC SS 8wk	100.0%	83.5%	16.9%	4.2	115.0
OST+RDC , 8wk, ignore no fcst	100.0%	77.8%	25.7%	11.8	110.5
OST+RDC , 8wk, 30% no fcst	**100.0%**	**65.0%**	**38.3%**	**11.8**	**110.7**

Figure 21: Performance Metrics for Supply-Driven Inventory Deployment IV

Figure 21 showcases the significant enhancements in key performance indicators resulting from this policy's enactment. Mandating that 30% of PO volumes be diverted to the OST for products lacking a demand forecast at the CxDF resulted in an uplift in non-PDC inflows to the RDCs to 38.3% and a reduction in CxDF overage to 65%. The inventory turnover rates at the RDCs were maintained at 11.8, and the costs stabilized at 110.7.

In comparison to the baseline demand-driven scenario, the supply-driven deployment scenarios marked a significant departure, improving the percentage of non-PDC inflows to the RDCs from a high of 5.4% in demand-driven scenarios to 38.3% in the best-performing supply-driven scenario, edging closer to goal of maimizing the target. The saturation at 38.3%, despite numerous network design modifications, implies that PDC will continue to face significant strain as it is utilized to supply inventory to the RDCs. It is also crucial to acknowledge that this boost in network efficiency was accompanied by a 10% cost increase, primarily due to heightened middle-mile expenses from increased transfers between the CxDF and OST, among other factors.

4.3. Scenario Analysis: Demand-Driven Deployment With Synthetic Data

In this section, we return to exploring demand-driven inventory deployment, supplementing our analysis with synthetic data to tackle the challenges of lumpy supply and inadequate forecast coverage. As a reminder, our transition to a supply-driven approach was a deliberate strategy to navigate the hurdles posed by sparse forecast data while making divert decision at the port in Asia and by the lumpy nature of the POs. This shift towards a supply-driven methodology significantly improved network performance, as evidenced by increases in non-PDC inflows to RDCs and reductions in CxDF overage, but it also resulted in a noticeable rise in variable network costs. Additionally, concerns regarding the CxDF's capacity to handle all POs allocated to digital orders from Asian factories were momentarily overlooked. Consequently, we pivot back to a demand-driven deployment strategy, incorporating synthetic data to effectively tackle the identified issues and assess the network's performance.

4.3.1. Demand-Driven Deployment With Synthetic Data: Synthetic Forecast

The first issue addressed was the absence of comprehensive forecast data. In the original long-range forecast dataset, merely 24% of products had a forecast signal. This scarcity of forecast data stems from the reliance on a single internal data source in the digital twin, despite the existence of multiple forecast data sources internally. These varied sources would need to be integrated to enhance product coverage. However, due to time constraints and the discovery of this issue relatively late in the process, integrating these disparate datasets — managed by different teams — into a unified forecast source was unfeasible. Consequently, I opted to create a synthetic dataset to improve the forecast coverage.

Synthetic forecasts were generated for products lacking projections for weeks 13 and 14 from the decision date for diversion. This process involved applying an imputation factor, derived from a scientifically established range, to actual sales data for these weeks. Sales figures were adjusted by a factor within the range [0.50, 1.87], corresponding to the 25th and 75th percentiles, respectively, of the forecast-to-sales ratio for products with both available forecasts and actual sales data for the simulation period. The selection of the imputation factor for a given product was random but remained unchanged across different decision dates and regions.

For example, illustrated in Figure 22, by applying an imputation factor of 1.2 to the sales figures for weeks 13 and 14, we generated a forecast of 96 and 144 respectively. This imputation factor of 1.2 remained unchanged across all the decision dates in the simulation for this product.

Decision Date	Product Code	Demand Week	Sales	Factor	Synthetic Forecast
2022/02/05	X08720-A10	13	80	1.2	80x1.2 = 96
2022/02/05	X08720-A10	14	120	1.2	120x1.2 = 144

Figure 22: Synthetic Forecast Creation Example

The creation of synthetic forecast data markedly expanded the forecast coverage for products, enhancing the dataset's comprehensiveness. This approach effectively increased the proportion of products with a non-zero forecast signal from 24% to 47%. It is important to remember that without a positive forecast signal, the IAS weekly model defaults to routing products from the factory directly to the PDC. As a result, in the simulation, 53% of the products had their entire PO volume dispatched to the PDC from the Asian port due to zero demand forecast. Additionally, this percentage reflects the count of products that lacked a non-zero forecast signal in each decision period divided by the total number of products in each decision period. It could be that a product without a non-zero forecast signal at later decision points might possess a positive forecast signal either in the original long-range forecast data set or as a result of the sales-based imputation method.

Although the expansion in product forecast coverage is notable, it is important to recognize the limitations of this approach. Specifically, the fixed assumption of a fixed 13-week lead time can occasionally fail to generate a synthetic forecast signal. For instance, as depicted in Figure 23, a forecast signal is present in the first scenario where actual sales data are available for weeks 13 and 14. However, in the second scenario sales for the product begin only several weeks beyond the assumed 13-week lead time, leading to a zero forecast. Thus, 53% of the products had no sales in both weeks 13 and 14, which resulted in a zero forecast.

Adopting a variable lead time could serve as a solution, yet its effectiveness remains uncertain, as demonstrated in the third scenario. Suppose a variable lead time model determined a 15-week lead time. Despite this adjustment, as shown within the red box, the forecast would still be zero because sales commenced 16 weeks after the decision date, as indicated in the green box. Consequently, the adoption of a variable lead time model does not necessarily ensure improved forecast coverage. Given these considerations, I chose to maintain a consistent 13-week lead time assumption to sidestep additional complexities.

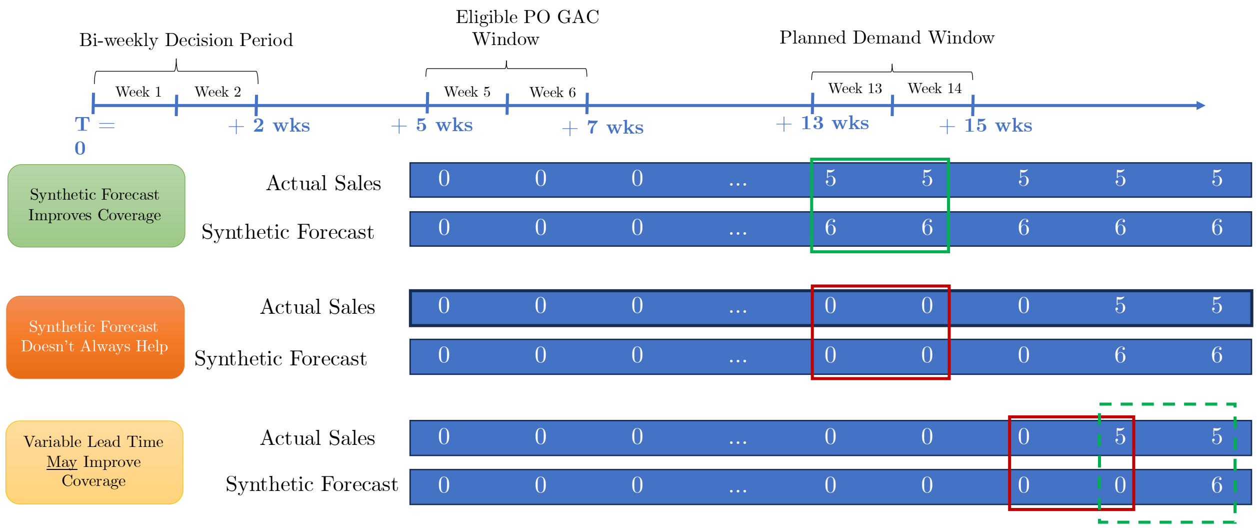

Figure 23: Example Showing Limitations of Synthetic Forecast Creation

As illustrated in Figure 24, the introduction of synthetic forecast data increased the percentage of non-PDC inflows to the RDCs from 5.4% to 8.0%. While this represents an improvement, it remains modest, and the proportion of PO volume redirected to the CxDF is still minimal at 2.2% under a demand-driven deployment strategy enhanced with synthetic forecasts. It is crucial to highlight that forecast coverage only reached 47% due to the reliance on sales-based imputation. With a forecast coverage of 100%, a greater number of products could have been diverted to the CxDF directly from the factory.

Scenario	% of PO Volume Diverted to CxDF	CxDF Overage	% of Non-PDC Inflows to RDCs	RDC Inventory Turns	Variable Network Costs (Scaled)
Baseline Flow Strategy Model	1.6%	58.0%	5.4%	11.6	100.0
Baseline Syn Fcst	**2.2%**	**61.3%**	**8.0%**	**11.4**	**99.7**

Figure 24: Performance Metrics: Demand-Driven Deployment with Synthetic Data I

4.3.2. Demand-Driven Deployment With Synthetic Data: Synthetic PO Smoothing

Although incorporating synthetic forecast data mitigated one of the challenges, addressing another significant issue remains crucial: lumpy supply patterns characteristic of Iota's purchasing behavior. To tackle this, I implemented a strategy that involved smoothing out the placement of POs by altering purchasing behaviors.

This modification applied to all POs containing more than 240 units, which were classified as "large" POs. Each large PO was uniformly distributed over the remaining "eligible period" of the simulation. The remaining eligible period was determined considering the actual arrival date of POs and the period covered in the simulation with the goal of ensuring that the total PO quantity remained unchanged across scenarios with original and synthetic POs. The details of the PO smoothing algorithm are explained in Appendix A.

Figure 25 offers a straightforward example of the PO smoothing process. In this example, an initial PO consisting of 486 units of a SKU is divided into seven smaller POs: six of them with 70 units each, and the final one with 66 units, thereby preserving the total quantity. Each new PO is assigned a unique label and is scheduled at two-week intervals, beginning from the decision date of the original PO. It is important to note that the arrival date for the first new PO has been adjusted to the actual arrival date (6/25), diverging from the estimated arrival date (6/22), to account for IAS's daily decision-making, which is based on the actual arrival date of POs. The determination of the exact number of splits takes into account other factors, such as the time span between the decision date for the given PO and the simulation's final decision date, as well as the arrival date for the last split PO. Comprehensive explanations of these considerations are provided in Appendix A. The fundamental principle maintained is that the scope of inventory should remain constant at 486 units, both before and after the application of PO smoothing.

Decision Date	Product Code	PO #	PO Qty	Estimated Arrival Date
3/19/22	X08720-A10	4508506611	486	6/20/22

Decision Date	Product Code	PO #	PO Qty	Arrival Date
3/19/22	X08720-A10	4508506611_1	70	6/25/22
4/2/22	X08720-A10	4508506611_2	70	7/9/22
4/16/22	X08720-A10	4508506611_3	70	7/23/22
4/30/22	X08720-A10	4508506611_4	70	8/6/22
5/14/22	X08720-A10	4508506611_5	70	8/20/22
5/28/22	X08720-A10	4508506611_6	70	9/3/22
6/11/22	X08720-A10	4508506611_7	66	9/17/22

Figure 25: Example of Smoothing of a Lumpy Purchase Order (PO)

Note that the smoothing strategy adopted here was conservative in nature. An approach to smoothing without constraints detailed in Appendix A could have led some parts of the POs to extend beyond the simulation's end period. Such an extension would alter the total volume of units considered in the simulation, affecting cost calculations and hindering a fair comparison between scenarios. Therefore, the chosen method is intentionally restrained. To this end, it is important to recall that the smoothing of POs results in the creation of additional PO lines, necessitating the application of appropriate PO splitting costs. These costs reflect the increased expenses associated with dividing a PO into multiple new ones.

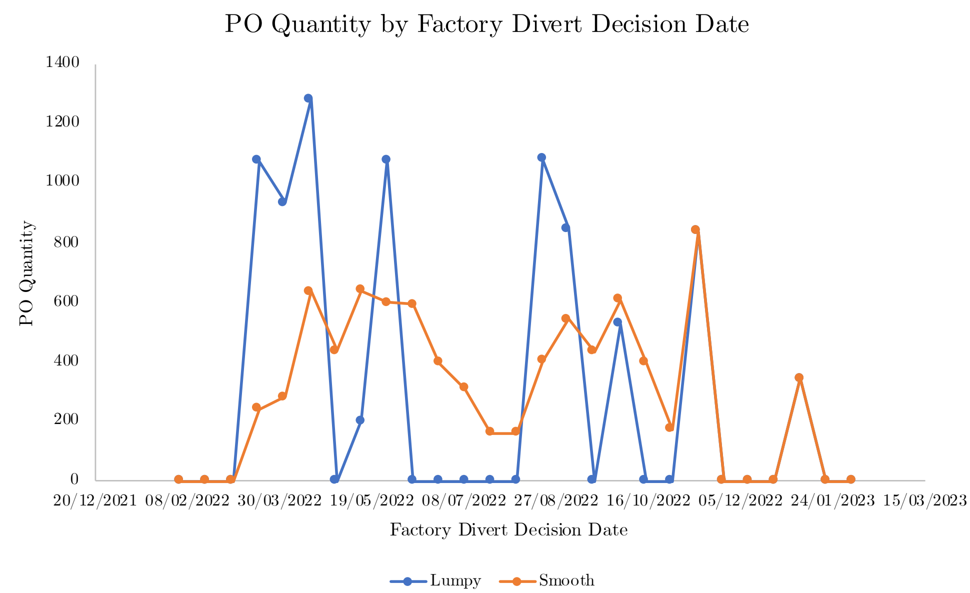

Figure 26: Graphical Representation of Smoothing of Multiple Lumpy Purchase Orders

Figure 26 provides a visual illustration of the impact of PO smoothing on supply. It displays the POs for a product subject to diversion decisions, organized by decision date. The implementation of synthetic smoothing has resulted in more frequent and smaller POs, as evidenced by the comparison between the blue (representing the original lumpy supply) and orange (indicating the smoothed supply) lines.

Scenario	**% of PO Volume Diverted to CxDF**	**CxDF Overage**	**% of Non-PDC Inflows to RDCs**	**RDC Inventory Turns**	**Variable Network Costs (Scaled)**
Baseline Flow Strategy Model	1.6%	58.0%	5.4%	11.6	100.0
Baseline Syn Fcst	2.2%	61.3%	8.0%	11.4	99.7
Baseline Syn POs	**2.4%**	**53.7%**	**9.4%**	**11.3**	**99.9**

Figure 27: Performance Metrics: Demand-Driven Inventory Deployment with Synthetic Data II

When evaluating the effect of smoothing POs on the percentage of non-PDC inflows to RDCs, there is a notable increase from the baseline figure of 5.4% to 9.4% as seen in Figure 27. This enhancement surpasses the improvement achieved through the implementation of synthetic forecasts. Costs remain in line with those of the baseline scenario. However, the proportion of PO volume diverted to the CxDF continues to be modest, standing at 2.4%.

4.3.3. Demand-Driven Deployment With Synthetic Data: Synthetic Forecast and PO Smoothing

The last scenario implemented the baseline demand-driven deployment strategy, enhanced by synthetic forecast data to broaden forecast coverage and synthetic POs to even out lumpy supply. A notable advantage of supply smoothing was the further extension of forecast coverage. Since products might display forecast signals in the original dataset after a lag, dividing large POs into several smaller orders allowed more products to eventually exhibit forecast data. This smoothing increased non-zero forecast coverage from 47% to 66%.

Scenario	% of PO Volume Diverted to CxDF	CxDF Overage	% of Non-PDC Inflows to RDCs	RDC Inventory Turns	Variable Network Costs (Scaled)
Baseline Flow Strategy Model	1.6%	58.0%	5.4%	11.6	100.0
Baseline Syn POs	2.4%	53.7%	9.4%	11.3	99.9
Baseline Syn Fcst	2.2%	61.3%	8.0%	11.4	99.7
Baseline Syn POs and Fcst	**3.6%**	**57.6%**	**14.8%**	**11.2**	**99.3**

Figure 28: Performance Metrics: Demand-Driven Deployment with Synthetic Data III

As illustrated in Figure 28, there was a nearly threefold increase in the percentage of non-PDC inflows to the RDCs when compared to the baseline scenario, rising from 5.4% to 14.8%. CxDF overage stayed relatively stable at around 58%. Despite enhancements in forecast coverage and the smoothing of supply, the percentage of PO volume rerouted to the CxDF remained modest at 3.6%. However, the percentage of PO volume redirected to the CxDF from the Asian port saw a significant improvement relative to the baseline scenario, increasing by a factor of 2.25. The persistence of no forecast data for 34% of the products, along with the conservative approach to PO smoothing, contributes partly to this outcome.

Given that the synthetic demand forecast employs random sampling of imputation factors from a scientifically established range, assessing the robustness of the results for this crucial scenario was essential. To accomplish this, I conducted a sensitivity analysis by executing 10 iterations with varied samplings of imputation factors within this key scenario. The outcomes of this analysis indicate that the findings—pertaining to network performance and costs—are stable, showing minimal variation across all 10 iterations. The details are presented in Appendix B.

4.4. Scenario Analysis: Bringing It All Together

The preceding three sections have elaborated on three distinct sets of scenarios: (i) demand-driven inventory deployment, (ii) supply-driven inventory deployment, and (iii) demand-driven inventory deployment with synthetic data. Utilizing the digital twin, these scenarios were analyzed to advance towards achieving the goals set for improving Iota's supply chain in support of D2C expansion, specifically to maximize the percentage of non-PDC inflows to RDCs. Although this metric peaked at a modest 38% and 15% in the best performing supply-driven and demand-driven scenarios respectively, several important insights have been gleaned for reconfiguring Iota's supply chain to better accommodate the burgeoning demands of D2C growth. The analysis of network performance relative to costs, as shown in Figure 29, illustrates the trade-offs between enhancing network performance and the associated costs across all examined scenarios. The network costs presented in the figure are relative to the baseline scenario labeled *Baseline Flow Strategy Model*.

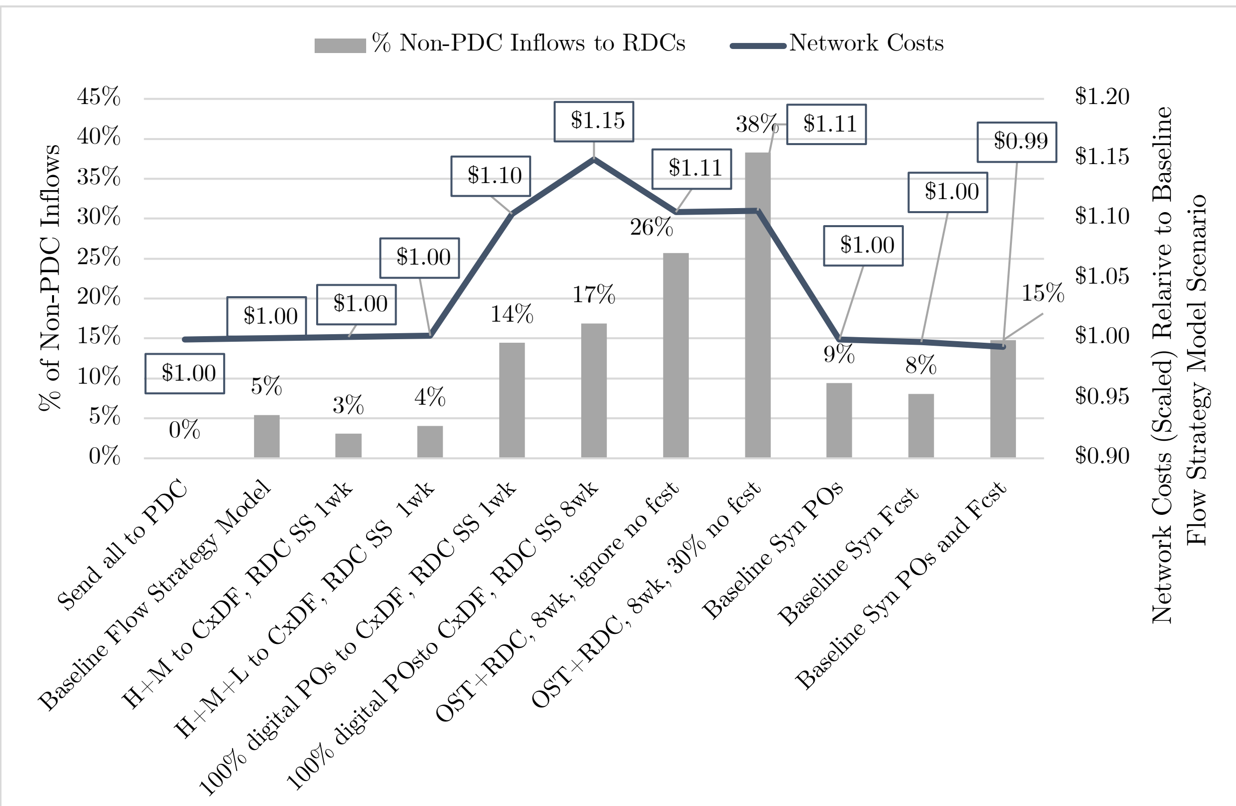

Figure 29: Supply Chain Network Performance vs Network Costs (Scaled) Across Scenarios Simulated Using Digital Twin

Comparing the initial set of demand-driven scenarios (without synthetic data) to the supply-driven scenarios reveals a notable network performance boost, increasing the % of non-PDC inflows from a maximum of 5% in demand-driven scenarios to 38% in supply-driven scenarios. However, the supply-driven approach, despite achieving the highest network performance at 38%, also incurred a 11% increase in costs. This cost surge stemmed primarily from higher processing and transportation expenses in middle-mile operations, with all inventory routed to the CxDF from the Asian port, regardless of forecasted demand. This poses a significant dilemma for Iota, balancing enhanced performance against increased costs. Moreover, diverting all inventory to the CxDF from factories is theoretical and improbable due to the considerable risk of overburdening the facility.

When comparing supply-driven scenarios, both with and without OST, it is evident that merely deploying a CxDF cannot singularly address the existing challenges. Although raising the safety stock target at RDCs boosted performance slightly from 14% to 17%, it detrimentally impacted both inventory turns at the RDCs and slighly increased overall costs. The introduction of OST significantly enhanced network performance to 26% with a less substantial cost ($1.11) increase than when safety stock targets at RDCs were raised ($1.15). Ultimately, without either adapting bulk buying behavior and/or integrating a OST to handle surplus inventory at the CxDF, the sought-after gains in supply chain efficiency may remain elusive.

Further analysis between demand-driven scenarios, with and without synthetic data, indicates that improved forecast accuracy and smoother POs can substantially boost network performance—from 5% in the *Baseline Flow Strategy Model* scenario to 15% in *Baseline Syn POs and Fcst* scenario—tripling performance and decreasing costs by 1%. Even if 1% may seem modest, this cost reduction could translate into significant savings given Iota's extensive operational scale. Notably, while the 15% network performance is less than half of the peak 38% observed in the top supply-driven scenario, it contrasts with a cost reduction of 1% rather than the 11% cost increase in the supply-driven model. A more aggresive supply smoothing and forecast coverage improvement could lead to additional performance gains using demand-driven deployment, without the steep cost escalations observed in the supply-driven scenarios.

Likewise, smoothing the supply chain in conjunction with enhancing forecast data availability has demonstrated additional benefits, as shown by the demand-driven scenarios using synthetic data. Individually, improved forecast coverage boosts performance from 5% to 8%, and smoother supply results in an increase from 5% to 9%, marking isolated gains of 3% and 4% points, respectively. However, when combined, these adjustments lead to a cumulative increase of 10 percentage points, elevating performance from 5% to 15%.

These outcomes from the scenario analysis underscore the need for a singular, reliable forecast data source. The digital twin's efficacy hinges on such data for precise decision-making. In its absence, inventory without a forecast signal will consistently revert to the PDC, countering the IAS engine's targets. Furthremore, it highlights the need for Iota to reconsider and possibly adjust its bulk purchasing practices and/or establish an off-site storage in conjucation with the cross-dock facility to foster more agile supply chain operations.

Chapter 5: Conclusion and Recommendations

This project undertook an empirical assessment to support the transition of Iota's supply chain from a wholesale-centric to a D2C model, a strategic pivot aimed at meeting contemporary consumer expectations in service, cost, and sustainability. The potential enhancements considered by Iota involved substantial investments in both physical and digital infrastructures, notably the establishment of a cross-dock facility and the development of a decision engine – IAS – to determine inventory routing, which signify the commitment to a more decentralized and responsive supply chain network.

The study's fundamental goal was to understand the potential impact of the IAS/CxDF solution on achieving Iota's strategic objectives for its supply chain network, notably the aim for maximizing in-region fulfillment and non-PDC inflows to the RDCs. Since the IAS/CxDF solution was not operational yet during the research period, the project conducted a retrospective analysis using a digital twin that simulated the possible outcomes of these significant infrastructure investments. Various demand- and supply-driven scenarios were analyzed using the digital twin to understand the impact of potential levers on the network performance.

The digital twin served as a robust platform for scenario analysis, revealing key insights into the future of Iota's supply chain under various network configurations. While the scenario analysis yielded a modest 38% non-PDC inflows to RDCs in the best push-driven scenario and 15% in the best pull-driven scenario, the exploration yielded a deeper understanding of the potential levers within Iota's supply chain.

Based on the findings from the scenario analysis, the following recommendations are proposed for Iota:

1. **Integrate Comprehensive Forecast Data**: Iota should endeavor to consolidate and utilize comprehensive long-range and short-range forecast data, enhancing the IAS engine's ability to make informed decisions across the supply chain. In the absence of a comprehensive forecast data, the IAS will simply route the inventory to the PDC, defying the objective of minimizing burden on the PDC.

2. **Smoothing Supply Patterns**: Iota might consider modifying its bulk purchasing behavior to align with a more continuous demand-driven supply chain model. This could involve adopting smoother, more frequent POs, a change that the simulations have shown can improve network performance and cost efficiency. Smoothing can also be achieved by establishing a distribution center in Asia that serves as an intermediary buffer,

moderating the lumpy production volumes and enabling a more streamlined transportation flow to the US.

3. **Invest in CxDF and OST**: The simulations underscored that the CxDF alone will not suffice to meet the challenges of Iota's D2C growth. In the absence of modifying purchasing behavior, an additional storage facility like the OST is instrumental in managing excess inventory.

4. **Review Target Feasibility**: The scenarios suggested that significantly reducing the PDC burden might be more challenging than anticipated. As observed in the scenario analysis, the % of non-PDC inflows to the RDCs reached a maximum of 38%. In light of this, Iota may need to reconsider and potentially recalibrate its specific targets, taking into account the inherent challenges, and accordingly adjust its strategies.

5. **Cost-Benefit Analysis of Supply-Driven Deployment**: While supply-driven deployment has shown promise in improving network performance, it also incurs higher costs. A more detailed cost-benefit analysis than what was employed in this study should be conducted to weigh the trade-offs of this approach versus demand-driven strategies.

6. **Integration with E2EDT and Continued Scenario Analysis**: IAS emulation should be enhanced by relaxing simplifying assumptions concerning lead time, capacity, and safety stock. Following these improvements, the enhanced IAS should be integrated into the E2EDT to conduct further simulations, thereby providing a deeper understanding of network performance given the evolving nature of the IAS/CxDF solution.

Through these recommendations, Iota can continue to refine its supply chain operations to ensure agility and responsiveness in the dynamic global retail landscape, thereby maintaining its competitive advantage during the era of D2C growth. While the challenges highlighted—such as the lack of adequate forecast coverage and bulk purchasing behavior—are discussed within the context of Iota, they arguably extend to many retailers undergoing similar transformations. The insights uncovered like the need to modify bulk purchasing behavior or to explore alternative methods of supply smoothing through the establishment of new storage facilities are lessons that have broader applicability. Additionally, the use of digital twins and synthetic datasets can offer critical insights in a timely and cost-effective manner, enhancing decision-making across the supply chains of many companies similar to Iota in the industry.

Appendix A

In the effort to address the lumpiness of POs within the digital twin simulation, a methodical approach was adopted to generate smoother, more frequent POs by distributing the order quantities across the "eligible period" remaining in the simulation. This process, referred to as PO smoothing, was executed using synthetic data.

The PO smoothing process involved splitting large, infrequent POs into smaller, more regular orders. This was done by staggering the original PO quantity over the "eligible period" in the simulation. To ensure that the total supply quantity remained unchanged—thereby permitting an equitable comparison across different scenarios—the split was constrained. This meant that as the simulation progressed and neared its end period, the algorithm applied a less aggressive approach to PO smoothing.

Algorithm Overview:

1. **Identification**: The algorithm begins by identifying POs at the style-color-size level with quantities of 240 units or more.
2. **Calculation of Last Eligible Divert Date**: For each identified PO, the algorithm calculates the last eligible date for diverting the PO, defined as the minimum of the last decision date in the simulation and the PO's actual arrival date.
3. **Divert Decision Timing**: A divert decision on a PO must be taken before its estimated arrival date or before the simulation period ends.
4. **Estimation of Splits**: The algorithm estimates the number of splits by dividing the difference between the last divert date and the decision date by 14 (the two-week period).
5. **Division of POs**: The original PO is divided into the number of splits calculated above. For each newly created PO resulting from the split, the decision date and arrival date are staggered by increments of 14 days.
6. **Arrival Date of the Last Split**: The arrival date for the last split of the PO is calculated.
7. **Adjustment for Splits Estimated**: If the arrival date of the last split extends beyond the simulation's end date (February 25, 2023), while the original arrival date was within the simulation period, the number of splits is adjusted so that all split POs arrive within the simulation timeframe.
8. **Final Split**: The PO is then split according to the adjusted number of splits.

This smoothing technique aimed to replicate a more smooth flow of POs that would align better with a demand-driven supply chain model, which is particularly beneficial for meeting the variable demands of D2C channels.

Appendix B

The synthetic forecast was developed by applying an imputation factor to each product, randomly selected from a scientifically determined range. The range captured the 25th and 75th percentiles of the forecast-to-sales ratio for products with both sets of data available. Considering the extensive runtime of the model, I executed 9 additional iterations of the baseline scenario, incorporating synthetic forecasts and POs, with each iteration employing a different set of imputation factors. As illustrated in Figure 30, the performance metrics across these iterations remained consistently stable, demonstrating that the model's outcomes are not significantly influenced by the variation in randomly generated imputation factors.

Scenario	% of PO Volume Diverted to CxDF	CxDF Overage	% of Non-PDC Inflows to RDCs	RDC Inventory Turns	Variable Network Costs (Scaled)
Baseline Syn POs and Fcst	3.56%	57.62%	14.80%	11.20	0.99
Baseline Syn POs and Fcst Iter 1	3.57%	58.73%	14.80%	11.20	0.99
Baseline Syn POs and Fcst Iter 2	3.58%	58.75%	14.84%	11.20	0.99
Baseline Syn POs and Fcst Iter 3	3.54%	59.17%	14.87%	11.20	0.99
Baseline Syn POs and Fcst Iter 4	3.60%	58.73%	14.84%	11.20	0.99
Baseline Syn POs and Fcst Iter 5	3.55%	58.88%	14.88%	11.20	0.99
Baseline Syn POs and Fcst Iter 6	3.59%	58.16%	14.88%	11.20	0.99
Baseline Syn POs and Fcst Iter 7	3.57%	58.43%	14.86%	11.20	0.99
Baseline Syn POs and Fcst Iter 8	3.58%	58.45%	14.89%	11.20	0.99
Baseline Syn POs and Fcst Iter 9	3.54%	58.49%	14.82%	11.20	0.99

Figure 30: Performance Metrics for Iterations of the Baseline Scenario with Synthetic Forecast and POs

Appendix C

This section describes the underlying algorithm within the flow strategy model in detail. The flow strategy model is a two-part framework for allocating products to segments and defining the percentage split across high, medium, and low confidence for each segment.

The first part of the model, the segmentation model, divides products into groups based on forecast volume and forecast accuracy. The second part of the model, the demand proportion/confidence model, defines the percentage split across high, medium, and low confidence for each segment.

1. Segmentation Model

1.1. Metrics Calculation

The segmentation model uses forecast volume and forecast accuracy and specifies thresholds on percentiles for these two metrics to segment the products.

Forecast Volume: This is the average weekly forecasted volume in weeks 15 to 20 from a decision date. This is a forward looking metric as it looks at the forecasted volume for a SKU from the decision date.

$$Average\ Forecasted\ Volume_{SKU} = \frac{\sum_{t=15}^{20} Forecasted\ Volume_{SKU,t}}{5}$$

Forecast Accuracy: This is a retrospective measure that evaluates the accuracy of sales forecasts by comparing the forecasted volume to the actual sales. It utilizes the Mean Absolute Percentage Error (MAPE) to assess accuracy. Specifically, the comparison is made for weeks 13 to 16, starting from the forecast creation date (not the same as decision date). To ensure sufficient forecast data, the earliest eligible forecast creation date must be at least 16 weeks prior to the decision date. The calculation of forecast accuracy incorporates the weighted average of actual sales. The equation below demonstrates the computation of forecast accuracy using a single forecast creation date (*FD*).

$$Forecast\ Accuracy_{SKU,FD} = \frac{\sum_{t=13}^{16} \left| \frac{Actual\ Sales_{SKU,FD} - Forecasted\ Sales_{SKU,FD}}{Actual\ Sales_{SKU,FD}} \right| * Actual\ Sales_{SKU,FD}}{\sum_{t=13}^{16} Actual\ Sales_{SKU,FD}}$$

Recognizing that relying solely on a single forecast creation date introduces significant variability, the model mitigates this issue by averaging the accuracy across the 13 most recent forecast creation dates. The following formula illustrates the calculation of the final forecast accuracy metric at the SKU level.

$$Forecast\ Accuracy_{SKU} = \frac{\sum_{FD=1}^{13} \sum_{t=13}^{16} \left| \frac{Actual\ Sales_{SKU,FD} - Forecasted\ Sales_{SKU,FD}}{Actual\ Sales_{SKU,FD}} \right| * Actual\ Sales_{SKU,FD}}{\sum_{FD=1}^{13} \sum_{t=13}^{16} Actual\ Sales_{SKU,FD}}$$

1.2 Segment Thresholds

For forecasting volume, products are categorized based on their forecasted weekly volume. This categorization is determined by percentiles. Products with forecasted weekly volume falling in the bottom 1 percentile are classified as Low (L). Those with forecasted volume between the 2nd percentile and the 80th percentile are categorized as Medium (M). Finally, products with forecasted volume in the top 20 percentile are classified as High (H).

When it comes to forecast accuracy, a separate categorization is applied based on the forecast error. This categorization is also determined by percentiles. Products with forecast errors in the top 20 percentile are classified as Low (L), indicating a higher level of forecast inaccuracy. Products with forecast errors between the 20th percentile and the 80th percentile are categorized as Medium (M), representing a moderate level of forecast accuracy. On the other hand, products with forecast errors in the bottom 20 percentile are classified as High (H), indicating a relatively higher level of forecast accuracy.

Finally, seven segments are defined by according to the matrix in Figure 31 using High (H), Medium (M), and Low (L) categorizations defined above.

Flow Segment Code	Forecast Volume Segment	Forecast Accuracy Segment
A	H	H
B	H	M
C	H	L
D	M	H
E	M	M
F	M	L
T (Tail)	L	H, M and L

Figure 31: Thresholds for Flow Segment Codes by Forecast Volume and Accuracy

2. Demand Confidence Model

The concept behind the demand confidence model is to determine which portion of the forecasted demand we have a high level of confidence in and can prioritize for the fastest routing to the final node, giving it the highest priority. Conversely, it helps us identify the portion of the forecasted demand where our confidence is not as high, allowing us to postpone decision making on the portions of demand that have medium and low confidence levels. The outcome of the demand confidence model provides the percentages of high, medium, and low confidence demand for each flow segment. To achieve this, the MVP demand proportion model analyzes historical data from the past three non-holiday weeks.

The high, medium, and low demand proportions are calculated for each of the seven segments using the algorithm below. Note that since the proportions add up to 1, we just need to calculate the proportions for high- and medium-confidence demand.

1. Gather historical sales data for each product over the past 39 weeks, excluding holiday sales.
2. Calculate the sales threshold for each product corresponding to the 20th percentile, denoted as **H**. This value represents the minimum sales level that was met or exceeded in 80% of cases.
3. Determine the sales threshold for each product corresponding to the 50th percentile, denoted as **M**. This value indicates the minimum sales level that was met or exceeded in 50% of cases.
4. Retrieve the average forecasted volume for products based as calculated in the segmentation model, referred to as **F**.
5. Compute the percentage of High Demand Confidence (R1) for a product as:
$$High\ Demand\ Confidence\ (R1) = \frac{H}{F}\ , if\ H > F, otherwise\ 20\%$$
6. Calculate the percentage of Medium Demand confidence (R2) for a product as:
$$Medium\ Demand\ Confidence\ (R2) = \frac{(M - H)}{F}\ , if\ (H - M) > F, otherwise\ 50\%$$
7. Calculate the percentage of Low Demand confidence (R3) for a product as:
$$Low\ Demand\ Confidence\ (R3) = 100\% - R1 - R2$$
8. After estimating R1, R2, and R3 for each product, calculate the average of these ratios to obtain segment-level demand confidence percentages. Each segment will now have corresponding percentages for High, Medium, and Low demand confidence.
9. Note that all products within a particular flow segment will share the same demand confidence percentages.