Ca' Foscari
University
of Venice

Master's Degree

in International Management

Double Degree Program in International Business in cooperation with SKEMA Business School

Final Thesis

FROM REACTIVE TO PREDICTIVE: AN ECONOMIC ANALYSIS OF ANTICIPATORY LOGISTICS IN THE LUXURY INDUSTRY

Supervisor
Ch. Prof. Raffaele Pesenti

Graduand
Riccardo Brasi
Matriculation Number 888544

Academic Year
2025 / 2026

DECLARATION CONCERNING THE USE OF ARTIFICIAL INTELLIGENCE (AI) TOOLS

I, the undersigned: Riccardo Brasi

Student ID Number: 888544

Enrolled at Ca Foscari University of Venice in the Master's Degree in: International Management

Author of the final thesis: From Reactive To Predictive: An Economic Analysis Of Anticipatory Logistics In The Luxury Industry

HEREBY DECLARES that generative artificial intelligence technologies and tools have been used as a limited support during the research and writing of the final thesis. The tool was employed to assist in software implementation and coding, supporting the design, development and structuring of software components, including specific portions of Python (under my supervision), as well as for debugging the project prototype. In addition, AI has been used for brainstorming, improving the grammar, readability and stylistic clarity of the text originally drafted by me. AI was also used to generate a schematic diagrams that visually renders, in graphical form, the network configurations already defined and described by me; the content, structure and labels of the diagram were defined by me, and the AI-generated graphic was reviewed, adjusted and verified before inclusion Every output generated by the AI tool was subjected to rigorous critical review, verified for correctness and modified where necessary. AI was not used to falsify data, invent results, or generate unverified bibliographic references, and every cited source has been personally checked and verified.

FURTHER DECLARES that the intellectual and scholarity contribution to this thesis remains entirely my own. I have independently defined the research problem, formulated the research questions, designed the synthetic dataset and interpreted the empirical results. I assume full and sole personal responsibility for the quality, accuracy and methodological validity of the final work

INDEX

INTRODUCTION **1**

1 INVENTORY MISALLOCATION IN LUXURY RETAIL **2**

1.1 SCARCITY VS AVAILABILITY 2

1.2 THE COST OF REACTIVE LOGISTICS 8

1.3 PERFORMANCE INDICATORS FOR MEASURING INVENTORY EFFICIENCY 12

2 ANTICIPATORY LOGISTICS: AN OPERATIONAL SOLUTION **17**

2.1 FROM REACTIVE TO ANTICIPATORY LOGISTICS 17

2.2 WHERE ANTICIPATORY LOGISTICS CREATES VALUE IN LUXURY RETAIL 21

2.3 A DECISION-MAKING FRAMEWORK FOR STOCK PRE-POSITIONING 27

3 PYTHON SIMULATION: ECONOMIC EVALUATION **34**

3.1 SIMULATION DESIGN: TRADITIONAL VS ANTICIPATORY 34

3.2 SIMULATION INPUTS AND ECONOMIC CALCULATION FRAMEWORK 43

3.3 RESULTS COMPARISON AND DECISION MATRIX 48

CONCLUSION **59**

BIBLIOGRAPHY **61**

APPENDIX — SIMULATION CODE AND INTERACTIVE INTERFACE **63**

INTRODUCTION

The luxury industry sells scarcity and depends on availability. Pricing power rests on restricted access, yet no sale is recorded unless the article is present in the boutique when the client enters, within the window in which the client wants it. The tension is structural. The problem analyzed here is not how much inventory a brand holds, but where and when it is committed. Under r*eactive logistics,* allocation follows demand rather than precedes it and the resulting failures (stock-outs, overstock, markdowns and urgent shipments) carry an indirect cost structurally larger in luxury than elsewhere: a stock-out can damage the relationship with the *Very Important Clients on whom a disproportionate share of revenue depends.* The alternative is to commit in advance, on the basis of an early demand signal. This is *anticipatory logistics.* Its appeal is intuitive; its economic justification is not. Anticipation must pay for a forecasting capability and absorb the cost of committing wrongly, and the literature that establishes its feasibility does not establish when it is worth its price. The question this thesis addresses is therefore narrower and more practical than whether anticipatory logistics works: under what conditions does it generate an economic advantage large enough to cover what it costs, and how can a brand tell in advance which of its products, boutiques and selling windows satisfy those conditions?

The work is organized in three chapters. Chapter 1 defines the problem: the paradox between exclusivity and availability, the four failure modes of the reactive model and the five indicators through which the loss becomes measurable. Chapter 2 presents anticipatory logistics as the operational response, identifies the conditions under which it creates value and derives a decision rule for stock pre-positioning from the newsvendor formulation. Chapter 3 subjects that rule to a quantitative test. A simulation prices three stocking policies against identical demand across eight combinations of product, location and selling window, measures how far the outcome depends on each assumption behind it, and converts the result into two instruments a brand can apply: the maximum cost at which a forecasting capability still pays for itself, and a decision matrix that assigns a recommendation to combinations the simulation did not cover. Two clarifications are required at the outset. The dataset is synthetic, so the analysis demonstrates a mechanism and quantifies the conditions under which it operates; it is not an empirical estimate of any brand’s returns.

Chapter 1

1 INVENTORY MISALLOCATION IN LUXURY RETAIL

Luxury is by definition a business built on contradiction. This chapter establishes the analytical foundation of the thesis. Section 1.1 defines the luxury concept and examines the structural paradox between exclusivity and availability. Section 1.2 analyses the cost of reactive logistics by identifying four measurable failure modes. Section 1.3 translates these failures into five key performance indicators. Together, these three sections define the problem that the rest of this thesis sets out to solve.

1.1 SCARCITY VS AVAILABILITY

The starting point of any rigorous analysis of luxury retail operations must be a precise definition of what "luxury" means as an economic and strategic category. This is not a philosophical digression. It is a mandatory step, which must be taken in order to understand some critical elements related to the luxury industry. Based on the luxury concept, it will be possible to comprehend how the supply chain has to be designed, how inventory has to be allocated, and why the rules that are often used to govern mass-market distribution fail when applied to luxury goods. In common language, 'luxury' is often used as a word to describe 'expensive' or 'high quality' products and services. This improper use of the word can mislead any individual. For instance, a business class airplane ticket represents a premium-priced service, whereas a handcrafted Swiss watch depicts high-quality manufacturing. Both of these cases, however, do not capture the relevance of luxury as a strategic dimension. According to the *Luxury Strategy*, luxury is defined as:

> *<<the highest level of premium brands, characterized by a combination of heritage, craftsmanship, exclusivity, and a high price point>> [1]*

The elements presented in this definition are called the four pillars of luxury, and they are inseparable. Removing any one of them can alter the strategic relevance of the term. For instance, a product that has heritage and craftsmanship but lacks exclusivity is what we properly call a 'premium artisanal good'. Furthermore, a product that has exclusivity and it is placed on the market at a high price, while lacking heritage, is a 'speculative asset'. Luxury requires all four dimensions to function simultaneously. This distinction between luxury and premium is not merely academic, it has direct operational consequences. Based on the definition above, it becomes clear that luxury is not simply perceived as 'more premium'. The logic that governs premium brands is the following: produce more when demand rises [1], adjust the price to properly reflect market conditions, and maximize distribution coverage whenever possible. There is a limit on the price that premium brands can place on their products, and any increase must be

justified. For luxury products, this specific rule does not apply, because those goods are perceived as art. Luxury brands operate under an entirely different strategic grammar, one in which the conventional rules of marketing must be inverted. This inversion is not accidental. Luxury brands compete not on quality but on a far more particular dimension: the latent social hierarchy of good taste and symbolic capital [1]. The difference relies on the brand's ability to generate social distinction. What is being sold, in other words, is not the object itself but the meaning embedded within it. A Hermès Birkin bag and a high-quality leather bag may be materially similar in function; yet they are worlds apart in what they communicate. This symbolic dimension is the engine of luxury value, and it is extraordinarily sensitive to how the product is distributed, allocated, and made available in the market. Historically, luxury has always served a stratifying social function [1]. This sociological dimension is critical for understanding inventory strategy. If luxury's role is to create social differentiation in a meritocratic society, then the moment a luxury product becomes genuinely accessible to everyone, it ceases to perform its primary function. The product may remain beautiful and well-crafted, but it loses its ability to signal distinction and with it, its capacity to command a luxury price.

Now that the structural definition of luxury has been established, it is possible to proceed by presenting the central operational paradox of the industry.

<<Luxury brands are simultaneously required to grow their revenues and to preserve the perception that their products are not widely available>>[1]

These two objectives are, in their pristine form, contradictory. Achieving a balance between those two goals, is the defining challenge of luxury retail operations. In this specific case, the contradiction is structural. Luxury conglomerates such as LVMH, Kering, and Richemont are publicly listed entities with shareholders and with growth targets. Their investors expect increase in terms of revenue year over year. Now, a question naturally arises: what is the impact of revenues' growth? The mechanism through which these revenues are generated (the perception of exclusivity and desirability) is weakened every time the brands enlarge their distribution footprint and enrich their stock levels. As a consequence, the product seems to be easier to obtain. Growth, left unchecked, destroys the foundation on which it is built. The well-known *Luxury strategy* aims to balance significant growth with sustained luxury positioning. The key lies in creating enduring, long-selling products rather than short-lived best-sellers. The Porsche 911 (in production since 1964) and Chanel N°5 (launched for the first time in 1921) are both concrete examples that can be used to explain this concept. Those products have sold , continuously, for decades without losing their luxury status. This result was possible because their availability was never maximized. In addition, their production was never scaled purely in response to demand spikes. This temporal dimension (the long-term orientation of luxury strategy) is in net contrast with the reactive, demand-driven logic of conventional supply chain management. In a standard retail supply chain, the goal is to have the right product in the right place at the right time to satisfy demand at the

lowest possible cost. In luxury, this principle must be modulated: having the right product in the right place at the right time for the right customer, in a quantity that preserves the perception that not everyone can obtain it. The addition of that final part, changes almost everything about how inventory must be managed. This counterintuitive principle can be summarized as: Do not respond to rising demand [1]. Instead of overextending to meet growing demand, luxury brands can focus on maintaining exclusivity by deliberately limiting their client base. It means that a sold-out product is not , automatically, a failure of inventory management. It may be a desired outcome, that could reinforce the brand's scarcity narrative and sustains the desire of those who could not obtain it.

For a long part of the luxury's history, scarcity was both: physical and natural. Small ateliers produced limited quantities of handcrafted objects and the supply was genuinely constrained by the number of skilled artisans. Furthermore, the time which was required in order to produce a good, was very high. In this complex context, exclusivity was not a marketing strategy, it was simply a fact. The last century was the turning point which has transformed this equation. Throughout globalization, luxury brands were able to consolidate into large conglomerates, and to expand across dozens of countries. As a direct consequence, their production necessarily scaled up. The luxury industry has moved from scarcity to rarity, produced by feelings of exclusivity and uniqueness [1]. This latest concept has been, then, called ‘abundant rarity’. The massive growth of today's luxury sector directly challenges the traditional definition of luxury; if physical rarity were still essential, the current market would fail to meet the criteria. This pivot (from physical rarity to perceived rarity) is arguably the most important strategic innovation in modern luxury management. It means that a brand can sell millions of units globally and still maintain a perception of exclusivity, as long as it is able to manage how, where, and to whom those units are made available. The tools which are commonly used, in order to spread a sense of perceived scarcity could include: limited editions that create urgency and reward loyal customers for waiting for iconic products such as the Hermès Birkin. Luxury brands abandoned product rarity as the precondition of luxury and often adopted the latest abundant rarity strategies [1]. Rarity , today, is driven by the perception of exclusivity instead of physical limitations. A large number of elements, should be taken into account while executing an abundant rarity strategy. It is not enough to produce fewer units. True luxury execution requires aligning product, channel, and customer perfectly. If a scarce item lands in the wrong market or targets the wrong demographic, the strategy has failed, regardless of sales success. This is precisely why inventory allocation (the choice of how many units of each product to send to which location, through which channel, at which point in the season) is a brand management decision as much as it is a logistics decision. Every allocation choice, implicitly communicates something about who the brand believes deserves access to its products. Get it wrong, and the brand either suffers stock-outs in key markets (lost revenue, frustrated top clients) or oversaturation in others (diluted exclusivity, damaged brand perception). The tolerance for error is narrow, almost close to zero. From a risk management point of view, the mismanagement of perceived scarcity can be interpreted as an

operational risk rather than only a branding issue. In this thesis, inventory misallocation is considered as a risk event because it can affect three essential goals at the same time: revenue generation, margin protection and brand positioning. The risk is real because it can materialize through two opposite outcomes (often connected). The first one is stock-out in high-value markets (where demand exists but the product is not locally available). This generates lost sales, lower customer satisfaction and possible deterioration of the relationship with top-tier clients. The second is oversaturation in low-demand markets, where excess inventory may weakens perceived scarcity and require markdowns or internal transfers. In both cases, the allocation error produces measurable economic consequences. For this reason, the scarcity-availability paradox can be shifted into a risk assessment issue. The question is not only whether the brand is perceived as exclusive, but whether the inventory is allocated in a way that minimizes the probability of stock-outs, overstock and corrective logistics costs.

The luxury paradox has been further complicated by the rise of digital retail. Online channels, based on their nature, maximize the total availability of a product. A good which is listed on a brand's e-commerce platform is , theoretically, accessible to any consumer in any market with an Internet connection and a credit card. This is structurally in contrast with the logic of selective distribution that underpins luxury positioning. The tension is not theoretical. Omnichannel retail imposes a cost on top-tier consumers by diluting exclusivity, while leaving mass markets unaffected [3]. For these high-value clients, whose purchases are driven purely by social signaling, unrestricted online availability directly erodes the prestige of ownership. Their satisfaction is partially related on others not having the same access that they do. For these consumers, the online channel is not simply a convenient new way to shop: it is a threat to the symbolic value of what they are purchasing. At the same time, those individuals are highly sensitive to physical stock availability. Once consumers' belief in inventory availability is higher, they have a higher utility [3]. This creates a nuanced and demanding operational requirement: luxury brands must ensure robust product availability in physical stores for their top-tier clients because unavailability in the boutique damages the relationship and the service experience. While simultaneously limiting the perception of broad, democratic online access. The same product must feel exclusive in the digital space and reliably available in the physical one. Managing this dual requirement is technically complex and strategically delicate. The 2024 Bain-Altagamma market study confirms that brands which have failed to manage this balance correctly are experiencing measurable consequences. Among the top-tier *Very Important Client* (VIC), who represent barely 2% of the luxury customer base but account for 45% of total luxury purchases, the reported satisfaction has declined. Those group of people describe their brand interactions as transactional rather than exclusive and personalized [2]. This erosion of the relationship with the most valuable segment is directly linked to the perception that access to luxury products has become too democratized, too available, and insufficiently differentiated.

At a first glance, it is possible to observe that the tension between scarcity and availability has created a financial consequence. The 2024 Bain-Altagamma data, provides a clear representation of what happens

when this balance breaks down. First of all, the personal luxury goods market has contracted for the first time in 15 years, reaching €363 billion, a decline of approximately 2% , if compared to 2023. The global luxury customer base has decreased of , approximately, 50 million people in just two years (going from 400 million to 350 million)[2]. Based on this data, it is apparent that only one-third of luxury brands has managed to grow their revenues in 2024. Average EBIT margins have slightly decline going from 21% to 18–19%. The entire industry's supply chains faced a significant contraction, absorbing volume drops of 20% to 25% from 2022 to 2024. This has created a scenario, which has forced brands to urgently reduce stock levels to free up liquidity [2]. These are not irrelevant numbers. This downturn reflects an industry-wide reckoning, exposing the consequences of mismanaging the delicate balance between exclusivity and availability during the post-Covid boom. The decisions made, in terms of pricing and availability, done during the period (going from 2022 to 2023), are the real cause of this situation. Brands decided to raise prices due to the inflation in raw material and the surge in demand of the period. However, they had not done investments in order to increase the perceived value, the craftsmanship or storytelling. The result obtained, was that more than 50% of luxury consumers, (by 2024), considered luxury brands overpriced. The dream had been monetized too bluntly, and consumer trust had eroded. Overexposure and overextension can influence negatively the brand's perceived value [1]. The data shown by the Bain report, are the empirical confirmation of this theoretical principle. Aggressive expansion, broad distribution, and high accessibility inevitably erode the psychological foundation of luxury desire, even when product quality is uncompromised. Scarcity is painful, but rarity, if well managed, can become a dream. The key concept on which the luxury strategy is based, is to never sacrifice long term to short term [1]. This principle is precisely what the 2024 market data illustrates in financial terms.

According to what has been discussed above, it is possible to say that, in luxury retail, stock allocation is not a logistics function, it is a brand function. Operational decisions regarding unit allocation, market focus, distribution channels, and client access are never purely logistical. They , directly, define the brand's identity, positioning, and long-term desirability. A luxury product is effective on an intangible dimension. It gives prestige and social distinction to the owner. Those perceptions depend heavily on how accessible a product feels. As a consequence, they are directly shaped by operational inventory allocation choices. From this perspective, the conventional supply chain rules (such as maximizing service levels, minimizing stock-outs, and optimizing turnover) must be adapted to the luxury context. While delivering every product on demand is ideal for fast-moving consumer goods. In luxury, it risks undermining the brand's strategic goals. Similarly, a zero stock-out rate for an iconic item is not a success metric. It signals overproduction and a failure to manage scarcity It is crucial for luxury firms to make decisions on several dimensions (product positioning, pricing, and inventory stock) especially, when selling luxury goods with consumers, in an uncertain market condition [3]. The integration of inventory decisions with brand positioning decisions (considering them as inseparable) is the strategic shift that

separates effective luxury supply chain management from reactive logistics management. In conclusions, the tension which can be identified between scarcity and availability (in luxury retail) is not a problem to be solved. It is a condition to be managed. Its existence is structural and permanent: as long as luxury brands must grow revenues while preserving exclusivity, the paradox will persist. Once scarcity is managed through geographically distributed retail networks, the key risks becomes misallocation: stock may be excessive in low-demand locations and unavailable in high-value locations. This is where reactive logistics becomes costly.

1.2 THE COST OF REACTIVE LOGISTICS

In the first place, it is essential to define what reactive logistics means as an operational model. In this case, we are not discussing a theoretical abstraction. Instead, we are presenting the dominant mode through which most retail supply chains, including luxury ones, still operates today. Reactive logistics is , a particular, supply chain model that responds to events after they have already occurred. Before the beginning of the selling season, inventory is produced and distributed in advance based on a static forecast. Once the initial forecast is translated into stock allocation and inventory is dispatched to stores, the supply chain operates with rigid adherence to the plan. This traditional system does not take into account real-time demand signals. It waits until a problem, such as a stockout, is already clearly visible. As a consequence, a corrective action can be made only with the purpose of mitigating a damage and not with purpose of avoiding it. The major cause of this behavior has been named as the *lead-time gap*. It can be defined as:

<<the difference between the time the supply chain needs to source, produce and deliver a product, and the time the customer is willing to wait>> [4]

When the lead-time gap is large, companies have to build inventory **buffers** based on forecast. The longer is the gap, the earlier those forecast must be made and the higher will be the probability of having a misalignment between reality and the plan. In luxury retail, where products' lifecycles last one or two seasons (with customers who prefer to shift rapidly from a product to another), this gap is particularly consequential. Inventory that has been committed four to six months before the selling window, can generate diverse scenarios. If demand does not materialize, a surplus in stock can be noticed. On the other hand, if demand exceeds expectations in specific location, a series of emergency corrections will be needed. The reactive model is not negligent. It is the natural consequence of a supply chain which has not been designed to analyze, read and respond to real-time demand. The problem relies on the structure.

A stock-out, generally, occurs when a product (that a customer is willing to purchase) is not available at the point of sale. It is the most visible symptom of reactive logistics and the one most extensively studied in the academic literature. Stock-out, often, generates important losses for both manufacturers and retailers. The magnitude of these losses depends on specific customers reactions, which can be different based on product or situation factors [6]. In response to a situation of stock-out, three possible behavioral responses can be identified:

1) Substitute: the customer purchase an alternative product within the same store
2) Delay: the customer postpones the purchase and returns later
3) Leave: the customer abandons the purchase entirely or goes to a competitor [6]

Each one of these reactions, can lead to diverse costs for the brand. For instance, substitution can preserve the revenue of the brand, but it can damage the relationship between the brand and customer

himself (especially if he has to accept a second-choice product). In addition, delay preserves the intention but it does not generate any immediate revenue. At the same time, through delay, a risk of non-return can be introduced. The third scenario, is the most damaging. Leaving is based on a combination of revenue loss and long-term customer attrition. In order to be clear, there are two variables which can exercise a positive impact on the customer reaction to stock-out. They are: loyalty and the presence of product alternatives [5]. In luxury, those two factors interact in an inherently disadvantageous manner. Loyalty has a threshold, for that reason repeated stockout experiences can erode it progressively. Meanwhile, the availability of alternatives is logically limited in the luxury field. A product that is considered to be irreplaceable within a brand, it will not have a soft landing when it is out of stock. As a matter of fact, the customer cannot be redirected to an equivalent good. The remaining options are delay or leave. Furthermore, the financial damage of a stock-out extends well beyond the immediate missed sale. Customers who have experienced stock-outs ordered 22% (on average) [5] fewer units in the following 13 months compared to those who did not. This effect has not diminished at the end of the period observed. It means that the true long run cost is likely larger than the one recorded. For luxury brands, losing a single top-tier relationship with a customer due to a stock out can be defined as an operational failure [8].

If stock-out depicts the most visible failure of reactive logistics, overstock is the most structurally widespread. Overstock occurs when inventory has been allocated in excess compared to the actual demand. An inevitable consequence of a forecast made too far in advance, without having the possibility of correct while demand signals were emerging during the season. Stock-out may alter each of the three consumer decisions: whether to buy, how much to purchase and what to choose [7]. Instead, overstocks can undermine the principle of exclusivity which we have discussed in the section 1.1. An excess in inventory can force brands to adopt strategies such as promotions or clearance. Those actions generate a signal for the market: the product was not as scarce or as desirable as originally positioned. In luxury, overstock is not merely an operational efficiency. A product that remains on the shelf beyond its intended selling window, or that circulates through outlet channels at reduced prices, communicates to the customer that the brand miscalculated demand and could not sell its products at full price. This is the mere opposite of the scarcity message that luxury brands want to share in order to sustain their pricing power. This dynamic has been framed within the concept of the total costs of logistics which is:

<< the full sum of costs across the entire pipeline, including inventory carrying costs, transportation costs, and the downstream cost to serve triggered when overstock must be corrected through markdowns or transfer>>[4]

In a reactive model, these costs are not exceptional, they are structural. They occur every season because the forecast method which is used it is not enough accurate to avoid them. The overstock problem is further amplified by the Bullwhip effect:

<< small fluctuations in retail demand cause increasingly distorted signals as they travel upstream through the supply chain, pushing suppliers and manufacturers to overproduce based on inflated orders>>[4]

When overstock has already accumulated or a stock-out has already occurred, a reactive supply chain responds with two corrective mechanisms: markdowns and internal transfers. Both are expensive and both are solutions that can be adopted only when the mistake has been identified rather than preventing it. Markdowns are price reduction which are applied to slow-moving or excess inventory to accelerate its clearance before the ending of the selling season. In luxury, they carry costs that go beyond the margin erosion of the discount itself. A customer who observes that a luxury product that they previously purchased at full price is now available at a discount, can prove a sense of fairness violation [8]. Their confidence in the brand's pricing integrity is damaged. On the other hand, we have internal transfers which are

<<the movement of stock between stores or distribution centers to correct a mismatch between where inventory is and where demand is>> [4]

They are triggered when the initial allocation decision was wrong and, as we know, no anticipatory mechanism exists to prevent the mismatch. This dynamic is part of a broader failure. Each transfer carries direct costs which space from transportation, handling, processing time to temporary unavailability of stock during transit. In luxury, when a product is in transit it means that it is unavailable to a client who may be standing in front of an empty shelf (the cost of transit is amplified). The cumulative effect of markdowns and internal transfers is not just financial, it is organizational. Reactive logistics is self-reinforcing: the cost of reacting consumes the resources that would otherwise fund the transition to anticipation. The last element analyzed is the well-known urgent shipment. In order to have a clear view about this topic, a proper definition is needed:

<< an expedited delivery triggered by a stock-out that has already occurred, requiring faster and therefore more expensive transportation to correct a misallocation>>

Urgent shipments are a clear signal that the supply chain is operating reactively. They are by definition exceptional costs that are treated as normal operating expenses rather than as indicators of a structural failure in the inventory management model. The full cost of fulfilling an order from source to customer, including emergency actions, is called cost-to-serve [4]. This concept is crucial for luxury: the cost of serving a top-tier client via an emergency transfer shipment, may easily exceed the margin earned on the sale. Once the supply chain correction costs are included, transactions that appear profitable on their face may in fact be margin negative. The paradox of urgent shipments in luxury is that they are paid twice. The first one is observed on a financial dimension, the second one instead is reputational (relationship with the brand tested because of a lack in terms of availability). Both payments are avoidable. The four failure models analyzed in this section are not isolated incidents. They are the

predictable output of a supply chain that commits inventory too early and corrects too late. Section 1.3 will introduce the five key performance indicators that allow a brand to quantify inventory inefficiency in precise economic terms, turning a qualitative problem into a diagnostic tool.

Table 1: ***Overview of the four failure modes***

FAILURE MODE	DEFINITION	PRIMARY COST	BRAND COST IN LUXURY
STOCK-OUT	Demand cannot be met at the point of sale	-22% future purchasing per affected client	Erosion of customer's loyalty
OVERSTOCK	Inventory exceeds actual demand at the location	Carrying cost + obsolescence	Damage on exclusivity and scarcity
MARKDOWN	Price reduction to clear excess inventory	Margin erosion on cleared units	Price fairness/power violation
URGENT SHIPMENT	Expedited delivery triggered by stock-out	Premium freight + handling	Client's trust damaged

Source: own elaboration

1.3 PERFORMANCE INDICATORS FOR MEASURING INVENTORY EFFICIENCY

According to section 1.2, reactive logistics produces four measurable failure modes: stock-outs, overstock, markdowns, and urgent shipments. However, it is crucial to quantify those failure modes before designing any type of corrective strategy. A problem that cannot be measured cannot be managed and in inventory management, imprecise measurement is a certain source of cost. This section introduces the five key performance indicators that translate the qualitative failure of reactive logistics into concrete, comparable numbers. Each Key Performance Indicator (KPI) as a role and answers a specific operational question. Together, they generate a diagnostic toolkit that any luxury brand must deploy. First of all, the analysis will start from the stock-out rate. It measures the frequency with which customer demand cannot be met due to the absence of the requested product at the point of sales. This concept is simplified by this formula:

$$\text{Stock-Out Rate} = \frac{\text{Number of stockout events}}{\text{Total demand opportunities}} \times 100$$

A demand opportunity (on the denominator) is any instance in which a customer requests a product. Instead, a stock-out event is any instance in which that request cannot be fulfilled thanks to the available inventory at the location. Commonly, the stock-out rate is expressed in percentage. A rate of 10% means that one purchase (on a maximum of ten) has failed due to unavailability [8]. A stock-out rate must be read, in conjunction, with the product. Not all stock-outs carry the same weight. A 5% stock-out rate focused on entry-level products (or secondary products) represents a moderate operational inefficiency. Instead, the same rate applied on iconic items, such as a specific bag, represents a strategic failure. The reason lies on what happens to the customer when the product in unavailable. In addition, there is a complementary metric which should be taken into account in order to better analyze the stock-out. This direct complement is named as Service level. Its formula is the following:

$$\text{Service Level} = \frac{\text{Orders fulfilled without stockout}}{\text{Total orders received}} \times 100$$

Service level is defined as the probability of meeting all customer demand during the time window between one inventory replenishment and the next [10]. A service level of 95% coincide to a stock-out rate of 5%. In the luxury field, where a failed service event can damage a long term relationship with the customer, a service level below 95% should be perceived as a critical issue (requiring immediate

intervention). For luxury brands, a stock-out rate below 5% represents the operational target. This threshold is derived from the service level requirement of 95% recommended by [10] for high-value product categories. The second KPI, which has been taken into consideration is the inventory turnover. It measures the number of times a company sells and replace its entire inventory stock within a given period (typically one fiscal year or one selling season). It is calculated through this method:

$$\text{Inventory Turnover} = \frac{\text{Cost of Goods Sold (COGS)}}{\text{Average Inventory Value}}$$

A high inventory turnover indicates that the stock is moving efficiently through the supply chain and capital is not being tied up in slow-moving units. On the other hand, observing a low inventory turnover can signal the opposite: inventory is accumulating (occupying warehouse space and consuming carrying costs). Inventory turnover is, often, identified as the primary measure of how effectively the capital invested in stock is being converted into sales [4]. Luxury typically operates with inventory turnover ratios between 2 and 4 times per year [9] (it is considered consistent with healthy luxury operations), reflecting the deliberate scarcity strategy discussed in the section 1.1. A ratio below 2, signals structural overstock accumulation and potential markdown risk. A ratio above 4, in luxury, may indicate insufficient inventory depth to serve top-tier clients [9]. By contrast, fast fashion retailers operate with a average turnover between 6 to 12. For a luxury brand, a prior target is not to maximize turnover, but to maintain it within a specific range which should be consistent with both: financial efficiency and brand positioning. A falling inventory turnover ratio across consecutive seasons is an early warning signal of overstock accumulation. It highlights that a company is buying more than what they are selling. A complementary metric is: Days of Inventory Outstanding (DIO). Based on this formula:

$$DIO = \frac{\text{Average Inventory Value}}{\text{COGS}} \times 365$$

DIO expresses the turnover in calendar days: how many days of sales are currently sitting in inventory. A DIO of 90 means the brand holds three months of stock. This indicator is useful for highlighting dead stocks (inventory that has been sitting in the network beyond the expected selling window and it is unlikely to sell at full price). A DIO between 90 and 180 days is therefore the operational healthy range for luxury brands (differences could be observed based on the product category) In order to have a clear view on this topic, it is essential to introduce the GMROI (Gross Margin Return on Inventory Investment). This KPI allows to understand the gross profit , which has been generated, for every unit

of currency invested in inventory. GMROI is a metric that embeds both the margin dimension and the turnover dimension into one number. It is calculated through this formula:

$$GMROI = \frac{\text{Gross Margin}}{\text{Average Inventory Cost}}$$

In few words, the GMROI is useful for understanding if an investment has been profitable or not. Having a GMROI higher than 1.0 means that the company is generating more gross margin than what they are spending in inventory. On the contrary, below 1.0 means that the investment is destroying value (the margin earned does not cover the cost of the inventory required to generate it). For instance, a product with a gross margin of €30,000 and an average inventory cost of €20,000 produces a GMROI of 1.5: for every euro invested in that stock, the brand earns €0,50 in gross margin. GMROI can be decomposed into:

$$GMROI = \text{Gross Margin \%} \times \text{Inventory Turnover}$$

This formula is quite powerful for the luxury industry. It makes explicit the structural tension between high margin (a luxury strength) and low turnover (a deliberate luxury constraint). A product with a 70% gross margin and with an inventory turnover of 0.8 produces a GMROI of 0.56. A loss-making inventory investment, despite the apparently attractive margin. This is the mathematical representation of the overstock problem: margin alone does not determine inventory profitability if the stock is not turning fast enough to justify the capital tied up. This KPI should be calculated at a Stock Keeping Unit (SKU) level. This granularity is helpful in identifying: high GMROI products (goods that should be prioritized in inventory allocation and never allowed to stock-out), low GMROI products (slow-moving items consuming capital and space without a proportionate return) and channel-level GMROI gaps (store or market that generates the highest return per unit of inventory invested). The minimum acceptable GMROI is 1.0 (below the investment is destroying value). A GMROI between 1.5 and 3.0 is considered to be a healthy range for luxury brands. A value above 2, proves the brand's immense pricing power, generating a high cash return for every euro invested. It is important to take into consideration the total costs of logistic. This is why, the Logistics Cost per Unit Sold will be discussed. It measures the total logistics cost incurred to bring one unit of product from the source to the final customer (including all standard and emergency logistics action).

$$\text{Logistics Cost per Unit} = \frac{\text{Total Logistics Costs in Period}}{\text{Units Sold in Period}}$$

Total logistics cost must include: inbound transportation, warehousing and storage, outbound distribution, order processing, returns handling and all emergency costs. This KPI (that has been previously defined as the cost-to-serve) is the operational measure that makes the hidden cost of reactive logistics visible. A company which is able to allocate inventory efficiently, incurs only on standard logistics cost per unit. A company operating reactively accumulates a layers of emergency logistics costs on top of its standard base. The financial impact is huge. If a product carries a gross margin of €200 and a standard logistics cost of €15 per unit, the net contribution will be €185. If a reactive misallocation forces the company to incur in an additional €40 of emergency transfer and expediting costs for 20% of units sold, the blended logistics cost per unit rises to €23, compressing the net contribution to €177. A high emergency logistics signals that initial allocation was systematically wrong and that the supply chain is spending money on corrections. According to [9], total logistics costs in retail operations typically represent between 8% and 15% of net revenue. For luxury brands, the target should sit at the lower of this range (between 8% and 10%) given the high unit margin characterizing the sector. The emergency logistics premium should not exceed 10% of the total logistics budget. In conclusion, it is possible to observe a KPI that allows to measure the true profitability of a product, category or channel, after all logistics costs have been fully deducted from gross margin. It is calculated as:

$$\text{Net Margin After Logistics} = \text{Gross Margin \%} - \frac{\text{Total Logistic Costs}}{\text{Net Revenue}} \times 100$$

It links the supply chain performance to the financial results. The net margin after logistics should be calculated at the level of the individual store, market and channel. A luxury brand that calculated net margin after logistics by location will be able to highlight the stores that are both the most productive and the largest, in terms of logistics cost generation. A net margin after logistics above 15% is considered healthy for luxury brands. A margin between 10% and 15% signals that logistics cost are eroding profitability. In conclusion, the five indicators work together as a diagnostic system in which each KPI explores different perspectives of the same underlying issue. No single KPI tells the full story of inventory efficiency. When these KPIs are read in combination, a pattern emerges. A company that is experiencing simultaneously: a rising stock-out rate, a declining inventory turnover, a deteriorating GMROI, a growing logistics cost per unit and a compressing net margin. It is dealing with one problem having 5 measurable symptoms. That problem is reactive logistics. The purpose of establishing this KPI baseline is not to generate numbers, but it is to create a diagnostic foundation. You cannot improve what

you do not measure. In conclusion, this chapter has mapped the structural inefficiencies and hidden costs generated by reactive logistics. The following chapter will introduce anticipatory logistics as the operational solution designed to mitigate these financial leaks through predictive data-driven stock pre-positioning.

Table 2: ***Overview of the different KPI's useful for the analysis***

KPI	FORMULA	HEALTHY RANGE
STOCK-OUT RATE	$\frac{\text{Number of stockout events}}{\text{Total demand opportunities}} \times 100$	<5% on core SKU
INVENTORY TURNOVER	$\frac{\text{Cost of Goods Sold (COGS)}}{\text{Average Inventory Value}}$	2-4x per year
GMROI	$\frac{\text{Gross Margin}}{\text{Average Inventory Cost}}$	1.5-3.0
LOGISTIC COST PER UNIT	$\frac{\text{Total Logistics Costs in Period}}{\text{Units Sold in Period}}$	8-10% of net revenue
NET MARGIN AFTER LOGISTICS	$\text{Gross Margin \%} - \frac{\text{Total Logistic Costs}}{\text{Net Revenue}} \times 100$	>15%

Source: own elaboration

Chapter 2

2 ANTICIPATORY LOGISTICS: AN OPERATIONAL SOLUTION

Chapter 1 has established what is the real problem in precise terms. Reactive logistics (the dominant operational model in the luxury industry) generates four measurable failure modes: stock-outs, overstock, markdowns and urgent shipments. Each of these failures carries both a direct financial cost and an indirect brand cost that is structurally more damaging in luxury than in any other retail segment. The question that follows is purely operational: are there alternatives? How they work? This chapter has the purpose of answering those questions. Section 2.1 will define anticipatory logistics and the transition from the old model to a new and more efficient one.

2.1 FROM REACTIVE TO ANTICIPATORY LOGISTICS

Anticipatory logistics is defined, in the academic literature, as the shift from a supply chain model driven by orders already received to a model driven by the prediction of future customer needs (acting before demand fully manifests, rather than responding after the fact) [11]. The most precise definition, which is available, is the following:

<< *Anticipatory logistics incorporates computer to analyze and search data, view purchase history and behavioral signals to predict what and when a customer will want something, and then position inventory accordingly- in some cases shipping products before the order has even been placed* >> [11]

Table 3: ***Reactive vs Anticipatory (model comparison)***

DIMENSION	REACTIVE MODEL	ANTICIPATORY MODEL
DECISION TRIGGER	Realized shortage or excess	Predicted demand imbalance
TIMING	After the mismatch emerges	Before the mismatch materializes
MAIN ACTION	Emergency correction	Planned repositioning
INFORMATION BASE	Orders and realized sales	Forecasts, demand signals and external data
MAIN RISK	Late and costly correction	Forecast-driven misallocation

Source: own elaboration

This definition is highlighting two essential factors. In the first place, anticipatory logistics is not a mere theorical concept: it is a concrete operational system that uses data as its primary input. In the second place, the unit of action is not the order but it is the prediction. The supply chain moves on the basis of what is expected to happen, not on the basis of what has already happened. This clear details are already putting in evidence the difference between anticipatory and reactive logistics (introduced in the Chapter 1).

Anticipatory logistics *and predictive analytics* are closely related, but they should not be treated as synonyms. There is a slightly difference between them. Anticipatory logistics refers to the strategic posture (the decision to act before demand manifest) while predictive analytics refers to the technical infrastructure that makes the anticipation possible (forecasting models and automated decision systems).

The intellectual roots of anticipatory logistics lie on the concept of the *agile supply chain*. An agile supply chain is, usually, defined by four properties: it is market sensitive (it reads and respond to real demand rather then relying exclusively on forecasts), it is information based (it uses data to reduce the need for physical inventory buffers), it is network integrated (all partner in the chain share information and coordinate actions) and it is virtually integrated [14]. The agile supply chain model depicts a sort of break from the push-based model that we have earlier identified as the main cause of reactive logistics costs. On the other hand, agility alone does not fully explain what modern data-driven supply chain are capable of. An agile supply chain responds quickly to demand signals once they are quite visible. An anticipatory supply chain instead detects the early, still partial signals that precede a demand acceleration and estimates its likely course before that acceleration is fully reflected in the sales data. The transition from agile to anticipatory logistics is a transition in the time horizon of action. Agility minimizes the lag between signal and response. Anticipation reduces the decision lag by acting on predicted demand before the demand imbalance becomes fully visible in sales data. In operational terms, the difference is the following. An agile luxury brand detects that a specific bag is selling faster than expected and initiates a inter-store transfer within 48 hours. Instead, an anticipatory luxury brand detects one week before the selling acceleration become visible in the sales' data, that the combination of a recent editorial feature and a surge in online wishlist additions for that specific SKU is likely to produce a stockout within two weeks. As a consequence, inventory is repositioned accordingly, before the acceleration begins.

The transition from reactive to anticipatory logistics is not achievable through organizational intent alone. It requires a specific technical infrastructure. For instance, point of sale data provides real time visibility of what is selling, where, and at what velocity. Client Relationship Management (CRM) data (such as purchase history, preferences, visit frequency…) allows brand to anticipate individual client demand at the level of the specific product variant. External data sources (like social media, weather patterns, local event calendars) provide early indicators of demand shifts before they appear in sales data. AI is empowering luxury and fashion brands to predict customer demand more precisely. This

predictive capability allows companies to efficiently manage inventory levels, reducing the risk of overstocking or stockouts [11].

This can be summarized in operational terms: the data sources (listed above) feed a predictive model whose output is not a single number but a small set of operationally usable estimates, which in turn change a specific logistics decision. Table 4 summarizes the correspondence. Chapter 3 implements this logic through a Python simulation rather than a trained machine-learning model; the correspondence below describes the principle, not a specific algorithmic architecture.

Table 4: ***From predictive output to operational decision***

INPUT DATA	PREDICTIVE OUTPUT	OPERATIONAL DECISION
HISTORICAL SALES	Expected demand	Base stock allocation
SELL-THROUGH VELOCITY	Probability of stock-out	Early reallocation
CRM SIGNALS	Expected client-specific demand	Boutique pre-positioning
EVENTS AND CALENDAR	Temporal demand uplift	Time-limited stock increase

Source: own elaboration

It is crucial to keep in mind two important distinction when AI-based forecasting systems (for luxury) are evaluated. In the first place, AI does not eliminate forecast mistakes (no system is capable of predicting demand without any error). What AI does is reduce forecast errors and reduce the time lag between a demand signal and the supply chain's response. In the second place, AI-based demand forecasting does not replace human judgment in luxury: it enriches it. The final decision remains under human responsibility, but it is supported by a broader and more timely information base.

The most advanced application of anticipatory logistics principles is depicted by *anticipatory shipping.* It stands for the practice of moving inventory toward its likely destination before the customer order is placed (based on demand predictions). As a consequence. the product is automatically shipped even before the order has been placed [11]. Amazon has patented a version of this concept in 2013 under the name of '*anticipatory package shipping*', in which products predicted to be ordered with high probability are pre-positioned in local distribution hubs to minimize delivery time once the order arrives. In luxury, anticipatory shipping acts on a different base and with a different client-rationale. The two variants must be kept terminologically distinct: anticipatory shipping in e-commerce moves the product near the individual customer before the order is placed, while anticipatory stock-positioning (in luxury retail) relocates the product to the point of sales judged most likely to generate it. Instead of shipping

products to a customer's home before they order, luxury brands use this method to pre-allocate inventory to specific stores or regions based on demand predictions (ensuring the product will be physically present in the right market before the client arrives, rather than transfers it after a stock out). The benefit seems to be the same: the client encounters a product that is available and the service is preserved. The client is not going to experience any type of friction.

The transition from reactive to anticipatory shipping is not an upgrade, it is a sort of structural evolution of how the supply chain processes and acts on information. It needs investments in three major areas: data infrastructure (ability of collecting data and integrate them in real time); analytical capability (AI and machine learning calibrated to the specific demand patterns of luxury products and clients); and organizational alignment (integration of demand intelligence into allocation and replenishment decisions).

2.2 WHERE ANTICIPATORY LOGISTICS CREATES VALUE IN LUXURY RETAIL

Section 2.1 has established the theorical and technical foundations of anticipatory logistics: what it is, how it works, and what infrastructure it requires. The question that follows is not about the technology, it is about selectivity. Anticipatory logistics is perceived to be a powerful tool, but applying it indiscriminately across an entire luxury portfolio would be both economically counterproductive and incoherent from a strategical perspective. The purpose of this section is to identify precisely where anticipatory logistics creates measurable economic value within the luxury retail context.

This identification is not arbitrary. The academic literature on luxury supply chain management provides a structural basis for segmentation. Research on supply chain strategy within the luxury industry has demonstrated that not all luxury companies face the same supply chain challenges. The most effective response depends on variables such as product fashion ability, product complexity and brand reputation [17]. The consequences for anticipatory logistics are direct: the value of acting before demand materializes depends entirely on whether demand is predictable enough to act upon, and whether the economic stakes justify the infrastructure investments. This section is going to translate that principle into a diagnostic framework identifying five specific conditions under which anticipatory logistics delivers its highest returns.

Table 5: ***Suitability matrix for anticipatory logistics deployment***

DIMENSION	VARIABLE	HIGH SUITABILITY	LOW SUITABILITY	ECONOMIC RATIONALE
PRODUCT	Demand predictability	Mature carry-over / iconic	New seasonal item	Lower forecast error
ECONOMICS	Unit contribution and GMROI	High	Low	Greater underage cost
LOCATION	Demand density	Flagship / tier 1	Peripheral / low-volume	Greater data and higher sales exposure
TIME	Demand concentration	Predictable peak	Irregular low-demand period	Higher probability of sell-through

CLIENT	Expected client value	VIC / high-frequency	Occasional	Greater relational underage cost

Source: own elaboration

Table 5 sets out, for each of the five dimensions, the variable which determines suitability, the high and low suitability poles and the underlying economic rationale. Each dimension is examined below in turn, together with the supporting evidence of the literature

The first and most crucial condition for the economic viability of anticipatory logistics is the product type. Within a luxury brand's assortment, a structural distinction exists between seasonal products and carry-over products. Seasonal items are designed for a single selling cycle, approximately twenty weeks [18]. Their demand trajectory is uncertain because each season represents a new product with no direct historical equivalent. On the other hand, carry-over products are sold across multiple seasons. After the initial launch phase, they enter a long maturity period in which sales stabilize at a predictable level, making it easier to forecast the end market demand [18]. The most extreme version of this category is the *permanent* or *iconic* product: an item that has been on the market for years, whose demand pattern is not only stable across seasons but also well-documented across geographies and client segments.

Empirical evidence reveal that predictability is not uniform. Analysis that have been conducted across six luxury fashion companies found out that the accuracy of demand predictions for carry-over products is affected by the number of years the product has been on the market [18]: items in their first season as carry-over retain a high forecast error, whereas products that have been on the market for multiple years generate reliable baselines. Anticipatory logistics generates its highest return precisely in this second sub-category: products with long market histories and stable demand patterns. For these items, the investment in pre-positioning inventory before the selling season is justified by two concurrent factors. First, the forecast error is structurally low, which means the risk of misallocating pre-positioned stock is contained. Second, the consequences of a stock-out on an iconic product are disproportionately severe, as established in Chapter 1. This finding is relevant for the present framework: anticipatory logistics should be applied selectively to the stable core of the assortment, not to the entirety of the portfolio [18].

The second condition for economic viability is purely financial. Anticipatory logistics requires a meaningful investment: data infrastructure, demand analytics capabilities and logistics pre-positioning costs. In order to allow this investment to generate a positive return, the margin profile of the products being managed must be high enough to absorb those costs and still deliver a net contribution improvement compared to the reactive baseline. The relevant metric here, as introduced in the Section 1.3 is the *Gross Margin Return on Investment* (GMROI). A product with a high gross margin percentage and a stable selling velocity generates a GMROI that rewards inventory pre-commitment. Conversely, a product with thin margins even if it sells consistently, may not justify the additional logistics and analytical costs associated with anticipatory pre-positioning. The economic logic is , in a sense,

directional: anticipatory logistics should be prioritized for the highest margins items in the portfolio, where the financial upside of avoiding a stock-out or overstock episode is the largest.

This principle finds additional support in the supply chain strategy, specific to the luxury sector. Research classifying luxury companies into distinct clusters according to product complexity, selling volume and brand reputation has shown that large firms managing high complexity products (having high margins) consistently invest in tighter supply chain control and more sophisticated replenishment logics than smaller firms with lower-margin positioning [17]. Entry-level products, whose margins are thinner and whose clients are less sensitive to a service failure, depicts a lower priority. Core and high-tier products (the ones which creates a large amount of the brand gross profit) depict the primary target for anticipatory pre-positioning, precisely because of the cost of a service failure at that level is both: financially and relationally largest.

The third condition concerns geography. Anticipatory logistics does not generate equal value across all retail locations. Its economic return is higher in boutiques where demand is concentrated, predictable and strategically significant. The distribution channel has been identify as a primary driver of supply chain strategy, with a specific emphasis on the relationship between channel type and inventory allocation complexity [18]. Luxury companies typically operate through a combination of directly operated stores (DOS), mono-brand franchise stores, and multi brand independent retailers [17]. Each channel type has a distinct demand profile. Directly operated stores, particularly flagship boutiques in primary luxury markets (Milan, Paris, London, Tokyo, Dubai), generate a disproportionate share of total brand revenue and serve the highest concentration of VIC. These locations are characterize by three factors that make them ideal targets for anticipatory pre-positioning: their historical demand data is both rich and reliable; their client base generates the highest reputational cost in the event of a stock-out and their logistical accessibility makes pre-positioning operationally feasible.

The geographic logic of anticipatory pre-positioning is reinforced by the demand concentration patterns documented in the luxury sector. As established in Chapter 1, this client concentration has a direct operational implication: an unavailability event affecting a single high-volume boutique in a primary market can generate a revenue loss that is structurally wider than a stock-out affecting multiple lower-volume locations simultaneously. Anticipatory logistics, applied at the level of the strategic boutique, addresses the highest-priority locations first.

On the other hand, secondary and tertiary markets locations present a different calculus. For these locations, the investment required to implement anticipatory prepositioning may not be recoverable through the economic gains it generates. The framework recommends a tiered approach: full anticipatory capability deployed at tier-one boutiques, and a more traditional demand-driven replenishment model maintained for tier-two and tier-three locations. This geographic tiering is consistent with the alignment between supply chain strategy and market requirements [17].

The fourth condition is temporal. Even within products and locations where anticipatory logistics is broadly applicable, the timing dimension matters. Certain periods of the calendar generate structurally higher and more predictable demand concentration than others. These windows represent the moments at which anticipatory pre-positioning delivers its highest operational return per unit invested. For luxury fashion and leather goods, the most crucial demand concentration periods embed: the weeks immediately following collection presentations, when editorial coverage and client interest peak simultaneously; the pre-holiday season (typically November to December), which drives a disproportionate share of gifting purchases across all luxury categories; and local high-spend events such as fashion weeks, cultural openings and VIC private presentations that generate predictable surges in specific locations.

This temporal predictability has a practical consequence for logistic planning. When a brand is able to observe in advance the periods during which demand for a specific SKU (at a specific location) is likely to exceed its standard selling velocity, it can pre-position additional inventory before the demand spike occurs (absorbing the incremental demand without triggering emergency transfer). In the fashion industry, it has been confirmed that the incorporation of external variables (including calendar effects, promotional periods and editorial events) into forecasting models significantly reduces prediction issues compared to models based on historical data alone [19]. The opposite situation also applies. During low-demand periods (off-season months, post holiday clearance windows, inter-seasonal transitions), the case for holding elevated inventory pre-positions weakens. The carrying cost of pre-positioned stock during these periods exceeds the expected benefit from improved availability, because that demand probability is low.

The fifth and final condition connects anticipatory logistics directly to the client relationship. As established in Chapter 1, not all stock-outs carry the same costs. The financial damage of a missed sale is computable; the relational damage of failing to serve a top-tier client is more complex to quantify but, in luxury, potentially far larger. This section makes the case that anticipatory logistics has a client-level justification that is independent of, and complementary to, the product, margin, location and calendar arguments developed above.

As established in Chapter 1, the demand effect of a stock-out is proportionally far greater for a VIC than for a standard client. The operational logic of client-driven anticipatory logistics follows from this evidence. A brand with access to CRM data can identify, for any given SKU and boutique location, which specific clients have historically purchased that product or product family, and when their purchase cycles typically happen. This precious information, if combined with inventory availability data, enables the brand to pre-position stock in advance of an expected VIC visit or purchase event. The client arrives and finds the product available. No emergency transfer is required. No service failure occurs. This client level application of anticipatory logistics is technically consistent with what is

described as demand sensing using CRM data as a predictive input. The use of client purchase history, visit frequency and preference profiles (all data captured in luxury CRM systems) as inputs to inventory positioning decisions is not a theoretical extension of the model: it is a practical application that transforms existing data into logistics value. Furthermore, the client dimension introduces an asymmetry that strengthens the case for investment. The cost of pre-positioning one additional unit of high margin SKU at a strategic boutique in advance of a VIC visit by definition (the holding cost of that unit for the relevant time window). The cost of failing to do so is the product of the VIC's expected lifetime value multiplied by the probability of churn. In luxury, where client lifetime values at the top segment are structurally high, this asymmetry is pronounced. The cost of over-positioning is bounded; the cost of under positioning is not.

A rigorous framework must identify not only where a model creates value but, also, where it does not. The five conditions described above are individually necessary but not sufficient: their value depends on how they are combined. Products that meet only one or two of the five criteria should not be subjected to anticipatory pre-positioning. The operational and financial costs of doing so would exceed the expected benefits. The most transparent counter-case is the seasonal product with low margin sold in a secondary market location to a client segment of moderate spending. For this combination, the demand uncertainty is high (seasonal product), the financial upside of correct positioning is limited (low margin), the geographic stakes are moderate (secondary market), the demand peaks are less pronounced and less predictable (no primary event calendar), and the client cost unavailability is contained. Applying anticipatory logistics to this combination would generate a pre-positioning cost without a proportionate benefit. The reactive model, in this case, is the correct choice from an economical perspective.

An additional limitation concerns the role of forecast accuracy. Section 2.1 has established that AI-based forecasting reduces forecast error but does not eliminate it. In the context of the present framework, this means that even within the five favorable conditions, a probability of forecast error must be priced into the investment decision. For products with very low demand volumes (limited editions, ultra-exclusive items), the small number of observation available for calibrating the forecast model limits the reliability of the prediction and therefore, the confidence with which inventory can be pre-positioned. The appropriate operational response is not to abandon anticipation for these products, but to maintain a smaller and more conservative pre-position quantities (supplement them with faster reactive corrections when the demand signal firms up).

The five conditions outlined in this section are not independent switches. They function as a combined diagnostic. A product that satisfies all five (iconic status, high margin, flagship boutique, identified demand peak, VIC client base) represents the maximum-priority target for anticipatory logistics deployment. A product that satisfies fewer than three conditions should, in most cases, remain within the reactive model until the economic case strengthens.

This graduated logic is consistent with the general principle established in the supply chain strategy literature:

<<there is no single supply chain approach that fits all product types within a complex luxury portfolio>> [17]

In the context of the KPI diagnostic framework introduced in Section 1.3, the five conditions described in this section can be mapped to the five performance indicators. A product-location combination with a high stock-out rate (signaling chronic unavailability), a declining GMROI (signaling underperformance relative to inventory investment), a growing logistics cost per unit (signaling reactive corrections), and a compressing net margin after logistics (signaling that the cost of reactivity is eroding financial performance) signals that a supply chain problem exists at the combination. It does not, by itself, prove that anticipatory pre-positioning is the correct solution: the same pattern of KPI's can also result from an incorrect assortment, inadequate forecasting, production delays or pricing policy. The KPIs identify where the five-condition framework (above) should be applied first. They are a signal to investigate, not a substitute for the eligibility check.

Table 6: ***Priority tiers by number of eligibility criteria met***

CRITERIA MET (OUT OF FIVE)	PRIORITY TIER	IMPLICATION
FIVE OUT OF FIVE	Maximum priority	Full anticipatory pre-positioning deserves formal economic evaluation
FEWER THAN THREE OUT OF FIVE	Low priority	The reactive model remains the correct choice, in most cases

Source: own elaboration

The next section of this thesis will translate this diagnostic logic into a formal decision-making framework for stock pre-positioning. Section 2.3 will compare the expected economic benefits of anticipatory logistics (recovered sales, lower emergency logistics costs, higher GMROI) against the costs of pre-positioning (holding costs, forecast error risk, infrastructure investment) and propose a structured decision calculus for determining when anticipatory logistics is economically superior to the reactive alternative.

2.3 A DECISION-MAKING FRAMEWORK FOR STOCK PRE-POSITIONING

Section 2.2 has identified where anticipatory logistics creates value: it has isolated five conditions (product type, margin profile, location, calendar window and client segment) that, when simultaneously satisifed, define the product-location-client combinations for which pre-positioning is a candidate strategy. That framework answers the question of eligibility. However, it does not answer the operational question that a logistics manager must resolve in practice: given an eligible combination, how much inventory should be pre-positioned? And at what point does the decision to pre-position stop to be economically convenient? As already said, the five conditions are necessary but not sufficient. A product may satisfy all of them and still not justify a specific pre-positioning quantity, because the quantity itself carries a cost. This section translate the qualitative diagnostic into a quantitative decision calculus. The objective is to provide a structured rule that compares the expected benefit of pre-positioning one additional unit, against its expected cost and identifies the quantity at which the two are balanced.

The starting point is to recognize that stock pre-positioning is, in its formal structure, a single-period inventory decision under uncertainty. The brand commits a quantity of stock to a location before the demand for the relevant window is known. Once the window opens, demand is realized and the commited quantity is revealed to have been either insufficient (a stock-out occurs) or excessive (unsold units remain). This is the classical structure of what the operations literature calls the newsvendor problem, and its mathematical foundation was first formalized by Arrow, Harris and Marschak in their static model with uncertainty [15]. In that model, the organization must chose the stock level S that a commodity should have at the beginning of a given period in order to provide for a demand that is a random variable with a known cumulative distribution [15]. The parallel with anticipatory pre-positioning is exact: the pre-position quantity is the stock level S, the selling window is the period and the demand for the specific SKU (at the specific location) is the random variable.

The economic tension in this decision is the balance between two opposing error costs. If the pre-positioned quantity is too low, the brand incurs an understock cost: the margin forgone on the sale (that they cannot fulfil), plus, the relational damage (related to luxury) documented in Chapter 1. If the quantity is too high, the brand incurs an overstock cost: the holding cost of carrying the unsold units through the window and the markdown loss required to clear them afterwards. The first scenario described can be captured by what has been named as the '*depletion penalty'*, defined as the amount the organization would be willing to pay to satisfy demand in excess of available stock [15]. The optimal stock level is the one at which the marginal cost of one additional unit is exactly offset by its marginal expected benefit. Denoting the per-unit understock cost as the underage cost and the per-unit overstock cost as the overage cost, the optimal quantity satisfies a single condition expressed though the critical ratio [10]:

$$\text{Critical Ratio} = \frac{\text{Underage Cost}}{\text{Underage Cost} + \text{Overage Cost}}$$

The critical ratio is a number between 0 and 1. It expresses the proportion of demand outcomes that should be covered by the pre-positioned stock. The optimality condition established by the model states that the pre-position quantity should be set at the level whose in-stock probability equals the critical ratio. Formally, if F is the cumulative distribution of demand, the optimal quantity S* is equal to: F(S*) = Critical Ratio [15]. The interpretation is direct and operational: the higher the underage cost relative to the overage cost, the higher the critical ratio, and the more generously the brand should pre-position. On the other hand, when the overage cost dominates, the critical ratio falls and the brand should pre-position conservatively. The entire decision therefore collapses into the estimation of two quantities (the underage cost and the overage cost) and the shape of the demand distribution, which the anticipatory infrastructure described in Section 2.1 is designed to estimate.

The next step is to populate the two terms with the specific economics of luxury retail, using the cost structure established in Chapter 1. The underage cost should not be perceived as the lost sale. It can be expressed as: $C_u = m + c_e + c_r$, where m is the contribution margin forgone, c_e is the incremental cost of the urgent correction (that the reactive alternative would require) and c_r is the expected cost of the compromised client relationship. For an occasional client this relational component is modest: for a Very Important Client, the relational component can exceed the immediate margin by an order of magnitude. The underage cost in luxury is therefore strongly client-dependent, and it is systemically wider for the high-value combinations that Section 2.2 identified as eligible. This raises an operational question that the Chapter 3 simulation must answer explicitly: how is c_r actually calculated here? As a fixed multiple of margin, as a function of client tier or through another parameterization? Without that operational definition, c_r cannot simply be added to the underage cost as an assumption.

On the other hand, the overage cost is bounded. It comprises the holding cost of the pre-positioned unit for the duration of the window and the expected markdown loss if the unit does not sell. Formally, $C_o = h + d$, where h is the holding cost of the unit for the duration of the window and d is the expected markdown loss if it remains unsold. Chapter 1 established that luxury logistics costs should sit between 8% and 10% of the net revenue [9], and that a markdown both erodes the margin on the discounted unit and damages price integrity. Even accounting for both effects, the overage cost per unit remains a fraction of the contribution margin, because the unit retains value and it can be transferred or in a later window rather than destroyed. The underage cost is large and, for iconic products served to top clients, effectively unbounded. Instead, the overage cost is small and bounded. This asymmetry pushes the critical ratio toward the upper end of its range.

It is possible to give this asymmetry a concrete magnitude by drawing on the empirical evidence from the apparel industry, which shares with luxury the structure of a concentrated selling season and long production lead times. In the accurate response study of a fashion firm (operating under a quick response

system), the underproduction (underage) cost was measured to be two to three times higher than the overproduction (overage) cost [21]. Applied to the critical ratio, an underage to overage ratio of two to three implies a critical ratio between 0.67 and 0.75. This is the so called floor for luxury: it corresponds to products where the stock-out carries only the lost margin and no significant relational damage. For the iconic (high margin products served to VIC clients) the relational factor raises the underage cost far above this floor. The service level target of 95% recommended for high value categories and adopted in Section 1.3 [10] corresponds, through the same critical-ratio logic, to an implied underage to overage ratio of roughly nineteen to one. A 95% target service probability is equivalent, within the newsvendor formulation, to assuming an underage cost nineteen times the overage cost. The 95% target is therefore the modelling assumption that drives the nineteen-to-one ratio, not an independent empirical finding about luxury cost structures.

The high relational penalty that luxury attaches to a stock-out is precisely what produces the low stock-out target that Chapter 1 recommended on qualitative grounds. The framework of this section closes the loop between the qualitative diagnostic of Chapter 1 and the quantitative optimization of the pre-position quantity. The critical ratio determines the target in-stock probability. However, the quantity that achieves that probability depends on the demand distribution. It is here that anticipatory logistics changes the economics rather than merely applying the classical rule. The purpose of the data and analytical infrastructure described in Section 2.1 is to reduce the uncertainty of the demand distribution: to make F tighter around its mean. The effect of a tighter distribution is that the same critical-ratio target can be reached with a smaller safety buffer, which lowers the overage cost of achieving a given service level. The empirical impact of this effect is documented. It has been observed that only the first 20% of a season's demand was found to improve forecast accuracy dramatically, because early and total demand were correlated at approximately 0.8 to 0.9 for most products, and the first 92% of orders exhibited a correlation greater than 0.99 with total season demand [21]. The demand signal that anticipatory logistics exploits, in quantitative terms, this early-demand information: it converts a wide pre-season distribution into a narrow in-season one and it is the narrowing that makes generous pre-positioning economically safe.

The economic value of this narrowing can be read directly from the same study which measured the total cost of the two error types. First, under no response (a reactive commitment made before any demand signal) the combined overproduction and underproduction cost reached 10.38% of sales, split into 6.80% of overproduction cost and 3.58% of underproduction cost [21]. Second, under the optimized responce model the same combined cost fell to 1.48% of sales with overproduction cost reduced to 1.3% and underproduction cost to 0.18% [21]. Both error types fell simultaneously: anticipation did not trade stock-outs for overstock, it reduced both. Given that profits in that sector averaged about 3% of sales, the reduction in mismatch cost was large enough to increase profits between 60% and 127%[21]. These figures are drawn from apparel rather luxury, and the absolute percentages will differ for luxury's higher

margins and lower volumes; but the direction and the mechanism transfer directly, because both sector face the same structural problem of committing inventory before a concentrated selling window.

The decision rule that the framework prescribes can now be stated in operational form. For a given SKU at a given location and window, the brand should pre-position the quantity at which the expected benefit of the marginal unit equals its expected cost. The expected benefit of the marginal pre-positioned unit is the underage cost multiplied by the probability that demand reaches or exceeds that unit. In economic terms, it is the recovered contribution margin, plus the avoided emergency logistics premium, plus the avoided relational damage. The expected cost of the marginal unit is the overage cost multiplied by the probability that demand falls short of that unit. It is the holding cost over the window plus the expected markdown loss. The optimal quantity is reached when these two expectations are equal, which is the same condition as F(S*) = Critical Ratio. In practice the calculation is performed unit by unit: the brand pre-positions additional units as long as the recovered value of an avoided stock-out exceeds the carrying and markdown cost of a potentially unsold unit and stops when the two converge.

A concrete application clarifies the mechanics. Consider an iconic leather good pre-positioned at a flagship boutique ahead of an identified peak window, using the cost parameters established in Section 1.3. The unit carries a gross contribution of €200 after standard logistics of €15. If the reactive alternative would require an emergency transfer costing an additional €40 per unit to correct a stock-out, then serving the demand reactively already erodes €40 of contribution, and failing to serve it at all forgoes the full €200 plus the relational damage. The underage cost of the pre-positioned unit is therefore at least €240 (the $m+ c_e$ portion of the decomposition above) and, once the VIC relational component is included, materially higher. The overage cost is the holding cost for the window plus the expected markdown loss on an unsold unit, which together remain a small fraction of €200 because the unit can be redeployed to a later window rather than cleared. With an underage cost of this magnitude against a modest overage cost, the critical ratio for this combination sits well above 0.90 and the framework recommends pre-positioning the full quantity required to cover demand up to that high percentile. The same calculation applied to an entry-level product with a €30 margin, no emergency premium and a client segment that substitutes readily would yield a critical ratio near the 0.67 floor. The framework would recommend little or no pre-positioning, confirming the selectivity established in Section 2.2 on independent quantitative grounds.

In addition, the framework clarifies the nature of the inventory being committed. Based on the definition, a resource can be of two types. The first ones, are the dedicated resources which are used to protect only against disruption. The second ones, are the dual-purpose resources. They are used to protect against uncertainty while also improving efficiency in normal times [22]. Pre-positioned luxury inventory is a dual-purpose resource: it raises availability for the anticipated demand of a specific window (an efficiency benefit) while simultaneously buffering against an unexpected demand surge (a resilience

benefit). This dual character is what justifies its cost in a sector where dedicated buffers held purely against rare events would be difficult to defend. Furthermore, it introduces a crucial qualification specific to luxury. The backlog cost (the cost of making the customer waiting) in luxury can be low. Clients frequently return for an iconic product even after a delay or a price increase, so that for the most desirable products a brand may rationally choose to post pone a sale rather than hold additional inventory [22]. Where this holds, the underage cost is partially recoverable, the effective critical ratio falls and the framework suggest a more measured pre-position quantity. The decision is not mechanical: the two cost terms must be estimated with the specific product and client view. The framework does not automatically recommend more inventory. It recommends a higher pre-positioned quantity only when the underage cost is higher than the overage cost.

A more practical dimension concerns where in the network the pre-positioned stock is held. Committing inventory downstream, close the boutique, maximizes availability but concentrates the overage risk at a single location. Holding it further upstream preserves the flexibility to redirect it but reintroduces the lead-time gap that anticipatory logistics is designed to close. This can be described as the trade off between the pooling benefit of centralized stock and the diversification benefit of decentralized stock and it finds that when the priority is to guarantee availability at the point of demand, holding inventory downstream is frequently optimal [22]. For luxury where the cost of unavailable iconic product in front of a VIC is precisely the large underage cost identified above, the framework leans toward downstream pre-positioning at the eligible tier-ones boutiques, with a smaller upstream reserve retained to serve less predictable locations through faster reactive correction. Table 7 summarizes this trade-off across the three configurations the framework considers.

Table 7: ***Network configuration for pre-positioned stock***

CONFIGURATION	POOLING FLEXIBILITY	LEAD TIME	KEY RISK
CENTRAL STOCK	High	Longer	Availability gap during transfer lead time
BOUTIQUE STOCK	Low	Short	Local overstock
HYBRID	Medium	Short for base stock; longer for the reserve	Balanced between the two

Source: own elaboration

Figure 1: ***Network configurations for pre-positioned stock***

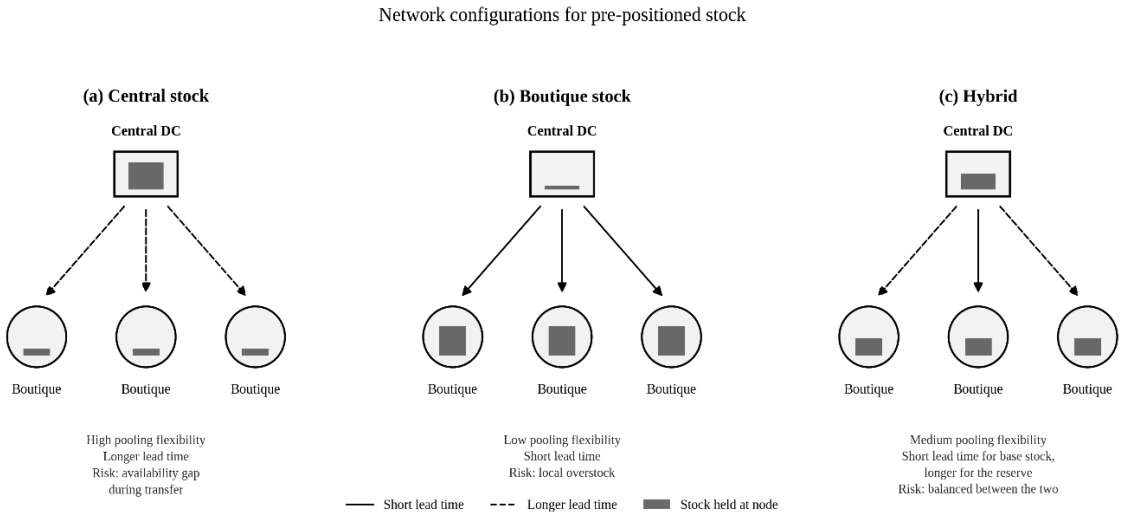

Source: own elaboration

This is a direct application of the concept of the anticipatory order, understood as the practice of moving part of an order to an earlier period in anticipation of a shortfall, whose optimal size increases as demand uncertainty increases [22].

This last point can appear to sit in tension with the earlier finding that a tighter forecast reaches the same service level with a smaller buffer. The two statements are compatible once four related quantities are kept distinct: the total stock required to protect a given service level against demand uncertainty; the share of that stock that is pre-positioned downstream in anticipation of the window rather than held centrally; the safety stock retained specifically as a buffer against residual forecast error; and the position of that stock within the network. A tighter forecast reduces the safety-stock component needed to reach a given service level — the effect described above. The anticipatory order specifically concerns the share of stock moved downstream ahead of the window, a quantity governed by the network trade-off just described rather than by forecast tightness alone

The final component of the framework is an explicit account for the boundary conditions under which the decision reserves, since these are precisely the parameters that determine whether pre-positioning remains convenient. Three parameters govern the outcome. First, forecast accuracy: as the forecast error rises, the demand distribution widens, the quantity required to reach the critical ratio percentile increases and the overage cost of achieving a given service level rises. Second, the markdown rate: a higher markdown loss raises the overage cost, lowers the critical ratio and reduces the optimal pre-position quantity. Third, the emergency shipping cost: a higher premium for reactive correction raises the underage cost, raises the critical ratio and increases the optimal pre-position quantity, because avoiding the stock-out becomes more valuable relative to carrying the buffer. These three parameters are the levers whose interaction determines when anticipatory logistics is economically superior to the reactive

alternative. Their sensitivity is summarized in the table below and will be examined empirically in Chapter 3.

Table 8: ***Effect of the key parameters on the pre-positioning decision***

PARAMETER	DIRECTION OF CHANGE	EFFECT ON OPTIMAL PRE-POSITION QUANTITY
FORECAST ACCURACY	Higher accuracy (tighter demand distribution)	Same service level reached with a smaller buffer; pre-positioning becomes economically safer
MARKDOWN RATE	Higher markdown loss	Overage cost rises, critical ratio falls, optimal quantity decreases
EMERGENCY SHIPPING COST	Higher reactive-correction premium	Underage cost rises, critical ratio rises, optimal quantity increases
UNIT CONTRIBUTION MARGIN	Higher margin	Underage cost rises, critical ratio rises, pre-positioning prioritized
CLIENT VALUE (VIC EXPOSURE)	Higher relational cost of a stock-out	Underage cost rises sharply, critical ratio approaches the service target, generous pre-positioning justified

Source: own elaboration

The framework which has been proposed in this section, converts the pre-positioning decision from a matter of judgement into a structured calculation. It reduces the choice to three inputs (the underage cost, the overage cost and the demand distribution) and a single output (the critical ratio target and the quantity that achieves it). It makes explicit why luxury, with its large and client dependent underage cost and its bounded overage cost, systematically favors generous pre-positioning for the eligible combinations identified in Section 2.2, and why the same calculus recommends against pre-positioning for thin-margin products in low-stakes locations. Above all, it identifies the parameters (forecast accuracy, markdown rate and emergency shipping cost) whose values determine whether the anticipatory model is economically superior to the reactive one. Chapter 3 operationalizes this framework in a Python simulation. It assigns explicit values to each of these parameters. Compares the traditional reactive model against the anticipatory model across the resulting scenarios and tests (through a sensitivity analysis) on exactly these three parameters, the conditions under which anticipatory logistics delivers a measurable economic advantage in luxury retail.

Chapter 3

3 PYTHON SIMULATION: ECONOMIC EVALUATION

Chapter 2 has proposed a solution and a rule for applying it. What the framework does not establish is whether following it is worth its cost. A rule can be internally coherent and still describe a mechanism of negligible economic effect, or one whose advantage disappears as soon as its operating condition change. This chapter subjects the framework to that test. Section 3.1 describes the design of the simulation built for the purpose and the assumptions on which it rests. Section 3.2 presents the input variables and the calculation logic through which the indicators defined in Section 1.3 are derived. Section 3.3 compares the results, examines their sensitivity to the parameters that govern them and identifies the conditions under which anticipatory logistics is economically convenient. One distinction runs through the whole chapter and is worth stating at the outset, since it is easily conflated. Where inventory is most worth holding and where predictive capability is most worth buying are two different questions and they do not have the same answer. Stock-positioning priority, which Section 2.2 establishes is not the same thing as the value of predictive information, which this chapter measures. The simulation separates the two and the separation proves to be among its more useful results.

3.1 SIMULATION DESIGN: TRADITIONAL VS ANTICIPATORY

Section 2.3 concluded with a decision rule. It reduced the pre-positioning problem to three inputs (the underage cost, the overage cost and the demand distribution) and to a single output, the target in-stock probability and the quantity that achieves it. That rule prescribes what a brand should do, but it does not establish what doing it is worth.

It is necessary to state at the outset what kind of evidence this exercise can and cannot produce. The dataset on which the simulation operates is synthetic. It has been designed to reproduce the economic structure described in Chapters 1 and 2 rather than to estimate the magnitudes prevailing in the luxury sector. The cost parameters are drawn from the worked examples of Section 1.3, the behavioral parameters are anchored to the literature where empirical evidence exits and are declared as assumptions where it does not. The results therefore demonstrate a mechanism and its sensitivity to the conditions under which operates. They do not constitute an empirical estimate of the return that a specific brand would obtain from adopting anticipatory logistics, and they should not be read as one. This limitation is not incidental to the method: no publicly available dataset reports, at the level of the individual product-location-window combination, both the demand realization and the full cost of the mismatch. A simulation is the appropriate instrument precisely because of the quantity of interest is unobservable in the available data.

The unit of analysis is the combination of a single Stock Keeping Unit (SKU), a single retail location and a single selling window. This is the same unit on which section 2.2 defined eligibility and Section 2.3 defined the optimal quantity. Preserving it across the three chapters ensures that the simulation tests the framework that was actually proposed. The selling window is the period within which the committed quantity must satisfy demand, and the commitment is made before the demand for that window is realized. This is the single-period structure of the newsvendor formulation discussed in Section 2.3 [15]. Each replication of the simulation represents one such window: demand within it is a random variable whose distribution is specified below.

The simulation compares three stocking policies rather than two. A comparison between anticipatory logistic and a reactive baseline is only informative if the baseline is specified fairly. A reactive model that is deliberately under-specified would guarantee the conclusion in advance, and the resulting advantage would measure the weakness of the comparison rather than the value of the anticipation. The reactive model is therefore represented twice: once as it plausibly behaves in practice, and once at its own best. The first policy, denoted the base-stock reactive model, does not observe the demand concentration window at all. It sizes inventory on routine demand at an ordinary service level and corrects any resulting shortfall after the fact through the emergency transfers described in Section 1.2. It depicts the reactive model as Chapter 1 characterized it. The second policy, the optimized reactive model, does observe the window and applies the critical ratio derived in Section 2.3, but it commits the quantity before any demand signal is available and therefore carries the full pre-season uncertainty. The third policy, the anticipatory model, observes an early demand signal and commits afterwards. Its quantity is consequently both better centered on the realized demand and less heavily padded against uncertainty. Table 9 summarizes the three configurations.

Table 9: ***The three stocking policies compared in the simulation***

POLICY	INFORMATION AVAILABLE AT COMMITMENT	QUANTITY RULE	ROLE IN THE COMPARISON
REACTIVE (BASE STOCK)	Routine demand history only; the peak window is not identified	Routine service level applied to ordinary demand	The reactive model as described in Chapter 1
REACTIVE (OPTIMISED)	The peak window is identified; no demand signal	Critical ratio applied to the pre-season distribution	The strongest fair baseline

ANTICIPATORY	The peak window and the early demand signal	Critical ratio applied to the updated distribution	The model proposed in Chapter 2

Source: own elaboration

Figure 2: ***Commitment timing and information set of the three policies***

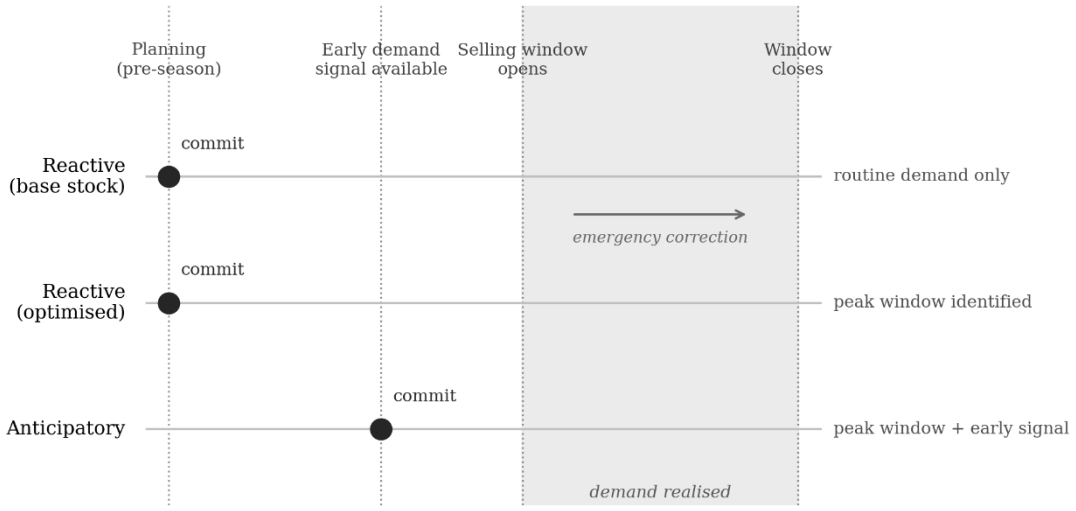

Source: own elaboration

Figure 2 places the same distinction on the time axis. Both reactive policies commit before any demand signal exists and differ only in the information brought to that decision; only the base-stock policy relies on emergency correction during the window as a systematic part of its operation. The anticipatory policy defers the commitment until the signal is available, which is the single structural difference the simulation isolates.

That difference is informational rather than quantitative. Demand and the early signal are generated jointly, so that the signal carries exactly the correlation with total window demand reported in the accurate response literature [21]: demand is expressed as the sum of the component explained by the signal and an independent residual, in the proportions that correlation determines. Under a linear-normal updating rule the residual standard deviation falls by a factor equal to the square root of one minus the square correlation, which is the quantitative content of the claim advanced in Section 2.3 that anticipation converts a wide pre-season distribution into a narrow in-season one. It should be declared that the cited study reports the correlation and not the joint distribution: the narrowing implemented here is a modelled consequence of a measured correlation under an assumption of joint normality.

$$D = \mu + \sigma \left(\rho \, z_s + \sqrt{(1 - \rho^2)} \, z_\varepsilon \right), \qquad S = z_s$$

Here D is total demand over the window, μ its mean and σ its standard deviation, ρ is the correlation reported in Ref. [21], and z_s and z_ε are independent standard normal variates. The early signal S is the

standardized component of demand observed before the commitment is made, so that demand and signal are jointly normal with correlation ρ by construction. Conditional on having observed S, demand is normally distributed with mean $\mu + \sigma\rho S$ and standard deviation $\sigma\sqrt{(1-\rho^2)}$. The anticipatory policy sets its quantity from that conditional distribution; the reactive policies set theirs from the unconditional one. The residual standard deviation therefore falls by the factor $\sqrt{(1-\rho^2)}$ stated above, and at $\rho = 0$ the two coincide, so that signal carrying no information leaves the anticipatory policy identical to the optimized reactive one.

Figure 3: ***Effect of the early demand signal on the demand distribution and on the committed quantity***

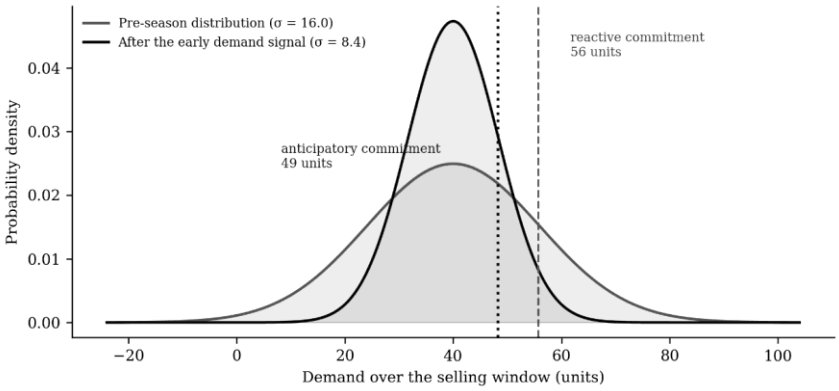

Source: own elaboration

Figure 3 applies the same target service level to both distributions. The area to the left of each vertical line is the same, since both policies are pointing to the same in-stock probability. What differs is the quantity required to reach it, and the gap between the lines is the inventory the reactive policy carries purely to compensate for the information it lacks. The narrowing of the buffer, however, is only one of two effects. The anticipatory quantity also moves with the signal instead of remaining fixed on the unconditional mean, and this repositioning of the central estimate is the larger of the two. Anticipation does not operate by committing more inventory: it operates by committing inventory that is located closer to the demand that will actually materialize.

The implementation required one correction to the cost structure proposed in Section 2.3, which must be reconciled in that section. The underage cost was defined there additively, as the sum of the contribution margin forgone, the incremental cost of the urgent correction and the expected cost of the compromised client relationship. These three outcomes, however, are not simultaneous. If the shortfall is corrected by an emergency transfer, the margin is earned and no relational damage occurs, and only the premium is paid. If it is not corrected, the margin and the client value are lost, but no premium is incurred. The additive expression therefore describes the case in which the emergency correction is attempted and fails, which is an upper bound rather than an expected cost.

The simulation adopts the expected-value formulation in which the underage cost is the probability-weighted average of the two outcomes, the weight being the share of shortfall recoverable through emergency transfer. At a recovery share of zero the expression reduces to the pure lost-sale case, and at a share of one it reduces to the emergency premium alone. Applied to the iconic article at the flagship boutique, the correction lowers the underage cost from 641.71 euro to 208.51 and the critical ratio from 0.9527 to 0.8674. The direction of the correction matters more than its magnitude: the additive formulation systemically overstates the case for pre-positioning, and adopting it would have inflated the advantage attributed to anticipation throughout the chapter.

The relational component of the underage cost requires an operational definition, since Section 2.3 established its existence without specifying its calculation and left the question explicitly open. It is modelled here as the product of three quantities: the proportional erosion of the client's future purchases following a stock-out, the share of that erosion which is not recovered when the client returns for the product at a later date, and the contribution margin the affected client is expected to generate over a one-year horizon. The first quantity is empirical, and corresponds to the reduction in future purchases documented in Chapter 1 [8]. The second reflects the observation that the blacklog cost in luxury can be low, because clients frequently return for an iconic product even after a delay [22]; its value is an assumption, and its influence is examined in the sensitivity analysis of Section 3.3. The third differentiates the client segments identified in Section 2.2.

The dispersion of that third quantity across client segments is not assumed. If the top segment comprises two per cent of clients and generates forty-five per cent of purchases [2], the implied purchase intensity per client is 22.5 for that segment and 0.561 for the remainder. As a consequence, a VIC is worth approximately forty times an occasional client. Only the average client value remains an assumption: the ratio between the segments is derived from the reported concentration. It must be stated that the cited source measures purchases rather than contribution margin, so that applying the same multiplier to margin assumes comparable margin rates across the two groups, an assumption which understates the top segment if its purchases are concentrated in higher-margin categories. This is the operational answer to the question left open in Section 2.3: the relational cost is neither a fixed multiple of margin, nor an unexplained assumption, but a function of three declared parameters, two of which are traceable to a source.

Demand over the window is modelled as normally distributed, consistently with the newsvendor formulation adopted in Section 2.3 [10]. The normal distribution admits negative values, which have no interpretation in this context and the model reports the probability mass falling below zero at every configuration. For the iconic article that mass is negligible at four hundredths of one per cent. For the seasonal article, whose demand is deliberately parameterized as far more variable, it reaches approximately three and a half percent. Negative drawns are truncated at zero, which introduces a small

upward bias in realized mean demand for that product. The bias applies identically to the three policies and therefore does not distort the comparison between them, but it means that the absolute figures reported for the seasonal article should be read as slightly conservative with respect to the mismatch costs incurred. This objection is not left open. Block 6 of the simulation evaluates all eight combinations a second time with demand drawn from a lognormal distribution, which cannot take negative values, calibrated to the same mean and the same standard deviation, with the order-up to level taken as the critical ratio quantile of that distribution. Every cost parameter is read from the same functions, so that the two runs are priced identically. The ranking of the eight combination is unchanged and the advantage of anticipation is larger under the lognormal in seven of the eight cases, rising om the seasonal article at the flagship boutique from 92.1 to 125.9 per cent. The reason is structural: a longer right tail makes a badly sized commitment more expensive so that information narrowing the distribution before the commitment is worth more. The figures reported throughout this chapter are therefore the conservative ones and the normal distribution is retained because it is the standard assumption of the newsvendor literature and keeps the optimal quantity in closed form not because the conclusion depends upon it.

The simulated context follows from the five conditions identified in Section 2.2. Two product categories carry both the product type and the margin profile; two retail locations carry the location and, through the client mix attached to each, the client segment; two window types carry the calendar condition. The resulting eight combinations exercise all five conditions while remaining few enough for every parameter to be individually inspectable. Table 10 reports the configuration of the two product categories.

Table 10: ***Configuration of the two simulated product categories***

PARAMETER	ICONIC CARRY-OVER	SEASONAL ENTRY-LEVEL
UNIT GROSS MARGIN	€200 (70% of retail price)	€80 (60% of retail price)
MARKDOWN DEPTH IF UNSOLD	10% — the article carries over	45% — the article must be cleared
RECOVERY FACTOR	0.70 — the client returns for it	0.15 — the client moves on
COEFFICIENT OF VARIATION OF DEMAND	0.30 — long sales history	0.55 — no sales history
ROUTINE DEMAND, FLAGSHIP	31 units	46 units
ROUTINE DEMAND, SECONDARY MARKET	9 units	19 units
TOP-TIER CLIENT SHARE, FLAGSHIP	55%	20%

TOP-TIER CLIENT SHARE, SECONDARY MARKET	15%	5%

Source: own elaboration

The two locations differ on routine demand volume, on the concentration of top-tier clients and on the cost of an emergency transfer, which is set at 40 euro for the flagship and 70 euro for the secondary market, on the ground that a location less well connected to the distribution network is more expensive to correct after the fact. The calendar dimension follows the characterization advanced in Section 2.2, where an identified peak is described as a period whose demand is both structurally higher and more predictable than an ordinary one. Both properties are modelled: the peak raises expected demand thirty per cent above the routine level and carries an early demand signal correlated at 0.85 with total window demand, while the ordinary window leaves demand at the routine level and carries a signal correlated at 0.55, since a period not identified in advance offers no comparison base against which early sales can be read. The correlation reported in the accurate response literature refers to identified selling seasons. The value adopted for the ordinary window is an assumption and is declared as such.

The distinction also governs the base-stock reactive policy, which commits on routine demand in both cases because by definition it does not identify the peak. In an ordinary window that policy is therefore correctly sized, and the reactive baseline stands at its strongest. This is deliberate: the design gives the reactive model the conditions most favorable to it in precisely those combinations where anticipation is expected to add least, so that any advantage measured there cannot be attributed to an unfair comparison.

The composition is treated as a property of the combination rather than of the location alone. Top-tier clients concentrate their purchases on iconic articles, so that an entry-level seasonal article sold at a secondary boutique serves a different clientele from an iconic article sold at a flagship. Attributing the same client mix to every article held at a given location would overstate the relational cost of a stock-out on the entry-level article and would mechanically inflate the value of anticipation where it is least warranted. Stock is allocated independently for each combination: the model does not simulate transfers between locations within the window, a simplification to be noted, since network operator could in principle rebalance across boutiques rather than only between the central warehouse and a boutique. For the same reason the model does not represent the network configurations compared in Table 7, which remain outside the scope of the simulation.

All values in Table 10 are assumptions. They have been chosen for internal consistency with the characterization advanced in Section 2.2 rather than drawn from a source, and their purpose is to reproduce the qualitative distinctions that section establishes, so that the simulation can test whether those distinctions carry the economic consequences attributed to them. Should the parameters fail to reproduce those distinctions, the failure would be informative in itself.

The design finally addresses reproducibility. Demand realizations are generated pseudo-randomly from a fixed seed, and each configuration is evaluated over fifty thousand replications, a number sufficient for the reported averages to be stable to the second decimal place. Because the same seed generates the demand realizations for all three policies within a configuration, the comparison is paired: each policy faces exactly the same sequence of demands, which removes sampling noise from the difference between them. The interactive interface described in section 3.2 is an interface onto the same calculation engine rather than a second model, so that any value reported in this chapter can be recovered by setting its controls to the parameter values declared here. Table 11 collects the full parameter set. Parameters common to every combination are given a single value; those distinguish one combination from another are marked as varying an their values appear in Table 10.

Table 11: ***Simulation parameters, values and provenance***

PARAMETER	VALUE	STATUS	SOURCE
UNIT GROSS MARGIN	Varies by product	Worked example	Table 10; Section 1.3
GROSS MARGIN AS SHARE OF PRICE	Varies by product	Worked example	Table 10; Section 1.3
STANDARD LOGISTICS COST PER UNIT	€15	Worked example	Section 1.3
EMERGENCY TRANSFER PREMIUM	€40 flagship, €70 secondary	Worked example	Section 1.3
ANNUAL INVENTORY HOLDING RATE	25% of unit cost	Assumption	—
SELLING WINDOW LENGTH	8 weeks	Assumption	—
MARKDOWN DEPTH	Varies by product	Assumption	Table 10
COEFFICIENT OF VARIATION OF DEMAND	Varies by product	Assumption	Table 10
DEMAND EROSION AFTER STOCK-OUT	22%	Empirical	[8]
RECOVERY FACTOR	Varies by product	Assumption	Table 10; mechanism from [22]
AVERAGE CLIENT ANNUAL MARGIN	€500	Assumption	—
VIC SHARE OF CLIENTS AND PURCHASES	2% and 45%	Empirical	[2]
EARLY-SIGNAL CORRELATION	0.85 peak, 0.55 ordinary	Empirical and assumption	[21]; ordinary assumed

PEAK UPLIFT OVER ROUTINE DEMAND	30% peak, 0% ordinary	Assumption	—
SHARE OF SHORTFALL MET BY EMERGENCY TRANSFER	0.70	Assumption	—
ROUTINE SERVICE LEVEL	90%	Assumption	—
MONTE CARLO REPLICATIONS	50,000, fixed seed	Method	—

Source: own elaboration

Ten of the seventeen parameters in Table 11 are assumptions. This is a substantial share, and it would be a serious weakness if the conclusions of the chapter depended on all of them equally. Section 3.3 establishes that they do not: a sensitivity analysis conducted parameter by parameter shows that the result is governed by three of them, two of which are traceable to a source or to the product characterization of Section 2.2, while the remainder move the outcome comparatively little. The purpose of declaring the full set here is to make that demonstration possible rather than to defend each value individually.

3.2 SIMULATION INPUTS AND ECONOMIC CALCULATION FRAMEWORK

Section 3.1 has established what the simulation compares and under which assumptions. This section sets out how a single replication is computed, how the two cost components are assembled from the input variables and how the five indicators introduced in Section 1.3 are derived from the simulated outcomes. The sequence is deterministic given the parameters and the demand realization: the only stochastic element is the demand itself and the early signal correlated with it.

Figure 4: ***Calculation sequence for a single replication***

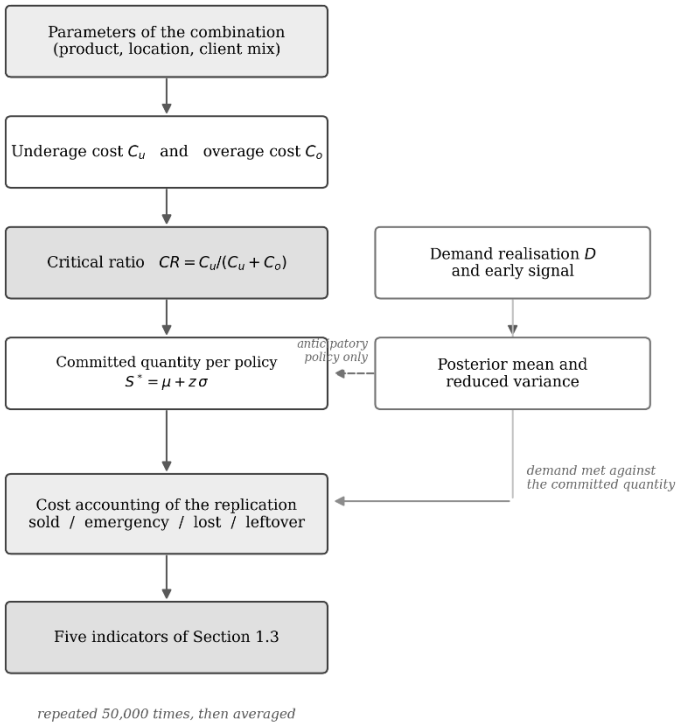

Source: own elaboration

The calculation is articulated in six stages, which are depicted in Figure 4. The parameters of the combination determine the two cost components; their ratio fixes the target in-stock probability; that probability, applied to the demand distribution, yields the committed quantity, which differs across the three policies. Demand is then realized and confronted with the committed quantity, the resulting mismatch is priced and the indicators are computed. The early demand signal enters at one point only, modifying the distribution to which the anticipatory policy applies its target; every other element of the calculation is common to the three policies, which ensures that the measured difference is attributable to the information available and not to a difference in accounting treatment.

The underage cost assembles three input variables. The unit contribution margin is the gross margin net of the standard logistics cost, and depicts what is forgone when a unit of demand is not served at all. The

emergency premium is the incremental cost of correcting the shortfall reactively and applies only to the share of the shortfall that the emergency channel actually recovers. The relational cost, defined in Section 3.1, applies to the residual share that is not recovered. The overage cost assembles two: the holding cost over the selling window, calculated as the annual holding rate applied to the unit cost and prorated to the length of the window, and the markdown loss, computed as the markdown depth applied to the retail price. Neither component is charged to units that sell.

The quantity these components produce is used throughout this chapter and has so far been described in words alone. It is stated here explicitly. For a given policy and a given demand realization, the expected economic contribution is:

$$EC = m\,Q_{\text{served}} - c_{\text{e}}\,Q_{\text{emergency}} - c_{\text{o}}\,Q_{\text{leftover}} - c_{\text{r}}\,Q_{\text{lost}}$$

Where m is the unit gross margin net of the standard logistics cost; Q_{served} is the number of units actually delivered to clients, whether from the quantity committed to the boutique or through an emergency transfer; c_{e} is the emergency premium and $Q_{\text{emergency}}$ the units supplied through that channel; c_{o} is the overage cost of holding and clearance and Q_{leftover} the units unsold at the close of the window; and c_{r} is the relational cost applied to Q_{lost}, the share of the shortfall the emergency channel does not recover. This is the expression the simulation evaluates, in a single function called unchanged by all three policies.

Three points guard against double counting. First, the standard logistics cost is already netted out in m, so that c_{e} carries only the increment over that cost and is charged on emergency units alone. Second, units served and units lost are disjoint by construction, so no unit contributes to both the first and the fourth term. Third, the margin forgone on lost units is not subtracted anywhere in the expression: it is simply not earned, because those units are excluded from Q_{served}. That margin does not appear in Table 15, but that table measures the cost of error against perfect foresight rather than restating this expression and the two are never added togheter.

The quantity is named expected economic contribution rather than net contribution, since it incorporates the relational cost, which is a loss of future client margin and not a cost of the current window. Where Chapter 1 speaks of net contribution per unit, it refers to the narrower quantity m above, which carries no relational component.

The commitment quantity follows from the target probability and the demand distribution, but the three policies read the pair differently. The optimized reactive policy applies the critical ratio to the pre-season distribution of window demand. The base-stock policy never identifies the window at all, and therefore applies the ordinary service level to the distribution of routine demand; its quantity is lower for that reason rather than because it targets a lower probability. Both commit a single number, fixed across replications. The anticipatory policy applies the same critical ratio as the optimized reactive one, but to

a distribution conditional on the observed signal. As a matter of fact, its quantity varies from replication to replication and is centered on the updated expectation rather than on the unconditional mean. This distinction matters for interpretation, since the anticipatory quantity should be reported as an average with its dispersion rather than as a single figure.

Once demand is realized, each replication is decomposed into four quantities, reported in Table 12. Their sum reconciles by construction: units served and units lost account for the whole demand, while units committed are exhausted between those sold from stock and those left over.

Table 12: ***Decomposition of a single replication***

QUANTITY	DEFINITION	COST CHARGED
SOLD FROM STOCK	The lesser of demand and the committed quantity	None
SERVED BY EMERGENCY TRANSFER	The recoverable share of any shortfall	Emergency premium
LOST	The residual shortfall	Margin forgone and relational cost
LEFTOVER	Committed quantity in excess of demand	Holding and markdown

Source: own elaboration

The five indicators of Section 1.3 are then computed from these quantities, as summarized in Table 13. Three of them translate without difficulty. The stock-out rate is the share of demand not met from local inventory. In addition, it is worth nothing that demand recovered by an emergency transfer still counts as a local stock-out, consistently with the definition adopted in Section 1.3: the request could not be satisfied from the inventory available at the location, whatever happened subsequently. The logistics cost per unit sold aggregates the standard cost on every unit served and the premium on those served urgently. Gross margin return on inventory investment relates the gross margin earned to the average inventory held over the window.

Table 13: ***Derivation of the five indicators from the simulated quantities***

INDICATOR	DERIVATION IN THE SIMULATION	COMPARABILITY
STOCK-OUT RATE	Unmet demand over total demand, including demand later recovered by emergency transfer	Direct
INVENTORY TURNOVER	Cost of units served over average inventory at cost, measured over the window	Window, not annual
GMROI	Gross margin earned over average inventory at cost	Direct
LOGISTICS COST PER UNIT SOLD	Standard cost on units served plus premium on units served urgently, over units served	Direct
NET MARGIN AFTER LOGISTICS	Gross margin less total logistics cost, over net revenue	Weakly discriminating

Source: own elaboration

Two of the indicators require a caveat. It is preferable to state both here rather than to allow the results of Section 3.3 to be misread. Inventory turnover is computed over the selling window and not over the year, whereas the benchmark range reported in Section 1.3 is annual. Converting between the two would require assuming that the window pattern repeats through the year, which the model does not simulate; the window figure and the annual benchmark should therefore be reported separately rather than compared directly. Net margin after logistics, as defined previously, is the gross margin percentage less logistics costs as a share of revenue. For a product carrying a gross margin of 70% this indicator settles close to 64% under every policy, because logistics costs move by fractions of a percentage point of revenue. It is reported for completeness, but it discriminates weakly between the scenarios and no conclusion rests on it.

Each configuration is evaluated over fifty thousand replications drawn from a fixed seed. The number is not arbitrary: the reported averages are stable to the second decimal place, so that differences between policies of the order of a few euro per window are not artefacts of sampling. Because the same seed generates the demand realizations for all three policies within a configuration, the comparison is paired: each policy faces exactly the same sequence of demands which removes sampling noise from the difference between them and makes small advantages measurable.

The calculation framework is finally exposed through an interactive interface, built on the same functions that produce every table in this chapter. It is not an auxiliary illustration but the calculation engine with controls attached: each input variable discussed above is exposed as a control, the outcome is recomputed on every change and any figure reported in this chapter can therefore be recovered by setting the controls to the values declared in Table 11. No formula is reimplemented, so the interface cannot diverge from the model it displays

Its purpose is to make the boundary conditions of the model inspectable rather than asserted. Ten of the parameters listed in Table 11 are assumptions, and a written sensitivity analysis, however systematic, still selects in advance which of them to vary and over what range. A reader who doubts a particular value can instead substitute their own and observe directly whether the conclusion survives. The interface also carries the diagnostic described above, reporting the probability mass falling below zero in the demand distribution and warning when the normal approximation ceases to be appropriate for the parameters selected. It is included as an appendix to this work rather than as a figure, since its content is the behavior of the model rather than any single state of it.

3.3 RESULTS COMPARISON AND DECISION MATRIX

Section 3.1 and 3.2 have established what the model compares and how each replication is priced. This section reports what the comparison produces. It proceeds in four steps. The first examines a single combination in detail, so that the mechanism generating the difference between the policies remains visible before it is aggregated. The second extends the comparison to the eight combination of the simulated context and confronts the outcome with the priority ranking advanced in Section 2.2. The third establishes which of the assumptions declared in Table 11 the conclusion actually depends upon, and identifies the condition under which anticipation ceases to pay for itself. The fourth summarizes the position in a SWOT analysis.

The detailed comparison is reported for the iconic carry-over article held at the flagship boutique during and identified peak window. This is the combination that Section 2.2 places in the highest priority tier and whose cost parameters are quoted in Section 3.1. It is therefore the case against which the framework should first be tested. Table 14 reports the five indicators of Section 1.3 for the three policies, together with the quantities committed and the composition of the resulting error.

Table 14: ***Policy comparison at the reference combination: iconic carry-over, flagship boutique, identified peak window***

INDICATOR	REACTIVE (BASE STOCK)	REACTIVE (OPTIMISED)	ANTICIPATORY
UNITS COMMITTED	43.0	54.0	47.9 (average)
EXPECTED ECONOMIC CONTRIBUTION (€)	6,494.36	6,819.00	7,113.64
STOCK-OUT RATE (%)	8.92	1.95	0.91
LEFTOVER UNITS	6.35	14.54	8.01
INVENTORY TURNOVER (WINDOW)	1.59	1.17	1.44
GMROI	3.70	2.72	3.35
LOGISTICS COST PER UNIT SOLD (€)	17.57	15.55	15.26
NET MARGIN AFTER LOGISTICS (%)	63.85	64.56	64.66

Source: own elaboration

Three readings follow from the table. The first concerns the baseline. The optimized reactive policy earns 324.64 euro per window more than the base-stock policy. It obtains that improvement without any forecasting capability whatever: it identifies the window and applies the critical ratio derived in Section 2.3 to information the brand already holds. A material share of the value available in this configuration is therefore accessible through better use of existing information rather than through anticipation. This is a managerially relevant result in its own right, and it establishes that the baseline against which anticipation is measured is not a straw man.

The second concerns the anticipatory policy. It adds a further 294.64 euro per window, an improvement of 4.3% over the stronger reactive policy, while committing 47.9 units on average against 54. It serves demand better with less inventory: the stock-out rate falls from 1.95% to 0.91% and the leftover quantity from 14.54 to 8.01 units. The two error types fall together rather than trading against one another, which is the central empirical finding of the accurate response literature [21] and the claim advanced in Section 2.3. The advantage does not come from stocking more: it comes from committing a quantity that moves with the demand it will face.

The third reading is a caution about the indicators themselves. The base-stock policy records the highest inventory turnover and the highest gross margin return on inventory investment of the three. While producing the lowest expected economic contribution and a stock-out rate nine times higher than the anticipatory policy. The result is not an artefact. Both indicators relate an outcome to average inventory held, so systematic under-commitment improves them and the demand that is never served is invisible to the denominator. Read in isolation they would endorse precisely the behaviour that Chapter 1 identifies as the costliest failure mode of the reactive model. This confirms the caution stated in Section 1.3: efficiency indicators built on inventory investment must be read jointly with the stock-out rate and never in place of it. Net margin after logistics, as anticipated in Section 3.2, settles between 63.85% and 64.66% under every policy and discriminates weakly. It is reported for completeness and no conclusion rests upon it.

Table 15 decomposes the cost of error into the four channels through which each policy loses value, so that the source of the difference can be located rather than inferred. Every figure in the table is an expected cost per selling window, on the same basis as the net contribution reported in Table 14. The relational channel is the one case in which a yearly quantity enters a per-window figure: the average client annual margin of Table 11 measures a year if future spending and what is charged here is the share of that year forgone on the units lost within this window. It is charged once, in the window in which the stock-out occurs and us not repeated in any later window.

Table 15: ***Decomposition of the cost of error per window at the reference combination (euro)***

COST CHANNEL	REACTIVE (BASE STOCK)	REACTIVE (OPTIMISED)	ANTICIPATORY
EMERGENCY PREMIUM	100.57	21.92	10.28
HOLDING AND MARKDOWN ON LEFTOVER	202.23	463.26	255.33
RELATIONAL COST OF LOST SALES	449.01	97.87	45.88
MARGIN FORGONE ON LOST SALES	199.34	43.45	20.37
TOTAL COST OF ERROR	**951.14**	**626.50**	**331.86**

Source: own elaboration

The decomposition shows that the two reactive policies fail in different ways and that anticipation does not simply do more of what the better of them does. The optimized reactive policy cuts the three shortfall channels sharply, from 748.92 to 163.24 euro in aggregate, but it more than doubles the cost of excess from 202.23 to 463.26 euro. Absent a demand signal, the only instrument available for reducing shortfall is to carry more inventory, so the policy relocates the error rather than removing it and retains only the difference between the two. The anticipatory policy reduces every channel simultaneously, including the cost of excess which falls below the level of both reactive policies. This is the quantitative expression of the mechanism described in Section 3.1: the committed quantity is both better centered on realized demand and less heavily padded against uncertainty. Furthermore, it should be noted that the relational cost is the largest single shortfall channel under every policy, which makes the comparison sensitive to the parameters governing it; that sensitivity is examined below and proves smaller than the decomposition would suggest.

The reference combination establishes the mechanism. It cannot establish whether the advantage is general and Section 2.2 argues explicitly that it should not be. Table 16 reports the eight combinations of the simulated context, in each case comparing the anticipatory policy with the better of the two reactive policies.

Table 16: ***Advantage of anticipation across the eight simulated combinations, per window***

COMBINATION	WINDOW	CRITICAL RATIO	BEST REACTIVE (€)	ANTICIPATORY (€)	GAIN (€)	GAIN (%)
ICONIC, FLAGSHIP	Identified peak	0.867	6,819.0	7,113.6	294.6	4.3
ICONIC, FLAGSHIP	Ordinary period	0.867	5,239.8	5,319.2	79.4	1.5
ICONIC, SECONDARY	Identified peak	0.817	1,998.5	2,071.8	73.3	3.7
ICONIC, SECONDARY	Ordinary period	0.817	1,533.6	1,553.6	20.0	1.3
SEASONAL, FLAGSHIP	Identified peak	0.750	1,327.2	2,549.6	1,222.4	92.1
SEASONAL, FLAGSHIP	Ordinary period	0.750	1,014.7	1,347.5	332.8	32.8
SEASONAL, SECONDARY	Identified peak	0.650	730.1	1,147.6	417.4	57.2
SEASONAL, SECONDARY	Ordinary period	0.650	556.6	672.6	115.9	20.8

Source: own elaboration

A note on method applies to this table and to Tables 14, 15 and 18. The figures are Monte Carlo estimates obtained from fifty thousand replications and are therefore subject to sampling error. The same quantities have been recomputed by numerical integration of the demand distribution, a procedure that involves no sampling at all and which underlies the interactive interface reported in the appendix. The two sets of figures agree to within one per cent on every combination and the ranking is identical; the largest divergence is 4.79 euro in absolute terms, on the seasonal article at the flagship boutique in an identified peak and 0.95 per cent in relative terms, on the same article at the secondary market in an ordinary window. The simulated estimates are reported here rather than the integrated ones because the simulation is the artefact from which the analysis is built, and the agreement between two procedures sharing no computational step is reported as evidence that neither carries the conclusion. One further difference should be recorded: the interface constrains the order-up-to level at zero, since a negative quantity has

no interpretation, whereas the simulation does not. The constraint binds in roughly half of one per cent of replications for the seasonal article and moves the measured advantage by at most 1.21 euro.

Figure 5: ***Advantage of anticipation across the four product-location pairs, by window type***

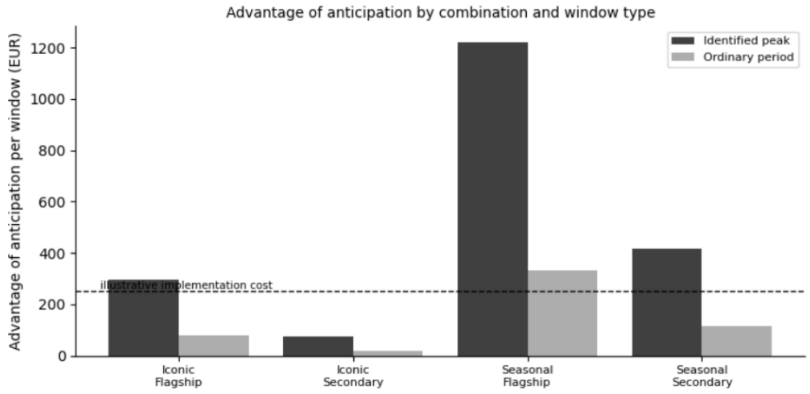

Source: own elaboration

The calendar condition advanced in Section 2.2 is confirmed without exception. In each of the four product-location pairs the identified peak yields a larger advantage than the ordinary period and the ratio between the two lies in a narrow band between 3.6 and 3.7. However, the effect operates principally through the correlation of the early signal, 0.85 in an identified window against 0.55 in an ordinary one, and only secondarily through the volume of the peak: varying each alone attributes roughly four fifths of the difference to the correlation and the remainder to the demand uplift, which acts as a pure change of scale. A period identified in advance is therefore valuable mainly because it is more forecastable rather than because it is busier.

The advantage is, also, selective, as Section 2.2 argues qualitatively. It ranges from 20.0 to 1,222.4 euro per window, a spread of nearly two orders of magnitude across combinations that differ only in the parameters of Table 10. Anticipating everywhere is not the correct choice and the question of where to anticipate is not secondary to the question of whether anticipation works. One result, on the other hand, separates two questions that Section 2.2 treats as one, and it is among the more important findings of this chapter. That section places the iconic carry-over article at the flagship boutique in the highest priority tier. Here that combination generates a relative gain of only 1.5% in its ordinary window and of 4.3% in its peak window, against 92.1% for the seasonal entry-level article at the same location. The finding was verified against alternative assumptions for the signal correlation, including reversing its value between the two product categories and the ranking does not change: it is driven by the difference in demand variability rathr than by the parameterization of the signal.

The two findings do not conflict, and setting them side by side clarifies both section. The two exercises rank the same combination by different criteria. Section 2.2 ranks them by suitability for pre-positioning, which asks where committing inventory ahead of demand is safe and worthwhile in absolute terms. This section ranks them by the incremental value of anticipating, which asks how much s added by observing an early signal over and above a reactive policy that already applies the critical ratio correctly. The reason the two diverges is structural. The characteristics that make an article safe to pre-position, namely stable demand and low exposure to markdown, are the same characteristics that leave a little uncertainty for a signal to resolve and make the residual error cheap when it occurs. The iconic article carries a coefficient of variation of 0.30 and a markdown depth of 0.10; the seasonal article carries 0.55 and 0.45, so that a proportionally identical error costs it far more and the signal has far more to resolve.

Both rankings are correct within their own terms and the figures demonstrate that they are not in competition. The iconic article at the flagship in a peak window return 7,113.64 euro of expected economic contribution against 2,549.59 euro for the seasonal article in the same window. A brand should still hold the iconic article at the flagship and Section 2.2 is not overturned. What the brand should not assume is that this is where a forecasting capability earns its cost. Section 2.2 answers where inventory should be placed. This section answers where it is worth investing in the capability to time that placement. In conclusion, the two answers are complementary and the practical implication is that the two decisions should be taken separately rather than derived from a single priority list. Put in the terms that matter to a decision-maker; stock-positioning priority is not the same thing as the value of predictive information, so the product a boutique most needs to have on hand is not necessarily the product on which it is most worth investing in predictive capability.

A symmetrical observation supports the same reading. At the seasonal article in the flagship peak window the optimized reactive policy earns almost nothing over the base-stock policy, 1.327.23 euro against 1,318.46 euro Where demand is highly variable, pre-season optimization has little to work with because the critical ratio can only reposition a distribution it cannot narrow. Only information helps. At the iconic article the position is reversed: reactive optimization captures the larger part of the value available and anticipation adds a modest remainder. The value of information is highest exactly where the alternatives to information are weakest

Section 3.1 declared that ten of the seventeen parameters in Table 11 are assumptions and undertook to demonstrate that the conclusion does not depend on them equally. Table 17 reports that demonstration. Each parameter is varied alone across its plausible range while every other is held at the base case and the swing it produces in the measured advantage indicates how far the conclusion rests upon it. The base case is the iconic article at the flagship because it is the highest-priority combination in Section 2.2 and because its advantage is positive but moderate, making it a conservative reference case for the sensitivity analysis.

Table 17: ***Sensitivity of the anticipatory advantage to each parameter, varied alone across its plausible range***

PARAMETER	RANGE TESTED	LOW END (€)	HIGH END (€)	SWING (€)
MARKDOWN DEPTH	0.05 – 0.25	187.1	522.1	335.0
EARLY-SIGNAL CORRELATION	0.60 – 0.95	121.4	426.2	304.9
COEFFICIENT OF VARIATION OF DEMAND	0.20 – 0.45	192.4	436.2	243.8
SHARE OF SHORTFALL MET BY EMERGENCY TRANSFER	0.40 – 0.90	338.9	229.9	109.0
AVERAGE CLIENT ANNUAL MARGIN	€250 – €2,000	264.0	370.0	106.0
RECOVERY FACTOR	0.30 – 0.90	338.7	250.1	88.6
DEMAND EROSION AFTER STOCK-OUT	0.10 – 0.35	259.3	315.1	55.7
SELLING WINDOW LENGTH	4 – 16 weeks	278.8	311.9	33.1
ANNUAL HOLDING RATE	0.15 – 0.40	281.2	303.2	22.0
EMERGENCY TRANSFER PREMIUM	€20 – €80	284.9	300.9	16.0
PEAK UPLIFT OVER ROUTINE DEMAND	0.10 – 0.50	290.6	290.6	0.0
ROUTINE SERVICE LEVEL	0.85 – 0.95	290.6	290.6	0.0

Source: own elaboration

Three parameters govern the result. The markdown depth, the correlation of the early signal and the coefficient of variation of demand produce swings of 335.0, 304.9 and 243.8 euro respectively. The remaining nine produce, individually, less than a third of the largest of these. Two of the three are not free assumptions. The signal correlation is anchored to the empirical evidence cited in Section 2.3 [21]. The markdown depth follows from the product characterization of Section 2.2, in which a carry-over article is by construction less exposed to clearance than one that must be cleared. Only the coefficient of variation is assumed outright. Its influence is monotonic and interpretable: the more variable the demand, the more anticipation is worth. Among the parameters varied one at a time, demand variability emerge sas the most influential parameter not directly anchored to an external source.

The most obvious objection to the model concerns the relational cost, whose three parameters are the least anchored in Table 11 and which Table 15 shows to be the largest shortfall channel. The objection does not survive the test. Varying the average client annual margin from 250 to 2,000 euro moves the result by 106 euro, the recovery factor across its full plausible range by 88.6 euro and the demand erosion by 55.7 euro, each of them less than the markdown depth alone. The reason deserves to be stated because it is not intuitive. The relational cost sets the level of the critical ration and therefore the quantity that every policy commits. It does not govern the difference between the policies which is a function of how far the early signal narrows and recenters the demand distribution. A parameter can dominate the cost decomposition and remain immaterial to the comparison and this one does.

Two parameters does not move the result at all. The peak uplift and the routine service level enter only the base-stock policy which is dominated by the optimized reactive policy at this configuration, so that neither reaches the reported advantage. This is a consequence of measuring the gain against the strongest of the two baselines rather than a defect of the model. It also means that this test, which holds expected demand fixed, does not measure the volume component of the calendar effect.

One divergence from Chapter 2 should be recorded. Table 8 of Section 2.3 identified forecast accuracy, markdown rate and urgent shipping cost as the three parameters determining whether pre-positioning is justified. The first two are confirmed here as the largest drivers of the result. The third is not: varying the emergency transfer premium from 20 to 80 euro moves the advantage by 16 euro. The second smallest swing recorded. The explanation is structural. The premium is charged only on the share of the shortfall the emergency channel actually recovers. Furthermore, the recovery share adopted here the volumes involved are small under every policy. Urgent shipping cost governs whether reactive correction is affordable which is the question Chapter 1 poses. It does not govern whether anticipation is worth its cost, which is the question posed here. The two are distinct and the divergence is reported rather than reconciled by adjustment.

Every result presented so far treats the early demand signal as free. Under that treatment information cannot harm, the anticipatory policy weakly dominates the reactive ones in every configuration and no break even exists. This is neither realistic nor consistent with Section 2.2, which argues that anticipatory logistics should be deployed selectively. Selectivity presupposes a cost, since in its absence the correct policy would be anticipate everywhere. Anticipation requires a forecasting capability, the personnel to operate it and the additional handling that pre-positioning entails.

Rather than assume value for that cost, the analysis inverts the question and asks, for each combination, the highest implementation cost at which anticipation still pays. The threshold is an output of the model rather than an input to it and it is directly interpretable by a decision-maker who knows what the capability costs their own organization. Table 18 reports it.

Table 18: ***Maximum sustainable implementation cost of the anticipatory capability, per combination per window, identified peak window***

COMBINATION	MAXIMUM SUSTAINABLE COST (€)
SEASONAL ENTRY-LEVEL AT THE FLAGSHIP BOUTIQUE	1,218.44
SEASONAL ENTRY-LEVEL AT THE SECONDARY MARKET	421.69
ICONIC CARRY-OVER AT THE FLAGSHIP BOUTIQUE	290.58
ICONIC CARRY-OVER AT THE SECONDARY MARKET	77.99

Source: own elaboration

The threshold are computed at the base configuration of the sensitivity analysis in which the expected demand is rounded and the replication count is reduced, so that they differ from the corresponding figures of Table 16 by a small margin. They are to be compared with one another rather than with Table 16. Their spread is the substantive result: the highest exceeds the lowest by a factor of approximately fifteen. A capability costing 300 euro per combination per window would be comfortably justified for the seasonal article at their location, marginal for the iconic article at the flagship and never justified for the iconic article at the secondary market. This is the selectivity of Section 2.2 expressed as a threshold

rather than as a judgment. It converts the qualitative eligibility test of that section into a figure a brand can compare with its own cost base.

Figure 6: ***Advantage of anticipation as a function of demand variability and signal strength, with the break-even contour***

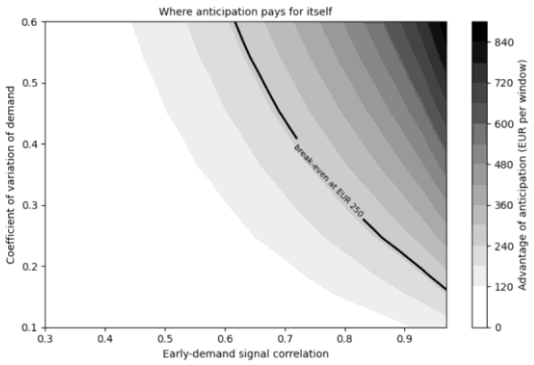

Source: own elaboration

Figure 6 traces the same result across the two parameters that jointly govern the value of the information, the variability of demand and the strength of the early signal. It marks the contour at which the advantage equals a given implementation cost. Below that contour the capability does not pay itself. The shape of the surface confirms the interpretation advanced above: the advantage rises in both parameters and it approaches zero not only where the signal is weak but also where demand is stable enough that there is little for the signal to resolve.

The simulation answers the economic question posed at the opening of this chapter. What it does not yet supply is a rule a manager can apply to a combination that was not simulated. Table 19 provides one. The sensitivity analysis of Table 17 identified three parameters that govern the result: the variability of demand, the quality of the early signal and the cost of being wrong. Those three are used here as the axes of a decision matrix, so that the recommendation for any combination follows from where it falls on them. Each row is anchored to one of the eight simulated combinations and to its measured advantage, which makes the matrix a summary of the results rather than a judgement laid over them.

Table 19: ***Decision matrix for the deployment of anticipatory logistics***

DEMAND UNCERTAINTY	SIGNAL QUALITY	COST OF ERROR	RECOMMENDATION	ILLUSTRATIVE COMBINATION	MEASURED GAIN
Low	Low	Low	Static policy	Iconic, secondary	€20.0 (1.3%)

				market, ordinary window	
Low	High	High	Evaluate case by case	Iconic, flagship, identified peak	€294.6 (4.3%)
High	Low	High	Improve forecasting before pre-positioning	Seasonal, flagship, ordinary window	€332.8 (32.8%)
High	High	High	Anticipatory policy strongly indicated	Seasonal, flagship, identified peak	€1,222.4 (92.1%)

Source: own elaboration

Two rows deserve comment. The second is the case a brand would address first on instinct and the one on which the matrix withholds a recommendation: the advantage of 294.6 euro is real, but it sits almost exactly on the sustainable cost of 290.58 euro reported in Table 18, so the decision turns on what the capability actually costs that organization. The third is the one the matrix redirects. There the cost of error is high and the advantage of 332.8 euro is already material, yet it is bounded by the quality of the signal rather than by the value at stake: raising the correlation from 0.55 to 0.85 multiplies the gain by 3.7. Pre-positioning more aggressively would not capture that; forecasting better would. The matrix therefore separates two failures that a single priority ranking would treat alike, one in which the opportunity is small and one in which it is large but not yet accessible.

The chapter opened by asking whether the framework proposed in Chapter 2 is worth following. The answer is affirmative but conditional, and the conditions are now stated in quantitative terms. Anticipatory logistics is economically convenient where demand is variable enough that information has something to resolve, where the early signal is strong enough to resolve it, and where the consequences of being wrong are asymmetric enough that the resolution is worth paying for. Where all three hold, as in the seasonal article at the flagship boutique during an identified peak, the advantage approaches a doubling of expected economic contribution. Where none holds, as in the iconic article at the secondary market in an ordinary period, the advantage does not cover a plausible implementation cost. The framework of Chapter 2 is therefore neither confirmed nor refuted as a general proposition; it is delimited, and the delimitation is the principal contribution of this chapter.

CONCLUSION

This thesis asked under what conditions anticipatory logistics generates an economic advantage large enough to cover what it costs and how a brand can identify in advance those conditions. The answer is affirmative but conditional and the conditions are narrower than the appeal of the idea suggests.

Anticipation pays where three things hold together: demand is variable enough that information has something to resolve, the early signal is accurate enough to resolve it and the consequences of committing wrongly are asymmetric enough that the resolution is worth paying for. Where all three hold, the advantage reaches 1,222 euro per selling window and more than doubles the contribution of the best reactive policy. Where none holds, it falls to 20 euro and would not cover a plausible implementation cost. That spread of a factor of sixty separates combinations differing only in their parameters which is the central practical finding: anticipating everywhere is not the correct policy and the question of where to anticipate is not secondary to the question of whether anticipation works.

The sensitivity analysis shows which parameters carry the result and the answer is not the intuitive one. Margin does not govern the advantage. Markdown exposure, signal quantity and demand variability do. Furthermore, two of the three are traceable to sources outside this work. The calendar effect behaves in the same counterintuitive way: an identified peak is valuable because it is more forecastable not because it is busier, The increase in volume contributes nothing once the improvement in predictability is accounted for. From this follows the distinction that the thesis contributes to the framework proposed in Chapter 2. Stock-positioning priority is not the same thing as the value of predictive information. The iconic article at the flagship boutique remains the combination that generates the largest absolute contribution and should be still held there. It is nonetheless among the combinations on which a forecasting capability earns at least because stable demand and low markdown exposure leave a signal little to resolve and make the residual error cheap. The product a boutique most needs on hand is therefore not necessarily the product on which predictive capability is worth buying. These are two decisions and they are better taken separately than derived from a single priority.

Two instruments follow from the analysis. The maximum sustainable cost of the capability, computed for each combination, converts the eligibility of Chapter 2 into a figure a brand can compare with its own cost base; across the cases examined it varies by a factor of fifteen. The decision matrix extends the recommendation to combinations the simulation did not cover using as its axes the three parameters the sensitivity analysis identified as decisive. The limits are those of the method. The dataset is synthetic, so the work demonstrates a mechanism and quantifies the conditions under which it operates rather than estimating any brand's returns; transfers between boutiques are excluded, and their inclusion would strengthen the reactive baseline. What the analysis does establish is that the conclusion does not rest on its distributional assumption, since repeating it under a distribution that cannot take negative values leaves the ranking unchanged and the advantage larger. Applying the framework to observed demand

data is the natural continuation, and the instruments developed here are the form in which that application would be carried out.

BIBLIOGRAPHY

[1] Kapferer, J.-N. & Bastien, V. (2012). The Luxury Strategy: Break the Rules of Marketing to Build Luxury Brands (2nd ed.), Kogan Page, London, 2012

[2] Bain & Company / Fondazione Altagamma (2024). Luxury in Transition: Securing Future Growth. 23rd Annual Luxury Goods Worldwide Market Study, 2024

[3] Wei, Y., & Li, F. (2020). Omnichannel supply chain operations for luxury products with conspicuous consumers. Transportation Research Part E: Logistics and Transportation Review, 137, 101918

[4] Christopher, M. (2016). Logistics & Supply Chain Management (5th ed.). Pearson Education

[5] Helm, R. & Hegenbart T., Explaining customer reactions to real stockouts. Review of Managerial Science, 7 (2013).

[6] Campo, K., Gijsbrechts, E. & Nisol, P., Towards understanding consumer response to stock-outs. Journal of Retailing, 76(2000), 219–242

[7] Campo, K., Gijsbrechts, E. & Nisol, P., The impact of retailer stockouts on whether, how much, and what to buy. International Journal of Research in Marketing, 20(2003), 273–286

[8] Anderson, E.T., Fitzsimons, G.J. & Simester, D., Measuring and Mitigating the Costs of Stockouts. Management Science, 52 (2006), 1751–1763

[9] Bowersox, D.J., Closs, D.J., Cooper, M.B. & Bowersox, J.C. (2020). Supply Chain Logistics Management (5th ed.). McGraw-Hill, New York, 2020

[10] Chopra, S. & Meindl, P., Supply Chain Management: Strategy, Planning and Operation (6th ed.). Pearson, Harlow, 2016

[11] Winkelhaus, S. & Grosse, E.H. (2020). Logistics 4.0: a systematic review towards a new logistics system. International Journal of Production Research

[12] Waller, M.A. & Fawcett, S.E. (2013). Data Science, Predictive Analytics, and Big Data: A Revolution That Will Transform Supply Chain Design and Management. Journal of Business Logistics, 34(2), 77–84

[13] Choi, T.M., Wallace, S.W. & Wang, Y. (2018). Big Data Analytics in Operations Management. Production and Operations Management, 27(10), 1868–1883

[14] Christopher, M. & Towill, D. (2001). An integrated model for the design of agile supply chains. International Journal of Physical Distribution & Logistics Management, 31(4), 235–246

[15] Arrow, K.J., Harris, T. & Marschak, J. (1951). Optimal Inventory Policy. Econometrica, 19(3), 250–272

[16] Sharma, R., Shishodia, A., Gunasekaran, A., & Munim, Z. H. (2022). The role of artificial intelligence in supply chain management: Mapping the territory. International Journal of Production Research, 60(24), 7527–7550

[17] Caniato, F., Caridi, M., Castelli, C. & Golini, R. (2011). Supply chain management in the luxury industry: A first classification of companies and their strategies. International Journal of Production Economics, 133(2), 622–633.

[18] Bindi, B., Bandinelli, R., Fani, V. & Pero, M.E.P. (2023). Supply chain strategy in the luxury fashion industry: impacts on performance indicators. International Journal of Productivity and Performance Management, 72(5), 1338–1367.

[19] Swaminathan, K. & Venkitasubramony, R. (2024). Demand forecasting for fashion products: A systematic review. International Journal of Forecasting, 40, 247–267.

[20] Al-Issa, N. & Thanasi, M. (2024). Mapping the Future of Tech-Infused Luxury: A Roadmap and Research Directions. International Journal of Consumer Studies, 48, e13103

[21] Fisher, M. & Raman, A. (1996). Reducing the Cost of Demand Uncertainty through Accurate Response to Early Sales. Operations Research, 44(1), 87–99

[22] Lücker, F., Timonina-Farkas, A. & Seifert, R.W. (2025). Balancing Resilience and Efficiency: A Literature Review on Overcoming Supply Chain Disruptions. Production and Operations Management, 34(6), 1495–1511

APPENDIX — SIMULATION CODE AND INTERACTIVE INTERFACE

The Python simulation described in Chapter 3 and the interactive interface built on it are available at the following address:

https://drive.google.com/drive/folders/1oSIw3NL5JeQKNXTCf0FMuit_TR0DczK-?usp=sharing

The folder contains two files. The first is the Jupyter notebook implementing the model: the cost decomposition, the critical ratio, the three stocking policies, the sensitivity analysis and the robustness check under a lognormal demand distribution. Every table reported in Chapter 3 is produced by this notebook and can be reproduced by running it from the first cell; it requires no installation and opens directly in Google Colab.

The second is a self-contained HTML page that exposes the same model through a graphical interface. The user enters their own products, boutiques and selling windows, either by typing them or by loading them from a spreadsheet, and the page generates the corresponding report: the advantage of anticipation for each combination, the channels through which each policy loses value, the target in-stock probability implied by the cost structure and the sensitivity of the result to each assumption. The report can be downloaded as an editable document. The page runs offline in any browser, requires no installation and transmits no data.

Expected values in the interface are obtained by integrating the demand distribution rather than by sampling it, which is the difference recorded in the note to Table 16; the economic model, the cost parameters and the decision rule are the same in both.